

Grounded Entity-Landmark Adaptive Pre-training for Vision-and-Language Navigation

Yibo Cui^{1,2}, Liang Xie^{*1,2}, Yakun Zhang^{1,2}, Meishan Zhang³, Ye Yan^{1,2}, Erwei Yin^{*1,2}

¹ Defense Innovation Institute, Chinese Academy of Military Science

² Tianjin Artificial Intelligence Innovation Center ³ Harbin Institute of Technology (Shenzhen)

<https://github.com/CSir1996/VLN-GELA>

Abstract

Cross-modal alignment is one key challenge for Vision-and-Language Navigation (VLN). Most existing studies concentrate on mapping the global instruction or single sub-instruction to the corresponding trajectory. However, another critical problem of achieving fine-grained alignment at the entity level is seldom considered. To address this problem, we propose a novel Grounded Entity-Landmark Adaptive (GELA) pre-training paradigm for VLN tasks. To achieve the adaptive pre-training paradigm, we first introduce grounded entity-landmark human annotations into the Room-to-Room (R2R) dataset, named GEL-R2R. Additionally, we adopt three grounded entity-landmark adaptive pre-training objectives: 1) entity phrase prediction, 2) landmark bounding box prediction, and 3) entity-landmark semantic alignment, which explicitly supervise the learning of fine-grained cross-modal alignment between entity phrases and environment landmarks. Finally, we validate our model on two downstream benchmarks: VLN with descriptive instructions (R2R) and dialogue instructions (CVDN). The comprehensive experiments show that our GELA model achieves state-of-the-art results on both tasks, demonstrating its effectiveness and generalizability.

1. Introduction

Vision-and-Language Navigation (VLN) [3] is an important task in the Embodied Vision community, which has gained great attention [66, 40, 50]. It aims to ask an agent to reach the target location inside photo-realistic environments by following natural language instructions. In VLN, cross-modal alignment is one critical step to accurately make the action decision for the agent [60], since matching the mentioned landmarks (objects or scenes) with visual observations can help the comprehensive understanding of the en-

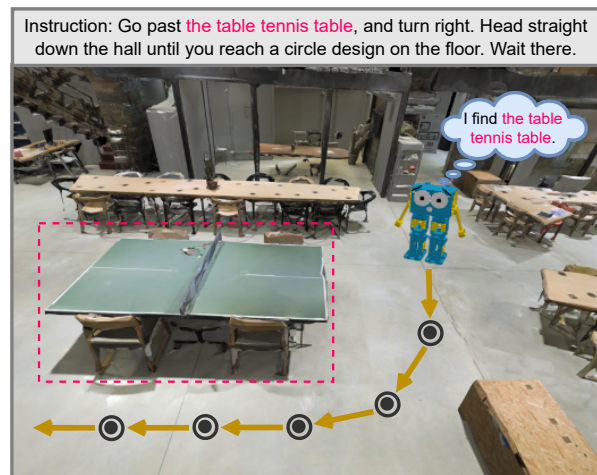


Figure 1. An illustration of an embodied agent navigating in a 3D photo-realistic environment. The agent is expected to navigate based on environment landmarks that correspond to the entity phrase in the instruction.

vironments and instructions [60, 15, 68]. However, most of the available datasets could only provide the coarse-grained text-image alignment signals [15], i.e., instruction-level correspondences to the complete trajectories, where finer ones are required to learn the cross-modality alignment for well-performed navigation [18, 67, 47, 43, 57].

A large part of previous works [66, 11, 40, 50] has concentrated on the global alignment of instructions and trajectories, which matches instructions to the overall temporal visual trajectory via reinforcement learning [60], auxiliary reasoning [66], or transformer-based matching pre-training [7, 50, 40]. Others [18, 67, 15] have attempted to align sub-instructions and sub-trajectories locally. At this granularity, the agents are designed to segment long instruction [67, 18] and determine which sub-instruction to focus on. To guide local cross-modal alignment, He *et al.* [15] introduced the Landmark-RxR dataset by human-annotating sub-instructions and sub-trajectories alignment. They ad-

*Corresponding author.

vanced the granularity of visual-textual alignment, but the agent learning is still supervised at the sentence level. The alignment within (sub-)instructions should be considered.

Identifying environment landmarks that correspond to entities in the instruction, as shown in Figure 1, is the next step to refine alignment and improve navigation. Recent studies [47, 43, 57] investigated the same problem, suggesting the alignment of entities and object regions [47] or the scene- and object-aware transformer model [43]. But the object regions were not under the direct supervision of the corresponding entity phrases in the instruction. To generate navigation instructions enriched with landmark phrases, Wang *et al.* [57] introduced grounded landmark annotations using a dependency parser and weak supervision from the pose traces, leading to the entity-landmark alignments of insufficient precision. Therefore, a dataset with high-quality entity-landmark level grounding annotations and powerful supervision for entity-landmark level cross-modal alignment is highly desired for VLN.

To address the above limitations, we first enhance the Room-to-Room (R2R) dataset [3] by introducing additional high-quality grounded entity-landmark human annotations, known as the Grounded Entity-Landmark R2R dataset (GEL-R2R). The GEL-R2R dataset is annotated with abundant precisely matched entity-landmark pairs (as magenta text and bounding box illustrated in Figure 1), which could provide the VLN models with direct supervision of fine-grained cross-modal alignment.

To verify the value of our GEL-R2R dataset, we propose a novel Grounded Entity-Landmark Adaptive (GELA) pre-training paradigm to improve the learning of entity-landmark level alignment for the VLN pre-trained model. Specifically, we suggest three grounded entity-landmark adaptive pre-training objectives: 1) Entity Phrase Prediction (EPP), locating entity phrases that refer to environment landmarks from the instruction; 2) Landmark Bounding box Prediction (LBP), predicting the bounding box of environment landmarks that match with entity phrases; and 3) Entity-Landmark Semantic Alignment (ELSA), aligning the matched pairs of landmark patches and entity tokens in the feature space by contrastive loss. These three tasks explicitly equip the model with the ability to comprehend the entity-level grounding between human instructions and visual environment observations.

Finally, we conduct extensive experiments on two downstream tasks to evaluate our proposed dataset GEL-R2R and adaptive pre-training methods GELA: Room-to-Room (R2R) and Vision-and-Dialog Navigation (CVDN). The instructions in R2R are fine-grained descriptions for the navigation trajectory, whereas the instructions in CVDN are multi-turn dialogs between the agent and the oracle during navigation. We use HAMT [7] as the backbone VLN model for the GELA pre-training. The results demonstrate our

suggested GELA achieves state-of-the-art (SoTA) performance in both seen and unseen environments of the above two benchmarks: 62% SPL on R2R and 5.87 GP on CVDN.

To summarize, our contributions are three-fold:

- We construct a new dataset GEL-R2R, which is the first dataset with high-quality grounded entity-landmark human annotations in the VLN domain.
- We propose a novel Grounded Entity-Landmark Adaptive (GELA) pre-training paradigm for VLN, explicitly supervising the models to learn fine-grained cross-modal semantic alignment between entity phrases and environment landmarks.
- Our suggested GELA achieves state-of-the-art results on two challenging VLN downstream benchmarks, demonstrating its effectiveness and generalizability.

2. Related Work

Vision-and-Language Navigation. Since the R2R benchmark was proposed by [3], studies of the VLN have made considerable progress in various aspects. VLN datasets have become increasingly diverse, existing in indoor [4, 30, 49, 55, 65] or outdoor [5, 16, 42] environments, and in discrete [3, 49] or continuous [29, 28, 34, 19] environments. The architecture of VLN models has become more complex, progressing from LSTM-based [11, 39] to transformer-based [14, 43, 45, 7]. To develop the competence of getting closer to the destination and to the ground truth path [23], VLN agents have been enhanced mainly through action strategy learning [63, 8, 48, 17, 21] and multi-modal representation learning [22, 59, 27, 26, 1]. Many excellent VLN models use imitation and reinforcement learning-based training paradigms for action strategy learning [60, 44, 54]. Recently, multi-modal transformer-based models that use effective joint representations for instructions and visual observations have achieved promising performance in VLN [12, 7, 50]. In this passage, we focus on incorporating more accurate and fine-grained cross-modal grounding information to improve the performance of VLN models by supervising the alignment between entity phrases and environment landmarks.

Phrase-to-Region Grounding. Phrase-to-region grounding is an important task in the Vision-Language (VL) domain that involves localizing textual entities in an image, commonly abbreviated as *phrase grounding* [25, 33, 38]. This task has seen significant progress since the introduction of the Flickr30k Entities dataset [46] and has played a crucial role in learning fine-grained semantic visual representation. Various models such as VisualBERT [32], MDETR [25], and GLIP [33] have explored this task under different architectures, and grounded pre-training has shown to facilitate fine-grained semantic understanding. However, the task of fine-grained cross-modal semantic alignment in embodied VLN tasks is more challenging due

	All						Object			Scene		
	Trajectory	Instruction	Phrase	P/I	Box	B/I	Phrase	Box	P/B	Phrase	Box	P/B
Train	4675	14039	57788	4.12	121146	8.63	30756	63248	2.06	27032	57902	2.14
Val Seen	340	1021	4196	4.11	8741	8.56	1939	4065	2.10	2257	4676	2.07
Val Unseen	775	2325	9483	4.08	20296	8.73	4324	8965	2.07	5159	11331	2.20
Total	5790	17385	71467	4.11	150183	8.64	37019	76278	2.06	34448	73909	2.15

Table 1. Statistics on our GEL-R2R dataset. P/I (resp. B/I) denotes the average number of entity phrases contained (resp. landmark bounding boxes matched with entity phrases) in each instruction. And P/B represents the average number of landmark bounding boxes matched with each entity phrase.

to more complex scenes and various objects. In this paper, we introduce high-quality entity-landmark grounding human annotations into the R2R dataset and propose a grounded entity-landmark adaptive pre-training scheme for VLN pre-trained models to address this problem.

Adaptive Pre-training. In the field of natural language processing, a pre-trained language model in the general domain is continuously pre-trained to learn the knowledge appropriate for a particular task or domain. This process is known as adaptive pre-training [13], which is conducted between pre-training and fine-tuning. Previous work [41, 64] gains consistent improvements by continuously pre-training an adaptive language model with the Masked Language Model (MLM) [10] objective. Furthermore, through three pre-training objectives — the MLM, Span Boundary Objective (SBO) [24], and Perturbation Masking Objective (PMO) — Wu *et al.* [62] improved the overall performance of a dialogue understanding model. In this work, we adopt adaptive pre-training into the area of VLN. Specifically, based on our human-annotated GEL-R2R dataset, we continuously train the SoTA pre-trained model [7] with three different supervised entity-landmark grounding objectives.

3. GEL-R2R Dataset

3.1. Dataset Construction

To establish alignment between entity phrases in instructions and their corresponding landmarks in the surrounding environment, we introduce GEL-R2R, a dataset that includes human annotations of grounded entity-landmark pairs. The construction process consists of five stages:

(1) Raw data preparation. We collect panoramas of the R2R dataset from the Matterport3D (MP3D) simulator [4] for each viewpoint. To improve the accuracy and efficiency of annotating landmarks, we mark the next action direction in panoramas and match each panorama with the corresponding sub-instruction based on [18].

(2) Annotation tool development. We develop a web-based annotation platform using the label-studio [56] to facilitate annotations. The annotation interface is shown in the supplementary material. Annotators first mark the entity

phrases in the instructions and then identify the corresponding landmarks in the images using the same labels.

(3) Annotation guideline standardization. Through pre-annotation, we standardized the annotation guideline by establishing four rules to ensure the consistency of the annotations. The four rules are as follows:

- **Alignment Rule:** The entity phrase in instructions should match landmarks in panoramas precisely.
- **Free Text Rule:** Free text instead of the class should be annotated, for instance, “the white dining table” instead of “table”.
- **Text Coreference Rule:** Entity phrases referring to the same landmark are marked with the same label.
- **Unique Landmark Rule:** For an entity phrase, only one corresponding landmark bounding box should be annotated in a panorama.

(4) Data annotation and revision. We first select annotators by testing the participants using 50 test instruction-path pairs. The qualification process leaves us 43 qualified annotators to complete the annotation task. Following the annotation, our experts double-check the annotations and correct any errors to ensure compliance with the above four rules.

(5) Data processing. In order to ensure the quality of the annotations, we first eliminate the incorrect annotations violating the above rules and then correct a few incorrect words. Finally, we established the GEL-R2R dataset by introducing additional grounded entity-landmark annotations to the R2R dataset. More details about data collection are provided in the supplementary material.

3.2. Statistics and Analysis

Table 1 presents statistics for the train, validation seen, and validation unseen splits of GEL-R2R. The dataset contains a total of 71,467 entity phrases, with 57,788 in the train split, 4,196 in the validation seen split, and 9,483 in the validation unseen split. It also includes 150,183 landmark boxes, of which 121,146 are in the train split, 8,741 are in the validation seen split, and 20,296 are in the validation unseen split. On average, each instruction contains 4.1 entity phrases, which is roughly equal to the number of actions in a path. Furthermore, each instruction uses about

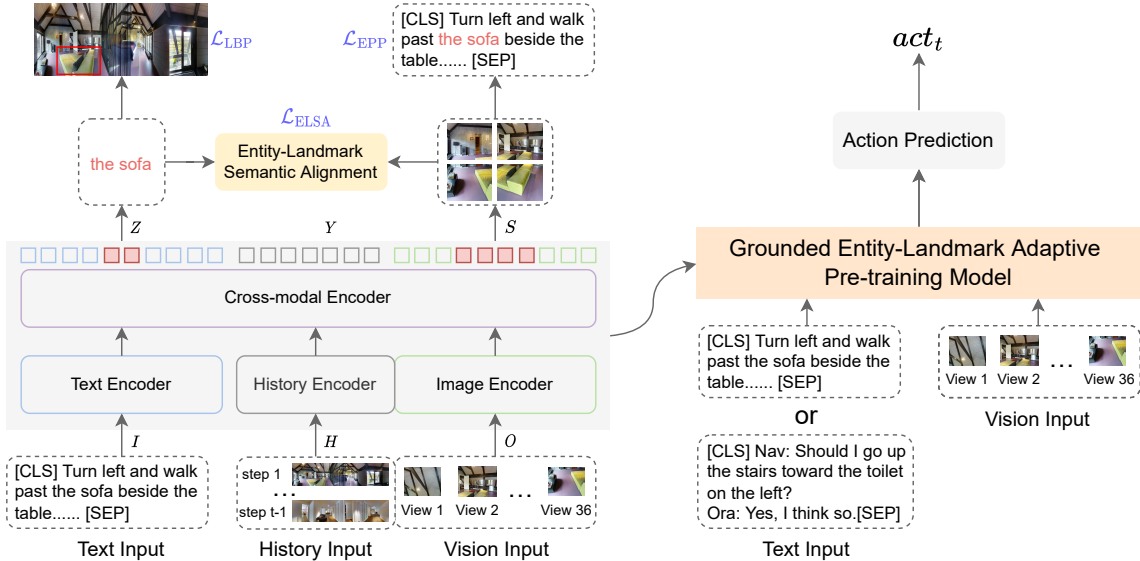


Figure 2. Overview of GELA. The three adaptive pre-training objectives served by our scheme encourage: 1) the annotated entity to predict the bounding box of its corresponding landmark in panoramas, 2) the annotated landmark to predict the positions of its corresponding entity phrases, and 3) the cross-modal encoder to establish the semantic alignment between landmarks and entities representations. After grounded entity-landmark adaptive pre-training, the GELA model is fine-tuned for the two VLN downstream tasks, R2R and CVDN.

8.5 landmark boxes on average. As a result, the agent has access to a diverse set of landmarks that can aid in decision-making during each navigation episode.

We compute statistics on the object and scene, taking into account their varying identification difficulties. In total, we gather 37019 object entity phrases and 34448 scene entity phrases, referring to 76278 object landmark boxes and 73909 scene landmark boxes, respectively. On average, there are about 2 landmark boxes corresponding to each exact entity phrase, whether for objects or scenes. This indicates that the same landmark appears in approximately two adjacent viewpoint panoramas from different sizes and angles, which enables the agent to perceive the landmark from multiple views. In summary, the GEL-R2R dataset is the first to provide high-quality entity-landmark alignment human annotations, which are essential for the area of VLN, such as cross-modal representation learning, data augmentation, and interpretable navigation.

4. Method

4.1. Problem Setup and Overview

Problem Setup. In VLN tasks, an embodied agent should move to a target location from a starting pose inside a 3D photo-realistic environment by following natural language navigable instructions. The instruction is a sequence of L words, denoted as $I = \{w_0, \dots, w_{L-1}\}$, and it guides the agent to traverse a connectivity graph G in order to reach the intended destination. At each step t , the agent ac-

quires a new panoramic visual observation $O_t = \{o_{t,i}\}_{i=1}^{36}$ from neighboring environment, where $o_{t,i} = [v_{t,i}, a_{t,i}]$ consists RGB image $v_{t,i}$ and orientation angle $a_{t,i}$ of i -th view. The relative angles of n navigable viewpoints to the current viewpoint and the stop action constitute the action space at step t , denoted as $A_t = \{act_{t,1}, \dots, act_{t,n}, [STOP]\}$. The agent selects an action to move to a navigable viewpoint in G or stop at the current location. Once the agent stops, the navigation episode is completed.

Overview. In this study, we propose a novel grounded entity-landmark adaptive pre-training scheme for a VLN pre-trained model, whose architecture is shown in Figure 2. The training scheme introduces explicit supervision on entity-landmark grounding to enhance fine-grained cross-modal representation learning. The scheme comprises two distinct stages: GELA pre-training on several proxy objectives and fine-tuning on the VLN downstream tasks.

4.2. Pre-trained Model

As illustrated in the bottom-left of Figure 2, our pre-trained model is a fully transformer-based architecture for multi-modal decision-making, which is modified from the classical cross-modal model LXMERT [53]. The pre-trained model takes three inputs: a global instruction I , history information H_t , and current panoramic visual observation O_t , which are fed in a language encoder, a history encoder, and a vision encoder, respectively. And then the textual and visual modalities exchange the signals through cross-attention layers in the cross-modal en-

coder. Specifically, the visual modality is the concatenation of history and visual observation. Finally, the representation of tokens in instruction, history, and visual state is $Z = \{z_{\text{cls}}, z_1, \dots, z_L\}$, $Y_t = \{y_{\text{cls}}, y_1, \dots, y_{t-1}\}$, and $S_t = \{s_1, \dots, s_{36}, s_{\text{stop}}\}$, respectively.

To learn effective uni-modal and multi-modal representation, transformer-based models for VLN commonly undergo pre-training on in-domain datasets using several proxy tasks [14, 50, 7]. These tasks include common vision-language pre-training tasks as well as VLN-specific auxiliary tasks. In this study, we utilize five proxy tasks: Masked Language Modeling (MLM) [10], Masked Region Classification (MRC) [37], Instruction Trajectory Matching (ITM) [40], Single-step Action Prediction (SAP) [7], and Spatial Relationship Prediction (SPREL) [7]. Details are demonstrated in the supplementary material.

4.3. GELA Pre-training

As shown in the top-left of Figure 2, we adopt three adaptive pre-training objectives for learning entity-landmark level alignment, which are Entity Phrase Prediction (EPP), Landmark Bounding box Prediction (LBP), and Entity-Landmark Semantic Alignment (ELSA). The three proxy objectives are designed by mimicking the objectives used in pre-training vision-language Transformers, particularly with the goal of learning phase-level grounding information. LBP is adapted from the training objective of TransVG [9]. EPP and ELSA are extended from the soft token prediction and contrastive alignment of MDETR [25], respectively.

Entity Phrase Prediction (EPP). In this objective, we predict the positions of entity phrases that correspond to annotated environment landmarks. We first transfer the human-annotated entity location to a mask vector M_z of $L+1$ dimensions, which is the same as Z . Similarly, we transfer the human-annotated landmark bounding box to a 37-dim (the same as S_t) mask vector M_s . Then we train the model to infer a uniform distribution over all token positions that refer to the corresponding landmark patches and supervise this process by the mask vector M_z . Specifically, we average the patches representation of the landmark: $S_t \times M_s^\top$, where \times denotes the matrix product, obtaining a 768-dim vector. Then we put it into a two-layer feedforward network (FFN) to predict a distribution over the token positions in the instruction sequence:

$$\text{logits} = \text{Softmax}(\text{FFN}(S_t \times M_s^\top)), \quad (1)$$

where the FFN maps a vector of 768-dim to a vector of $L+1$ dimensions. Finally, we minimize the cross-entropy loss between *logits* and M_z :

$$\mathcal{L}_{\text{EPP}} = \text{CrossEntropy}(\text{logits}, M_z). \quad (2)$$

Landmark Bounding Box Prediction (LBP). In this objective, we predict the bounding box of the landmark that matches with annotated entity phrases. We train the model to directly predict a 4-dim vector $\text{box}' = (x', y', w', h')$ as the coordinates of the bounding box for each entity phrases and supervise this process by the human-annotated bounding box $\text{box} = (x, y, w, h)$. Specifically, we first average the token embeddings of the entity phrase: $Z \times M_z^\top$ and then we predict the coordinates of the box through a two-layer FFN and sigmoid function:

$$\text{box}' = \text{Sigmoid}(\text{FFN}(Z \times M_z^\top)), \quad (3)$$

where the FFN maps a vector of 768-dim to a vector of 4-dim. Finally, we apply the smooth L1 loss and generalized IoU loss (GIoU loss) to optimize the box coordinates M_z :

$$\mathcal{L}_{\text{LBP}} = \mathcal{L}_{\text{smooth-l}_1}(\text{box}, \text{box}') + \lambda \mathcal{L}_{\text{GIoU}}(\text{box}, \text{box}'), \quad (4)$$

where λ is the weight coefficient of GIoU loss to balance these two losses.

Entity-Landmark Semantic Alignment (ELSA). While the above two unidirectional prediction tasks use positional information to match the entity and landmark, the entity-landmark semantic alignment loss enforces alignment between the hidden embeddings of the landmark and entity at the output of the cross-modal encoder. This additional contrastive alignment loss ensures that the representations of the landmark patches and the corresponding entity tokens are closer in the feature space compared to representations of unrelated tokens. This constraint is stronger than the above two unidirectional prediction losses as it directly operates on the representations and is not single-handedly based on the positional information. Specifically, inspired by InfoNCE loss [36], the objective is the mean of two contrastive losses as follows:

$$\mathcal{L}_s = \sum_{i=0}^{36} \frac{1}{|Z_i^+|} \sum_{j \in Z_i^+} -\log \left(\frac{\exp(s_i^\top z_j / \tau)}{\sum_{k=0}^L \exp(s_i^\top z_k / \tau)} \right), \quad (5)$$

$$\mathcal{L}_z = \sum_{i=0}^L \frac{1}{|S_i^+|} \sum_{j \in S_i^+} -\log \left(\frac{\exp(z_i^\top s_j / \tau)}{\sum_{k=0}^{36} \exp(z_i^\top s_k / \tau)} \right), \quad (6)$$

$$\mathcal{L}_{\text{ELSA}} = (\mathcal{L}_s + \mathcal{L}_z) / 2, \quad (7)$$

where Z_i^+ is the token set to be matched with an annotated patch s_i , and S_i^+ is the landmark patch set that should be matched with an annotated token z_i , and τ is a temperature parameter regulating attention to negative samples.

Therefore, the full adaptive pre-training objective is:

$$\mathcal{L}_{\text{GELA}} = \alpha \mathcal{L}_{\text{EPP}} + \beta \mathcal{L}_{\text{LBP}} + \gamma \mathcal{L}_{\text{ELSA}}. \quad (8)$$

4.4. Fine-tuning for VLN Tasks

The fine-tuning stage is illustrated in the right of Figure 2, where we generalize the GELA model for two VLN downstream tasks under the scheme of imitation learning (IL) and reinforcement learning (RL) by following previous work [20, 7]. IL supervises the agent to clone the behavior of the expert while RL encourages the agent to explore the trajectory according to the learning policy. Firstly, the GELA model navigates in the environment following the ground-truth action and generates gradients by IL. Secondly, using the same instruction, the model samples the action space to make decisions and uses RL to generate gradients. Finally, we combine the gradients and optimize the pre-trained model. For the CVDN task, our setting is the same to [7]. For the R2R task, we apply three settings with different data augmentation following [31]:

- `none`: without data augmentation on environments.
- `st`: with data augmentation by style transformation on original environments. Specifically, each discrete view of panoramas is transferred with a random style.
- `smo`: with data augmentation by semantic class masking on synthesis environments, which are generated by image synthesis with the same style as the original environments. One semantic class is randomly masked out during image generation.

5. Experiment and Results

5.1. Experimental Setup

Datasets. We evaluate our proposed method on the two VLN datasets: Room-to-Room (R2R) [3] and Vision-and-Dialog Navigation (CVDN) [55], which are all based upon the MP3D indoor environments. The instructions in R2R are detailed descriptions of the navigation trajectory. There are four splits in 21,558 trajectory-instruction pairs of the R2R dataset: training (14,039), validation seen (1,021), validation unseen (2,325), and test unseen (4,173). Our grounded entity-landmark human annotations are based on the training, validation seen, and validation unseen splits. The instructions in CVDN are multi-turn dialogs between the agent and the oracle during navigation. These kinds of instructions are frequently vague and unspecific. Therefore, the CVDN task is more challenging, serving to evaluate the generalization ability for new downstream tasks.

Evaluation Metrics. For R2R, we report four evaluation metrics: Trajectory Length (TL); Navigation Error (NE \downarrow) - the average distance in meters between the agent’s final position and the goal viewpoint; Success Rate (SR \uparrow) - the proportion of paths where the agent stopped within 3 meters of the goal viewpoint; Success rate weighted by Path Length (SPL \uparrow) [2]². SR and SPL are the recommended

²We denote a \downarrow to demonstrate lower is better and an \uparrow to demonstrate

Model	Feature	Validation Unseen			
		TL	NE \downarrow	SR \uparrow	SPL \uparrow
HAMT[7]	<code>none</code>	11.46	2.29	65.7	60.9
	<code>st</code>	11.78	3.42	67.3	62.6
	<code>smo</code>	12.13	3.22	67.9	62.9
GELA (ours)	<code>none</code>	11.75	3.33	69.2	63.4
	<code>st</code>	11.56	3.26	69.3	64.2
	<code>smo</code>	11.73	3.11	71.1	65.0

Table 2. Comparison with the HAMT baseline on the R2R validation unseen split. **Black** denotes the best results.

key metrics. For CVDN, we use the average progress in meters of the agent towards the goal viewpoint as the primary evaluation metric, denoted as Goal Progress (GP \uparrow).

Implementation Details. We typically conform to the architecture of [7] and its associated hyper-parameters. We conduct adaptive pre-training on our GEL-R2R training dataset and the augmented dataset from [14]. In the adaptive pre-training stage, we inherited the pre-training tasks to prevent knowledge forgetting. The GEL-R2R training set serves for both pre-training tasks and adaptive pre-training tasks. Due to the lack of entity-landmark grounding annotations, the augmented dataset serves only for pre-training tasks. We set $\lambda=1.0$ in Eq. (4), and $\alpha=1.0$, $\beta=1.0$, $\gamma=1.0$ in Eq. (8). And we set τ to 0.07 in Eq. (5) and Eq. (6) following [25]. And We adopt adaptive pre-training for 200k iterations with a learning rate of $5e-5$ by using 2 NVIDIA RTX 3090 GPUs, and the batch size for each GPU is set to 64. We adopt fine-tuning on the R2R and CVDN tasks respectively. For both R2R and CVDN tasks, we train the GELA model for 100k iterations with a learning rate of $1e-5$ on a single NVIDIA RTX 3090 GPU. We set the batch size to 16 for the R2R task and 8 for the CVDN task.

5.2. The Effect of GELA Pre-training

In this section, we demonstrate the effects of our GELA pre-training paradigm for VLN models in terms of navigation performance and cross-modal representation.

The effect of GELA pre-training on navigation performance. Table 2 presents a comparison with the HAMT baseline model on the R2R validation unseen split under three fine-tuning settings. The results of the first three lines are obtained from [31]. The results show that our GELA model outperforms HAMT by a large margin under all three settings. Specifically, GELA achieves absolute improvements of 3.5%, 2%, and 3.2% in SR under `none`, `st`, and `smo` respectively. Furthermore, we show the single-step action prediction (SAP) accuracy of GELA and HAMT on the R2R validation unseen split during adaptive pre-training in the left panel of Figure 3. As the training progresses, we

higher is better.

Methods	Validation Seen				Validation Unseen				Test Unseen			
	TL	NE↓	SR↑	SPL↑	TL	NE↓	SR↑	SPL↑	TL	NE↓	SR↑	SPL↑
Seq2Seq[3]	11.33	6.01	39	-	8.39	7.81	22	-	8.13	7.85	20	18
SSM [58]	14.70	3.10	71	62	20.70	4.32	62	45	20.40	4.57	61	46
EnvDrop[54]	11.00	3.99	62	59	10.70	5.22	52	48	11.66	5.23	51	47
AuxRN[66]	-	3.33	70	67	-	5.28	55	50	-	5.15	55	51
SEvol[6]	11.97	3.56	67	63	12.26	3.99	62	57	13.40	4.13	62	57
PREVALENT[14]	10.32	3.67	69	65	10.19	4.71	58	53	10.51	5.30	54	51
AirBERT[12]	11.09	2.68	75	70	11.78	4.01	62	56	12.41	4.13	62	57
RecBERT[20]	11.13	2.90	72	68	12.01	3.93	63	57	12.35	4.09	63	57
HOP[50]	11.26	2.72	75	70	12.27	3.80	64	57	12.68	3.83	64	59
REM[35]	10.88	2.48	75	72	12.44	3.89	64	58	13.11	3.87	65	59
HAMT[7]	11.15	2.51	76	72	11.46	2.29	66	61	12.27	3.93	65	60
GELA (ours)	11.19	2.39	76	73	11.73	3.11	71	65	12.99	3.59	67	62

Table 3. Comparison with SoTA methods on the R2R dataset. The methods in the top group are trained from scratch. The methods in the second group are based on pre-training.




Figure 3. The effect of GELA pre-training. Single-step Action Prediction (SAP) accuracy (left) and Effective Attention (EA) score (right) of GELA and HAMT on the R2R validation splits.

observe a gradual increase in the performance gap between GELA and HAMT, with GELA and HAMT achieving prediction accuracies of approximately 78% and 75%, respectively. These significant improvements demonstrate that fine-grained cross-modal alignment between entity phrases and environmental landmarks is highly effective in enhancing agents’ navigation performance.

The effect of GELA pre-training on cross-modal representation. To directly demonstrate the effect of GELA pre-training on cross-modal semantic alignment, we analyze cross-modal attention weights at the last Transformer layer during inference. First, we visualize the attention weights normalized to a scale of 0 to 1 in a heatmap. We then utilize the annotations in GEL-R2R to mark tokens, which are involved in the grounded entity-landmark pairs, and calculate the attention weights received by each token from the corresponding tokens in the other modality, defined as the Effective Attention (EA) score. In the right panel of Figure 3, we compare the average EA score of each marked token in

GELA and HAMT on the R2R validation seen and unseen splits, where *e2l* (resp. *l2e*) represents the attention of the entity token to the landmark token (resp. the attention of the landmark token to the entity token). Among all comparisons, GELA achieves significantly higher EA scores than HAMT (with a p-value in the two-sample t-test approaching 0), indicating that the grounded entity-landmark pairs have a stronger correspondence in GELA. While HAMT has a strong ability for uni-modal representation based on the large-scale transformer and end-to-end pre-training, it can be confused about matching the mentioned landmarks with visual observations. The explicit supervision for entity-landmark level alignment in GELA can improve the cross-modal representation and substantially enhance the agent’s decision-making ability during navigation. Qualitative examples are provided in the supplementary material.

5.3. Comparison to SoTA

Room-to-Room: R2R. Table 3 reports the performance comparison of various models on the R2R benchmark. The ensemble models, SE-Mixed [51] and EnvEdit [31], are excluded for a fair comparison. Our proposed GELA model outperforms all the other models in primary metrics (SR and SPL) across all the dataset splits. This indicates that GELA agents can navigate more accurately and efficiently than the other models. Specifically, on the validation seen split, GELA performs comparably to the SoTA HAMT. However, on the validation unseen and test unseen splits, GELA outperforms HAMT by absolute 5% and 2% (SR) and 4% and 2% (SPL) improvements, respectively. This demonstrates that GELA exhibits better generalization ability to unseen environments, which can be attributed to its stronger capability in capturing the semantic features of landmarks after entity-landmark grounding learning.

Methods	Val Seen	Val Unseen	Test Unseen
Seq2Seq[55]	5.92	2.10	2.35
CMN[68]	7.05	2.97	2.95
PREVALENT[14]	-	3.15	2.44
VISITRON[52]	5.11	3.25	3.11
ORIST[61]	-	3.55	3.15
HOP[50]	-	4.37	3.31
MT-RCM+EnvAg[61]	5.07	4.65	3.91
HAMT[7]	6.91	5.13	5.58
GELA (ours)	8.57	5.86	5.87

Table 4. Comparison with SoTA methods on the CVDN dataset.

Methods	Task	TL	NE↓	SR↑	SPL↑
0	None	11.46	2.29	65.7	60.9
1	EPP	11.83	3.63	65.8	60.7
2	LBP	12.71	3.71	65.6	60.1
3	ELSA	11.73	3.58	66.3	61.7
4	EPP+LBP	12.10	3.59	67.0	61.4
5	EPP+LBP+ELSA	11.75	3.33	69.2	63.4

Table 5. Ablation study of adaptive pre-training objectives on the R2R validation unseen split.

Vision-and-Dialog Navigation: CVDN. Table 4 presents the results on the CVDN benchmark, which utilizes Goal Progress (GP) in meters as the key performance metric. During training, we use a mixture of two types of demonstrations (navigator and oracle) as supervision for the trajectory. The results indicate that GELA consistently outperforms the other models on both the validation and test unseen environments. Particularly, our model achieves up to 1.66-meter improvement over the SoTA model HAMT on the validation seen split. Additionally, GELA outperforms the SoTA results by 0.73 meters and 0.29 meters on validation unseen and test splits, respectively. These results clearly demonstrate the GELA model is effective and generalizable to different types of instructions in more challenging VLN downstream tasks.

5.4. Ablation studies

Adaptive Pre-training Objective. To evaluate the effectiveness of different adaptive pre-training objectives, we conduct an ablation study on R2R validation unseen split under `none` setting. The results are presented in Table 5. Model 0 represents the HAMT baseline, while Models 1-3 show the results of combining the original proxy tasks in [7] with EPP, LBP, and ELSA, respectively. The results demonstrate that ELSA can further improve navigation performance, while the other two unidirectional prediction objectives solely based on positional information do not enhance the pre-trained model. This indicates that ELSA pro-

Category	TL	NE↓	SR↑	SPL↑
None	11.46	2.29	65.7	60.9
Scene	11.84	3.44	66.7	61.2
Object	11.23	3.54	66.8	61.9
All	11.75	3.33	69.2	63.4

Table 6. Ablation study of landmark categories on the R2R validation unseen split.

vides more effective supervision of cross-modal alignment by directly operating on the feature representations. Subsequently, we evaluate the effect of combining EPP and LBP. The results show that the combined performance (Model 4) is much better than the separate performance of Model 1 and Model 2, demonstrating that these two objectives are complementary. Finally, we combined all three objectives, and the results show that this performance (Model 5) improves even further. These findings highlight the importance of effective cross-modal alignment and the complementary nature of different adaptive pre-training objectives in improving navigation performance.

Landmark Category. Table 6 presents our investigation into the impact of two category entities (object and scene) from GEL-R2R under `none` setting. Our results demonstrate that integrating both object and scene entities can enhance navigation performance. Of the two entity types, object entities have a greater influence on performance, as agents find it easier to comprehend individual objects compared to complex scenes. Furthermore, the best performance is achieved by combining both entity types, enabling agents to navigate more precisely by utilizing both object and scene landmarks. This improvement can be attributed to the fact that the amount of available data for each entity type is only half that of the complete dataset. Therefore, it is plausible that the performance of GELA has not reached saturation due to the limited volume of available data.

5.5. Limitations and Future Work

GELA is a preliminary investigation into cross-modal alignment at the entity-landmark level in VLN based on our human-annotated GEL-R2R dataset. Due to the high cost of human annotation, the volume of GEL-R2R is relatively limited. In the future, we plan to expand the dataset using two cost-effective approaches: 1) developing a data augmentation model by using GEL-R2R to produce grounded entity-landmark annotations, and 2) incorporating large-scale *phrase grounding* datasets in the VL domain. Additionally, we intend to explore interpretable navigation based on our introduced dataset and method.

6. Conclusion

A well-performed VLN agent should have powerful competency in fine-grained cross-modal semantic alignment between entities and landmarks. In this paper, we introduce the manually annotated grounded entity-landmark dataset GEL-R2R, which provides powerful cross-modal alignment for VLN at the entity-landmark level. We then adopt three grounded entity-landmark adaptive pre-training objectives based on GEL-R2R to facilitate cross-modal semantic alignment learning under explicit supervision. Comprehensive experimental results on two downstream VLN tasks, R2R and CVDN, demonstrate the effectiveness and generalizability of our proposed model, GELA.

7. Acknowledgement

We thank ICCV reviewers for their constructive suggestions. This work is supported in part by the National Natural Science Foundation of China under Grant 62076250, Grant 61703407, and Grant 61901505.

References

- [1] Dong An, Yuankai Qi, Yan Huang, Qi Wu, Liang Wang, and Tieniu Tan. Neighbor-view enhanced model for vision and language navigation. In *MM '21: ACM Multimedia Conference*, pages 5101–5109, 2021. 2
- [2] Peter Anderson, Angel X. Chang, Devendra Singh Chaplot, Alexey Dosovitskiy, Saurabh Gupta, Vladlen Koltun, Jana Kosecka, Jitendra Malik, Roozbeh Mottaghi, Manolis Savva, and Amir R. Zamir. On evaluation of embodied navigation agents. *CoRR*, abs/1807.06757, 2018. 6
- [3] Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian D. Reid, Stephen Gould, and Anton van den Hengel. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *CVPR*, pages 3674–3683, 2018. 1, 2, 6, 7
- [4] Angel X. Chang, Angela Dai, Thomas A. Funkhouser, Maciej Halber, Matthias Nießner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from RGB-D data in indoor environments. In *3DV*, pages 667–676, 2017. 2, 3
- [5] Howard Chen, Alane Suhr, Dipendra Misra, Noah Snaveley, and Yoav Artzi. TOUCHDOWN: natural language navigation and spatial reasoning in visual street environments. In *CVPR*, pages 12538–12547, 2019. 2
- [6] Jinyu Chen, Chen Gao, Erli Meng, Qiong Zhang, and Si Liu. Reinforced structured state-evolution for vision-language navigation. In *CVPR*, pages 15429–15438, 2022. 7
- [7] Shizhe Chen, Pierre-Louis Guhur, Cordelia Schmid, and Ivan Laptev. History aware multimodal transformer for vision-and-language navigation. In *NeurIPS*, pages 5834–5847, 2021. 1, 2, 3, 5, 6, 7, 8
- [8] Shizhe Chen, Pierre-Louis Guhur, Makarand Tapaswi, Cordelia Schmid, and Ivan Laptev. Think global, act local: Dual-scale graph transformer for vision-and-language navigation. In *CVPR*, pages 16516–16526, 2022. 2
- [9] Jiajun Deng, Zhengyuan Yang, Tianlang Chen, Wengang Zhou, and Houqiang Li. Transvg: End-to-end visual grounding with transformers. In *ICCV*, pages 1749–1759, 2021. 5
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, pages 4171–4186, 2019. 3, 5
- [11] Daniel Fried, Ronghang Hu, Volkan Cirik, Anna Rohrbach, Jacob Andreas, Louis-Philippe Morency, Taylor Berg-Kirkpatrick, Kate Saenko, Dan Klein, and Trevor Darrell. Speaker-follower models for vision-and-language navigation. In *NeurIPS*, pages 3318–3329, 2018. 1, 2
- [12] Pierre-Louis Guhur, Makarand Tapaswi, Shizhe Chen, Ivan Laptev, and Cordelia Schmid. Airbert: In-domain pretraining for vision-and-language navigation. In *ICCV*, pages 1614–1623, 2021. 2, 7
- [13] Suchin Gururangan, Ana Marasovic, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. Don't stop pretraining: Adapt language models to domains and tasks. In *ACL*, pages 8342–8360, 2020. 3
- [14] Weituo Hao, Chunyuan Li, Xiujun Li, Lawrence Carin, and Jianfeng Gao. Towards learning a generic agent for vision-and-language navigation via pre-training. In *CVPR*, pages 13134–13143, 2020. 2, 5, 6, 7, 8
- [15] Keji He, Yan Huang, Qi Wu, Jianhua Yang, Dong An, Shuanglin Sima, and Liang Wang. Landmark-rxr: Solving vision-and-language navigation with fine-grained alignment supervision. In *NeurIPS*, pages 652–663, 2021. 1
- [16] Karl Moritz Hermann, Mateusz Malinowski, Piotr Mirowski, Andras Banki-Horvath, Keith Anderson, and Raia Hadsell. Learning to follow directions in street view. In *AAAI*, pages 11773–11781, 2020. 2
- [17] Yicong Hong, Cristian Rodriguez Opazo, Yuankai Qi, Qi Wu, and Stephen Gould. Language and visual entity relationship graph for agent navigation. In *NeurIPS*, 2020. 2
- [18] Yicong Hong, Cristian Rodriguez Opazo, Qi Wu, and Stephen Gould. Sub-instruction aware vision-and-language navigation. In *EMNLP*, pages 3360–3376, 2020. 1, 3
- [19] Yicong Hong, Zun Wang, Qi Wu, and Stephen Gould. Bridging the gap between learning in discrete and continuous environments for vision-and-language navigation. In *CVPR*, pages 15418–15428, 2022. 2
- [20] Yicong Hong, Qi Wu, Yuankai Qi, Cristian Rodriguez Opazo, and Stephen Gould. VLN BERT: A recurrent vision-and-language BERT for navigation. In *CVPR*, pages 1643–1653, 2021. 6, 7
- [21] Haoshuo Huang, Vihan Jain, Harsh Mehta, Alexander Ku, Gabriel Magalhães, Jason Baldrige, and Eugene Ie. Transferable representation learning in vision-and-language navigation. In *ICCV*, pages 7403–7412, 2019. 2
- [22] Gabriel Ilharco, Vihan Jain, Alexander Ku, Eugene Ie, and Jason Baldrige. General evaluation for instruction conditioned navigation using dynamic time warping. In *NeurIPS Workshop*, 2019. 2

- [23] Vihan Jain, Gabriel Magalhães, Alexander Ku, Ashish Vaswani, Eugene Ie, and Jason Baldridge. Stay on the path: Instruction fidelity in vision-and-language navigation. In *ACL*, pages 1862–1872, 2019. [2](#)
- [24] Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. Spanbert: Improving pre-training by representing and predicting spans. *Trans. Assoc. Comput. Linguistics*, 8:64–77, 2020. [3](#)
- [25] Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. MDETR - modulated detection for end-to-end multi-modal understanding. In *ICCV*, pages 1760–1770, 2021. [2](#), [5](#), [6](#)
- [26] Liyiming Ke, Xiujun Li, Yonatan Bisk, Ari Holtzman, Zhe Gan, Jingjing Liu, Jianfeng Gao, Yejin Choi, and Siddhartha S. Srinivasa. Tactical rewind: Self-correction via backtracking in vision-and-language navigation. In *CVPR*, pages 6741–6749, 2019. [2](#)
- [27] Jing Yu Koh, Honglak Lee, Yinfei Yang, Jason Baldridge, and Peter Anderson. Pathdreamer: A world model for indoor navigation. In *ICCV*, pages 14718–14728, 2021. [2](#)
- [28] Jacob Krantz and Stefan Lee. Sim-2-sim transfer for vision-and-language navigation in continuous environments. In *ECCV*, volume 13699 of *Lecture Notes in Computer Science*, pages 588–603, 2022. [2](#)
- [29] Jacob Krantz, Erik Wijmans, Arjun Majumdar, Dhruv Batra, and Stefan Lee. Beyond the nav-graph: Vision-and-language navigation in continuous environments. In *ECCV*, volume 12373 of *Lecture Notes in Computer Science*, pages 104–120, 2020. [2](#)
- [30] Alexander Ku, Peter Anderson, Roma Patel, Eugene Ie, and Jason Baldridge. Room-across-room: Multilingual vision-and-language navigation with dense spatiotemporal grounding. In *EMNLP*, pages 4392–4412, 2020. [2](#)
- [31] Jialu Li, Hao Tan, and Mohit Bansal. Envedit: Environment editing for vision-and-language navigation. In *CVPR*, pages 15386–15396, 2022. [6](#), [7](#)
- [32] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. *CoRR*, abs/1908.03557, 2019. [2](#)
- [33] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, Kai-Wei Chang, and Jianfeng Gao. Grounded language-image pre-training. In *CVPR*, pages 10955–10965, 2022. [2](#)
- [34] Xinghang Li, Di Guo, Huaping Liu, and Fuchun Sun. REVECE: remote embodied visual referring expression in continuous environment. *IEEE Robotics Autom. Lett.*, 7(2):1494–1501, 2022. [2](#)
- [35] Chong Liu, Fengda Zhu, Xiaojun Chang, Xiaodan Liang, Zongyuan Ge, and Yi-Dong Shen. Vision-language navigation with random environmental mixup. In *ICCV*, pages 1624–1634, 2021. [7](#)
- [36] Yue Liu, Junqi Ma, Yufei Xie, Xuefeng Yang, Xingzhen Tao, Lin Peng, and Wei Gao. Contrastive predictive coding with transformer for video representation learning. *Neurocomputing*, 482:154–162, 2022. [5](#)
- [37] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *NeurIPS*, pages 13–23, 2019. [5](#)
- [38] Panzhong Lu, Xin Zhang, Meishan Zhang, and Min Zhang. Extending phrase grounding with pronouns in visual dialogues. In *Proc. of the EMNLP*, pages 7614–7625, Abu Dhabi, United Arab Emirates, Dec. 2022. Association for Computational Linguistics. [2](#)
- [39] Chih-Yao Ma, Jiasen Lu, Zuxuan Wu, Ghassan AlRegib, Zsolt Kira, Richard Socher, and Caiming Xiong. Self-monitoring navigation agent via auxiliary progress estimation. In *ICLR*, 2019. [2](#)
- [40] Arjun Majumdar, Ayush Shrivastava, Stefan Lee, Peter Anderson, Devi Parikh, and Dhruv Batra. Improving vision-and-language navigation with image-text pairs from the web. In *ECCV*, volume 12351 of *Lecture Notes in Computer Science*, pages 259–274, 2020. [1](#), [5](#)
- [41] Shikib Mehri, Mihail Eric, and Dilek Hakkani-Tür. Dialogue: A natural language understanding benchmark for task-oriented dialogue. *CoRR*, abs/2009.13570, 2020. [3](#)
- [42] Dipendra Kumar Misra, Andrew Bennett, Valts Blukis, Eyvind Niklasson, Max Shatkhin, and Yoav Artzi. Mapping instructions to actions in 3d environments with visual goal prediction. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 2667–2678, 2018. [2](#)
- [43] Abhinav Moudgil, Arjun Majumdar, Harsh Agrawal, Stefan Lee, and Dhruv Batra. SOAT: A scene- and object-aware transformer for vision-and-language navigation. In *NeurIPS*, pages 7357–7367, 2021. [1](#), [2](#)
- [44] Khanh Nguyen, Debadeepta Dey, Chris Brockett, and Bill Dolan. Vision-based navigation with language-based assistance via imitation learning with indirect intervention. In *CVPR*, pages 12527–12537, 2019. [2](#)
- [45] Alexander Pashevich, Cordelia Schmid, and Chen Sun. Episodic transformer for vision-and-language navigation. In *ICCV*, pages 15922–15932, 2021. [2](#)
- [46] Bryan A. Plummer, Liwei Wang, Chris M. Cervantes, Juan C. Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. *Int. J. Comput. Vis.*, 123(1):74–93, 2017. [2](#)
- [47] Yuankai Qi, Zizheng Pan, Yicong Hong, Ming-Hsuan Yang, Anton van den Hengel, and Qi Wu. The road to know-where: An object-and-room informed sequential BERT for indoor vision-language navigation. In *ICCV*, pages 1635–1644, 2021. [1](#), [2](#)
- [48] Yuankai Qi, Zizheng Pan, Shengping Zhang, Anton van den Hengel, and Qi Wu. Object-and-action aware model for visual language navigation. In *ECCV*, volume 12355 of *Lecture Notes in Computer Science*, pages 303–317, 2020. [2](#)
- [49] Yuankai Qi, Qi Wu, Peter Anderson, Xin Wang, William Yang Wang, Chunhua Shen, and Anton van den Hengel. REVERIE: remote embodied visual referring expression in real indoor environments. In *CVPR*, pages 9979–9988, 2020. [2](#)

- [50] Yanyuan Qiao, Yuankai Qi, Yicong Hong, Zheng Yu, Peng Wang, and Qi Wu. HOP: history-and-order aware pre-training for vision-and-language navigation. *CoRR*, abs/2203.11591, 2022. 1, 2, 5, 7, 8
- [51] Wenda Qin, Teruhisa Misu, and Derry Wijaya. Explore the potential performance of vision-and-language navigation model: a snapshot ensemble method. *CoRR*, abs/2111.14267, 2021. 7
- [52] Ayush Shrivastava, Karthik Gopalakrishnan, Yang Liu, Robinson Piramuthu, Gökhan Tür, Devi Parikh, and Dilek Hakkani-Tur. VISITRON: visual semantics-aligned interactively trained object-navigator. In *Findings of ACL*, pages 1984–1994, 2022. 8
- [53] Hao Tan and Mohit Bansal. LXMERT: learning cross-modality encoder representations from transformers. In *EMNLP-IJCNLP*, pages 5099–5110, 2019. 4
- [54] Hao Tan, Licheng Yu, and Mohit Bansal. Learning to navigate unseen environments: Back translation with environmental dropout. In *NAACL-HLT*, pages 2610–2621, 2019. 2, 7
- [55] Jesse Thomason, Michael Murray, Maya Cakmak, and Luke Zettlemoyer. Vision-and-dialog navigation. In *CoRL*, volume 100 of *Proceedings of Machine Learning Research*, pages 394–406, 2019. 2, 6, 8
- [56] Maxim Tkachenko, Mikhail Malyuk, Andrey Holmanyuk, and Nikolai Liubimov. Label Studio: Data labeling software, 2020-2022. Open source software available from <https://github.com/heartexlabs/label-studio>. 3
- [57] Hanqing Wang, Wei Liang, Jianbing Shen, Luc Van Gool, and Wenguan Wang. Counterfactual cycle-consistent learning for instruction following and generation in vision-language navigation. In *CVPR*, pages 15450–15460, 2022. 1, 2
- [58] Hanqing Wang, Wenguan Wang, Wei Liang, Caiming Xiong, and Jianbing Shen. Structured scene memory for vision-language navigation. In *CVPR*, pages 8455–8464, 2021. 7
- [59] Hanqing Wang, Wenguan Wang, Tianmin Shu, Wei Liang, and Jianbing Shen. Active visual information gathering for vision-language navigation. In *ECCV*, volume 12367 of *Lecture Notes in Computer Science*, pages 307–322, 2020. 2
- [60] Xin Wang, Qiuyuan Huang, Asli Celikyilmaz, Jianfeng Gao, Dinghan Shen, Yuan-Fang Wang, William Yang Wang, and Lei Zhang. Reinforced cross-modal matching and self-supervised imitation learning for vision-language navigation. In *CVPR*, pages 6629–6638, 2019. 1, 2
- [61] Xin Eric Wang, Vihan Jain, Eugene Ie, William Yang Wang, Zornitsa Kozareva, and Sujith Ravi. Environment-agnostic multitask learning for natural language grounded navigation. In *ECCV*, volume 12369 of *Lecture Notes in Computer Science*, pages 413–430, 2020. 8
- [62] Han Wu, Kun Xu, Linfeng Song, Lifeng Jin, Haisong Zhang, and Linqi Song. Domain-adaptive pretraining methods for dialogue understanding. In *ACL/IJCNLP*, pages 665–669, 2021. 3
- [63] Yubo Zhang, Hao Tan, and Mohit Bansal. Diagnosing the environment bias in vision-and-language navigation. In *IJCAI*, pages 890–897, 2020. 2
- [64] Jeffrey Zhao, Mahdis Mahdieh, Ye Zhang, Yuan Cao, and Yonghui Wu. Effective sequence-to-sequence dialogue state tracking. In *EMNLP*, pages 7486–7493, 2021. 3
- [65] Fengda Zhu, Xiwen Liang, Yi Zhu, Qizhi Yu, Xiaojun Chang, and Xiaodan Liang. SOON: scenario oriented object navigation with graph-based exploration. In *CVPR*, pages 12689–12699, 2021. 2
- [66] Fengda Zhu, Yi Zhu, Xiaojun Chang, and Xiaodan Liang. Vision-language navigation with self-supervised auxiliary reasoning tasks. In *CVPR*, pages 10009–10019, 2020. 1, 7
- [67] Wang Zhu, Hexiang Hu, Jiacheng Chen, Zhiwei Deng, Vihan Jain, Eugene Ie, and Fei Sha. Babywalk: Going farther in vision-and-language navigation by taking baby steps. In *ACL*, pages 2539–2556, 2020. 1
- [68] Yi Zhu, Fengda Zhu, Zhaohuan Zhan, Bingqian Lin, Jianbin Jiao, Xiaojun Chang, and Xiaodan Liang. Vision-dialog navigation by exploring cross-modal memory. In *CVPR*, pages 10727–10736, 2020. 1, 8