This ICCV paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

The Road to Know-Where: An Object-and-Room Informed Sequential BERT for Indoor Vision-Language Navigation

Yuankai Qi¹ Zizheng Pan² Yicong Hong³ Ming-Hsuan Yang^{4,5,6} Anton van den Hengel¹ Qi Wu^{1*} ¹Australian Institute for Machine Learning, The University of Adelaide ²Monash University ³The Australian National University ⁴University of California, Merced ⁵Google Research ⁶Yonsei University {qykshr, zizhpan}@gmail.com yicong.hong@anu.edu.au mhyang@ucmerced.edu {anton.vandenhengel, qi.wu01}@adelaide.edu.au

Abstract

Vision-and-Language Navigation (VLN) requires an agent to find a path to a remote location on the basis of natural-language instructions and a set of photo-realistic panoramas. Most existing methods take the words in the instructions and the discrete views of each panorama as the minimal unit of encoding. However, this requires a model to match different nouns (e.g., TV, table) against the same input view feature. In this work, we propose an object-informed sequential BERT to encode visual perceptions and linguistic instructions at the same fine-grained level, namely objects and words. Our sequential BERT also enables the visual-textual clues to be interpreted in light of the temporal context, which is crucial to multiround VLN tasks. Additionally, we enable the model to identify the relative direction (e.g., left/right/front/back) of each navigable location and the room type (e.g., bedroom, kitchen) of its current and final navigation goal, as such information is widely mentioned in instructions implying the desired next and final locations. We thus enable the model to know-where the objects lie in the images, and to know-where they stand in the scene. Extensive experiments demonstrate the effectiveness compared against several state-of-the-art methods on three indoor VLN tasks: REVERIE, NDH, and R2R. Project repository: https://github.com/YuankaiQi/ORIST

1. Introduction

Vision-and-Language Navigation (VLN) offers the appealing prospect of more flexible interactions with robotic applications including domestic robots and personal assistants. One of the first VLN tasks to appear was Room-to-Room navigation (R2R) [4]. This task saw an agent initialised at a random location within a simulated environ-



Figure 1. Objects, rooms, and directions are important clues that can be inferred from visiolinguistic information. Our Object-and-Room Informed Sequential BERT (ORIST) is designed to learn to navigate by leveraging this information.

ment rendered from real images, and required it to navigate to a remote goal location according to natural-language instructions, such as "Leave the bedroom, and enter the kitchen. Walk forward, and take a left at the couch. Stop in front of the window." The actions available to the agent at each step are to investigate the current panorama, to move to a neighbouring navigable location/viewpoint, or to stop. An interactive version of the problem is introduced in [25, 33], while REVERIE [28] extends it to identifying remote objects, and TOUCHDOWN [8] introduces outdoor environments.

Numerous methods have been proposed to address indoor VLN tasks. Ma *et al.* [22] propose to learn textualvisual co-grounding to enhance the understanding of instructions completed in the past and the instruction to be executed next. Heuristic search algorithms for exploration and back-tracking are introduced in [16, 23]. Qi *et al.* [27]

^{*}Corresponding author

propose to disentangle object- and action-related instruction chunks for more accurate visual-textual matching. Another line of work exploits data augmentation to improve the generalisation ability in unseen environments. In [10], a speaker model is proposed to generate instructions for newly sampled trajectories, and in [32] visual information of seen scenarios is randomly dropped to mimic unseen scenarios. Most recently, it has been empirically demonstrated that hard example sampling and auxiliary losses [11, 41] are helpful for navigating unseen environments.

Although significant advances have been made, these methods take words in instructions and each discrete view of a panorama as the minimal unit for encoding, which limits the ability to learn the relationships between language elements and fine-grained visual entities. This is because the view feature, generally obtained by ResNet [13] which is trained for whole-image classification, mainly represents one salient object but it needs to be matched to many different natural-language landmarks mentioned in crowd-sourced navigation instructions. Figure 1 shows an example where the *coffee table, couch*, and *fireplace* are used as navigation landmarks in different instructions, but they need to match to the same input view feature in most existing VLN methods.

Motivated by the success of BERT-like models (Bidirectional Encoder Representations from Transformers) [9, 21, 30] in joint textual and visual understanding, we propose an object-and-room informed sequential BERT to encode instructions and visual perceptions at the same fine-grained level, namely words and object regions. This enables the model to better "know where" the objects referred to lie. We also introduce temporal context into our BERT-based model, which allows the model to be aware of completed parts of instructions and seen environments, resulting in more accurate next action prediction.

Relative directions (*i.e.*, left, right, front, and back) and room types (*e.g.*, living room, bedroom) are important clues for VLN tasks as they provide strong directional guidance and semantic-visual information for action selection. To take advantage of such information we incorporate a direction loss and two room-type losses into the proposed model. These predict the relative direction of navigable viewpoints, the type of room that need to be reached next, as well as the type of room at the final goal location. The relative direction is predictable from each viewpoint's orientation description (*i.e.*, heading and elevation angles), and the room types are identifiable through the presence of specific objects, such as a couch indicating living room or a microwave indicating kitchen. This provides the model with opportunity to "know where" it is, and where it is going.

To demonstrate the generalisation ability of our model, we evaluate on three different indoor VLN tasks: remote object grounding (REVERIE) [28], visual dialog navigation (NDH) [33] and room-to-room navigation (R2R) [4]. Our method achieves state-of-the-art results on all of the listed tasks: 18.97 SPL (9.28 RGSPL) on the REVERIE task, 3.17 GP on the NDH task, and 52% SPL on the R2R task.

2. Related Work

In this section, we briefly review closely related VLN methods and vision-and-language BERT-based works. Vision-and-Language Navigation. A large number of existing works have focused on cross-modality matching and generalisation from seen to unseen environments. For cross-modality matching, Ma et al. [22] propose a visualtextual co-grounding module and a progress monitor to distinguish completed instructions from those yet to be executed. In [38], a local navigation distance reward and a global instruction-trajectory matching reward are employed for reinforcement learning. Qi et al. [27] propose a multimodule framework to separately learn matching between action-related words and candidate orientations, and between landmark-related words and candidate visual features. In NvEM [1], visual features of navigable locations are adaptively enhanced by neighbor views leading to better visual-textual matching. To improve generalisation, data augmentation is introduced in [10] and [32] to generate new data from seen environments. Fu et al. propose to mine hard training samples via adversarial learning in [11]. In [37], reinforcement learning based on adaptively learned rewards is utilised to enhance generalisation. On the other hand, active exploration is designed in [36] to learn when, where, and what information to explore so as to form a robust agent.

In contrast to these methods, which take words and discrete images within each panorama as the minimal unit for encoding, our model is designed to handle fine-grained inputs, namely objects and words, to facilitate coordinated textual and visual understanding.

Vision-Language BERTs. VL-BERTs have been widely adopted to learn joint visual and textual models [30, 21, 31, 9, 18] for Vision-and-Language tasks, such as Visual Question Answering [5], Visual Commonsense Reasoning [40] and Referring Expressions [15]. In VLN, Majumdar et al. [24] predict whether a trajectory matches an instruction using a transformer. However, this method does not predict navigation actions. On the other hand, as pointed out in [12], most existing VL-BERTs learn to match textual and visual elements without considering their relationship to the actions available. To solve the problem, Hao et al. [12] propose a model that is pre-trained with image-text-action triplets generated from the R2R dataset. However, the pretrained model only acts as an instruction encoder and therefore represents a pre-processing step for other VLN methods. R-VLNBERT [14] enhances existing VLBERTs for VLN by representing agent states with the built-in [CLS] token, while it learns word-image level relationship.



Figure 2. An unrolled illustration of the proposed Object-and-Room Informed Sequential BERT (ORIST), which is composed of three main parts: the object-level initial embedding module (Section 3.2), the sequential BERT module (Section 3.3), and the room-and-direction multi-task module (Section 3.4). Note that all modules are reused between steps.

One of the key distinguishing features of the model proposed here is that, unlike most of existing VL-BERTs that are designed for one-round decision tasks (*e.g.*, VQA), our model is endowed with temporal context and thus is suitable for partially observed Markov decision processes.

3. Proposed Method

In VLN tasks, a robot agent is given a natural language instruction $\mathcal{X} = \{x_1, x_2, \cdots, x_L\}$, where *L* is the length of the instruction and x_i is a single word token. At each step, the agent is provided with a panoramic RGB image, which is divided into 36 discrete views $\mathcal{V} = \{v_{t,i}\}_{i=1}^{36}$ (12 horizontal by 3 vertical). A navigable location falls in one of these views and it has an orientation $\mathcal{D}_{t,i} = \langle \sin \theta_{t,i}, \cos \theta_{t,i}, \sin \phi_{t,i}, \cos \phi_{t,i} \rangle$, where θ and ϕ are measured relative to the current heading and elevation. At each step, the agent predicts the next navigable locations (including the current location). The agent stops if the current location is selected or it reaches the maximum number of steps.

3.1. Overview

Figure 2 shows an unrolled illustration of the proposed **O**bject-and-**R**oom Informed **S**equential BERT (ORIST). The ORIST is mainly composed of three parts: an object-level initial embedding module, a sequential BERT module and a room-and-direction multi-task module. All these modules are reused between steps. At step t, the initial embedding module takes as inputs an instruction \mathcal{X} , features of object regions \mathcal{O}_t , and navigable orientations \mathcal{D}_t . Then its outputs are fed into transformer layers of the sequential BERT module to adaptively fuse information from individual words and visual objects as well as each orientation via the self-attention mechanism. Finally, the fused features

 $\mathbf{H}_{t,U}$ are sent to two branches. One branch is the roomand-direction multi-task module, which predicts the room type $\mathbf{R}_{n,t}$ and $\mathbf{R}_{g,t}$, the direction \mathbf{D}_t , the next navigating action \mathbf{a}_t , and the navigation progress \mathbf{p}_t . Another branch is the LSTM of the sequtial BERT module, which produces a new temporal context \mathbf{h}_t and connects to the next navigation step. The whole model is trained end-to-end.

3.2. Object-level Initial Embedding Module

Existing VLN methods typically use ResNet [13] features of each discrete view to represent visual information. However, this may weaken vision-and-language matching because the ResNet is trained for image classification which may represent only one salient object, leading to the same image feature being matched to a variety of different objects mentioned in instructions (*e.g.*, laptop and table). To address this issue, we introduce an object-based representation to facilitate object-level cross-modality matching. The bounding boxes are adopted from REVERIE [28].

When designing the embedding module, we take into account the following three considerations. (I) The distinct characteristics of natural-language instructions and visual objects, *e.g.*, word tokens *vs.* object regions, word order *vs.* object position. (II) The type of an input token (*e.g.*, a visual or textual input). (III) The characteristic of VLN tasks, *e.g.*, the orientation feature of each navigable location. To this end, for an instruction \mathcal{X} and a candidate location in view $v_{t,i}$ composed of N^i objects $\mathcal{O}_{t,i} = \{o_{t,i}^1, \cdots, o_{t,i}^{N^i}\}$, we obtain the initial embedding via

$$\begin{aligned} \mathbf{E}_{t,w} &= \operatorname{Emb}(\boldsymbol{x}^{tok}) + \operatorname{Emb}(\boldsymbol{x}^{pos}) + \operatorname{Emb}(\boldsymbol{x}^{type}), \\ \mathbf{E}_{t,o}^{i} &= \operatorname{FC}(\boldsymbol{o}_{t,i}^{fea}) + \operatorname{FC}(\boldsymbol{o}_{t,i}^{pos}) + \operatorname{Emb}(\boldsymbol{o}_{t,i}^{type}), \\ \mathbf{E}_{t,d}^{i} &= \operatorname{FC}(\boldsymbol{d}_{t,i}), \end{aligned}$$
(1)

where $\mathbf{E}_{t,w} \in \mathbb{R}^{(L+2) \times d_h}$ denotes the embedding of the

expanded instruction $\boldsymbol{x}^{tok} = \{[\text{CLS}], \mathcal{X}, [\text{SEP}]\}$. In addition, $\boldsymbol{x}^{pos} = [0, 1, \cdots]$ is zero-based token position; $\mathbf{E}_{t,o}^i \in \mathbb{R}^{N^i \times d_h}$ denotes an embedding of N^i object regions; $\boldsymbol{o}_{t,i}^{fea} \in \mathbb{R}^{N^i \times 2048}$ represents objects' features extracted using FasterRCNN as described in [3]; and $\boldsymbol{o}_{t,i}^{pos} \in \mathbb{R}^{N^i \times 7}$ is the position feature of N object regions, of which each row is composed of the x, y location of the top-left and bottom-right points, and the region's height, width, and area. In Eq. (1), $\boldsymbol{x}^{type} = \mathbf{0}$ and $\boldsymbol{o}_{t,i}^{type} = \mathbf{1}$ are token type vectors, and $\mathbf{E}_{t,d}^i \in \mathbb{R}^{1 \times d_h}$ denotes the embedding of the candidate location's orientation $d_{t,i} \in \mathbb{R}^{1 \times 4} = \mathcal{D}_{t,i}$. Emb(\cdot) is a projection layer that embeds tokens to feature vectors. FC(\cdot) is a fully-connected layer.

For each candidate location *i*, we obtain its initial embedding matrix $\mathbf{H}_{t,0}^i = [\mathbf{E}_{t,w}, \mathbf{E}_{t,o}^i, \mathbf{E}_{t,d}^i] \in \mathbb{R}^{C_t \times d_h}$, where C_t is the maximum number of tokens at step *t* across all candidates. The subscript 0 in $\mathbf{H}_{t,0}^i$ denotes it is an initialisation before going through the sequential BERT module (detailed in the next section). The concatenation of all G_t candidates at step *t* constructs the initial embedding $\mathbf{H}_{t,0} = [\mathbf{H}_{t,0}^1; \cdots; ; \mathbf{H}_{t,0}^{G_t}] \in \mathbb{R}^{G_t \times C_t \times d_h}$.

3.3. Sequential BERT Module

BERT-like models have demonstrated their effectiveness in learning textual-visual entity matching [9, 30]. However, most of them are designed for one-round decision tasks (*e.g.*, VQA). VLN, as a multi-round task and partially observed Markov decision process, requires a model to be aware of which parts of an instruction have been completed and which parts have yet to be executed. As such, a direct application of BERTs to VLN does not perform well due to the loss of temporal context. To address these problems, in this work we propose a sequential BERT module.

The light blue panel in Figure 2 shows an unrolled illustration of our sequential BERT at step t and t + 1, which consists of two main components: transformer layers and an LSTM. Each transformer layer \mathcal{F}_j is a basic transformer [34], which is a multi-head self-attention

$$\mathbf{H}_{t,j} = \mathcal{F}_j(\mathbf{H}_{t,j-1}, \mathbf{M}_t), \qquad (2)$$

where $\mathbf{H}_{t,j-1} \in \mathbb{R}^{G_t \times C_t \times d_h}$ is the output from a previous transformer layer, \mathbf{M}_t is a mask matrix indicating whether a token can be attended, and $\mathbf{H}_{t,j} = [\mathbf{A}_{t,j,1}, \cdots, \mathbf{A}_{t,j,n}]$ is a concatenation of outputs of n attention heads. Concretely, the k-th head's output $\mathbf{A}_{t,j,k}$ is obtained by

$$\mathbf{A}_{t,j,k} = \operatorname{softmax}\left(\frac{\mathbf{Q}^{\top}\mathbf{K}}{\sqrt{d_{j,k}}} + \mathbf{M}_{t}\right)\mathbf{V}^{\top},$$
$$\mathbf{Q} = \mathbf{W}_{j}^{Q}\mathbf{H}_{t,j-1}^{\top}, \mathbf{K} = \mathbf{W}_{j}^{K}\mathbf{H}_{t,j-1}^{\top}, \mathbf{V} = \mathbf{W}_{j}^{V}\mathbf{H}_{t,j-1}^{\top}, \quad (3)$$

where \mathbf{W}_{j}^{*} is a learnable parameter matrix, and $d_{j,k}$ is the embedding dimension of $\mathbf{A}_{t,j,k}$. Assume we have U

transformer layers in total, as was demonstrated experimentally in [6, 9, 21], then the embedding of the first token (*i.e.*, [CLS]) of the last transformer layer, denoted as $\mathbf{H}_{t,U}[0] \in \mathbb{R}^{G_t \times d_h}$, is able to represent fused mutually attended information from textual and visual inputs.

To enable these transformer layers to model temporal context, an intuitive solution is to use a collection of object entities observed along the navigation path. In this way, one can expect that previously observed objects might be matched to completed instruction parts, and new visual perceptions could be matched to instruction parts to be executed. However, this will involve a huge computation cost (*e.g.*, GPU memory and computation time) as the agent navigates. To tackle this problem, we propose to employ an LSTM to adaptively learn how to summarise seen instructions and visual perceptions as the agent navigates. Specifically, we first encode $\mathbf{H}_{t,U}[0]$, which contains mutually attended instruction and visual information, to get a deeper representation $\mathbf{E}_t \in \mathbb{R}^{G_t \times d_h}$ via

$$\mathbf{E}_t = \operatorname{Tanh}(\mathbf{W}_e \mathbf{H}_{t,U}[0]). \tag{4}$$

Then, we aggregate surrounding information from all the G_t candidates via $\mathbf{I}_t = \text{AvgPool}(\mathbf{E}_t)$. Next, we obtain the new temporal context \mathbf{h}_t by

$$\mathbf{h}_t, \mathbf{c}_t = \mathrm{LSTM}(\mathbf{I}_t, \mathbf{h}_{t-1}, \mathbf{c}_{t-1}), \tag{5}$$

where \mathbf{h}_{t-1} and \mathbf{c}_{t-1} represents previous temporal context and LSTM cell state, respectively. \mathbf{h}_0 and \mathbf{c}_0 are initialised with zeros. Thus, \mathbf{h}_t contains the accumulated temporal context along the navigation and it is passed to the next navigation step as shown by the right panel of Figure 2.

3.4. Room-and-Direction Multi-task Module

Room types and directions can be useful information for navigation as they are often mentioned in instructions and provide strong navigation guidance. We note that a certain room type can usually be characterised by the presence of specific objects, such as bed for bedroom, TV/couch for living room, oven/microwave for kitchen. As our model takes objects as inputs, it is thus able to predict which room the agent should go to according to the current visual perception and instruction. Furthermore, it is also important for agents to understand the room that constitutes its final goal. This can help the agent make decisions based on the longterm goal. Additionally, we enable an agent to recognise the relative direction (i.e., left/right/front/back) of each candidate in order to better align with key direction cues (e.g., "turn left/right") in instructions. The progress monitor [41] is also adopted to further facilitate analysis of progress. We use MLP layers to implement the above goals:

$$\mathbf{D}_{t} = \mathrm{MLP}(\mathbf{E}_{t}),$$
$$\mathbf{R}_{n,t} = \mathrm{MLP}(\mathbf{I}_{t}), \ \mathbf{R}_{g,t} = \mathrm{MLP}(\mathbf{I}_{t}),$$
$$\mathbf{a}_{t} = \mathrm{MLP}(\mathbf{E}_{t}), \ \mathbf{p}_{t} = \mathrm{MLP}(\mathbf{I}_{t}),$$
(6)

where \mathbf{D}_t , $\mathbf{R}_{n,t}$, $\mathbf{R}_{g,t}$, \mathbf{a}_t , \mathbf{p}_t denote the relative direction prediction for each candidate (we evenly divide the surrounding 360° into four directions representing left/right/front/back), the room type prediction for the next step and the goal location, the action logit, and the navigation progress, respectively. The ground truth room types are extracted from the Matterport3D dataset [7]. \mathbf{E}_t is composed of embeddings of all candidate and \mathbf{I}_t is the fused embedding of all candidates as described in Section 3.3.

Loss Function. Similar to [32], we adopt the training strategy of mixed Imitation Learning (IL) and Reinforcement Learning (RL). Based on the above-mentioned tasks, for imitation learning we have the following losses:

$$\mathcal{L}^{IL} = \mathcal{L}^{\mathbf{D}} + \lambda_1 \mathcal{L}^{\mathbf{R}_n} + \lambda_2 \mathcal{L}^{\mathbf{R}_g} + \mathcal{L}^{\mathbf{a}} + \mathcal{L}^{\mathbf{p}}, \quad (7)$$

where $\mathcal{L}^{\mathbf{D}}$, $\mathcal{L}^{\mathbf{R}_n}$, $\mathcal{L}^{\mathbf{R}_g}$, $\mathcal{L}^{\mathbf{a}}$ are cross-entropy losses for the relative direction prediction \mathbf{D}_t , the room type of the next step $\mathbf{R}_{n,t}$, the room type of the goal location $\mathbf{R}_{g,t}$, and the navigating action \mathbf{a}_t , respectively. λ_1 and λ_2 are trade-off factors. $\mathcal{L}^{\mathbf{p}}$ is the progress loss adopted from [41], which is a BCELoss $\mathcal{L}^{\mathbf{p}} = \sum_t -p_t^* \log(\mathbf{p}_t)$ where p_t^* is the teacher progress at each time step t. As in [32], the A2C reinforcement learning method is adopted with the loss function

$$\mathcal{L}^{RL} = -\sum_{t} a_t^* \log(p_{a_t^*}) Z_t, \qquad (8)$$

where a_t^* is a sampled navigation action, $p_{a_t^*}$ is its probability and $Z_t = FC(\mathbf{I}_t)$ is an estimated reward. Finally, the total loss is

$$\mathcal{L} = \mathcal{L}^{RL} + \lambda_3 \mathcal{L}^{IL}, \qquad (9)$$

where λ_3 manages the trade-off between RL and IL.

4. Experiments and Analysis

In this section, we present extensive experimental evaluations on three VLN tasks: REVERIE [28], NDH [33], and R2R [4]. We first give a brief introduction to these tasks and the evaluation protocols, and then present two sets of experimental results: one is a comparison against several state-of-the-art VLN methods, and the other one is an ablation study on the proposed model.

4.1. Evaluation Tasks

REVERIE. This task [28] requires an agent to localise a remote target object within a photo-realistic indoor environment on the basis of concise human instructions, such as "Go to the massage room with the bamboo plant and cover the diamond shaped window". REVERIE contains two sub-tasks: (I) Vision-Language Navigation, where the agent needs to navigate to the target room; (II) Referring Expression Grounding, where the agent must identify the target object of an interaction from a provided set of objects (the interaction is not required). Here we compare primarily against the navigation sub-task, as our model has been developed for navigation rather than grounding.

NDH. The Navigation from Dialog History (NDH) task is derived from the Cooperative Vision-and-Dialogue Navigation (CVDN) dataset [33]. Specifically, each item of the CVDN dataset is a trajectory paired with a navigation dialogue between an Oracle and a Navigator, where the Oracle provides additional instruction towards navigating to the goal location when Navigator requests help. Each round of dialogue (together with previous dialogue if it exists) in CVDN is extracted, and forms an item within the NDH dataset. Based on which path is selected as the ground-truth, NDH has three settings: (I) the Oracle setting, which uses the shortest path observed by the Oracle; (II) the Navigator setting, which uses the path taken by Navigator; and (III) the Mixed setting, which takes the path of Navigator if it visits the target location, or the shortest path if not.

R2R. This task [4] requires an agent to follow natural language instructions to navigate from one room to a remote goal room in a photo-realistic indoor environment. Instructions in R2R contain rich linguistic information about the navigation trajectory, such as "Walk through the kitchen. Go past the sink and stove, stand in front of the dining table on the bench side."

4.2. Implementation details

We set parameters λ_1 and λ_2 to 0.2 to keep the associated losses at the same level with others, and set λ_3 to 0.2 following [32]. We use U = 12 transformers in our sequential BERT module. To take advantage of existing BERT-based works, we initialise the transformers in our model with the weights of UNITER [9]. UNITER is originally trained on four image-text datasets (COCO [19], Visual Genome [17], Conceptual Captions [29], and SBU Captions [26]) for joint visual and textual understanding. Note that these datasets have no overlap with the VLN data.

Our model is trained separately for each task. We use the AdamW [20] optimiser with a learning rate of 1×10^{-6} and the model is trained on 8 Nvidia V100 GPUs. For the R2R task, following common practice, augmentation data from Envdrop [32] are used. For the REVERIE and NDH tasks, only the original training data are used. It is worth noticing that although Thomason *et al.* [33] show that using the complete dialogue leads to much better results on the

	Val Seen						Val UnSeen						Test					
Methods	Navigation			PCSA	DCSDI +	Navigation			DCSA		Navigation				PCSA			
	SR↑	OSR↑	SPL↑	TL↓		KG51 L	SR↑	OSR↑	SPL↑	TL↓	KU3	NODI L	SR↑	OSR↑	SPL↑	TL↓	KUS I	KGSI L
RCM [38]	23.33	29.44	21.82	10.70	16.23	15.36	9.29	14.23	6.97	11.98	4.89	3.89	7.84	11.68	6.67	10.60	3.67	3.14
SM [22]	41.25	43.29	39.61	7.54	30.07	28.98	8.15	11.28	6.44	9.07	4.54	3.61	5.80	8.39	4.53	9.23	3.10	2.39
FAST-Short [16]	45.12	49.68	40.18	13.22	31.41	28.11	10.08	20.48	6.17	29.70	6.24	3.97	14.18	23.36	8.74	30.69	7.07	4.52
Nav-Pointer [28]	50.53	55.17	45.50	16.35	31.97	29.66	14.40	28.20	7.19	45.28	7.84	4.67	19.88	30.63	11.61	39.05	11.28	6.08
ORIST	45.19	49.12	42.21	10.73	29.87	27.77	16.84	25.02	15.14	10.90	8.52	7.58	22.19	29.20	18.97	11.38	10.68	9.28

Table 1. Results on the REVERIE dataset. MAttNet [39] is adopted for object grounding. The top two results are highlighted in red bold and blue italic fonts, respectively.

NDH task (rather than using just the oracle's answer), here we choose the oracle's answer as instruction to avoid the additional computation cost, and we achieve the best results documented, as is demonstrated later.

4.3. Evaluation Metrics

Metrics for REVERIE. Trajectory Length (TL), Success Rate (SR), Success weighted by Path Length (SPL) and Oracle Success Rate (OSR) are commonly used metrics for the navigation sub-task. Navigation is considered successful if the target object can be observed at the agent's final location, and considered as oracle-successful if the target object can be observed at one of its passed locations. TL is the average length of agent's navigation trajectories. SR is the percentage of successful tasks. SPL is the main metric for navigation [2], which represents a trade-off between SR and TL. Higher SPL indicates better navigation efficiency. Additionally, REVERIE uses Remote Grounding Success rate (RGS) and RGS weighted by Path Length (RGSPL) as whole-task performance metrics to measure the percentage of tasks that correctly locate the target object.

Metrics for NDH. Goal Progress (GP) is the main metric [33], which measures how much progress (in meters) the agent has made towards the target. A higher GP denotes a better performance.

Metrics for R2R. TL, SR, SPL and Navigation Error (NE) are four widely adopted metrics. A navigation is considered as successful if the agent stops within 3 meters to the target. NE is the average distance between the agent's final location and the target.

4.4. Comparison to State-of-the-Art Methods

Results on REVERIE. REVERIE is a newly proposed VLN task [28]. We evaluate our method against three baselines and a state-of-the-art model provided in [28]. Our method first performs navigation and then conducts object grounding as in [28] after the agent stops.

Table 1 presents the evaluation results. The proposed ORIST method achieves the best performance on both the Val UnSeen and Test splits. Particularly, our ORIST obtains nearly double the previous best performance in terms of the main SPL and RGSPL metrics on the Val UnSeen split. On the test split, it also achieves approximately 7% and 3% absolute improvement in terms of the navigation

Agent		Val Unseen		Test Unseen				
Agent	Oracle Navigator		Mixed	Oracle	Navigator	Mixed		
Random	1.09	1.09	1.09	0.83	0.83	0.83		
Skyline	8.36	7.99	9.58	8.06	8.48	9.76		
Seq2Seq [4]	1.23	1.98	2.10	1.25	2.11	2.35		
CMN [42]	2.68	2.28	2.97	2.69	2.26	2.95		
PREVALENT [12]	2.58	2.99	3.15	1.67	2.39	2.44		
ORIST	3.30	3.29	3.55	2.78	3.17	3.15		

Table 2. Results of the proposed method on the NDH datasets compared against state-of-the-art methods. Note that PREVALENT is pre-trained on a large scale augmented VLN dataset while ours not. The top two results are highlighted in red bold and blue italic fonts, respectively.

Agant		Val U	nseen		Test					
Agent	TL↓	NE↓	SR↑	SPL↑	TL↓	NE↓	SR↑	SPL↑		
FAST [16]	21.17	4.97	56	43	22.08	5.14	54	41		
SM [22]	17.09	-	43	29	18.04	5.67	48	35		
EnvDrop [32]	10.70	5.22	52	48	11.66	5.23	51	47		
AuxRN [41]	-	5.28	55	50	-	5.15	55	51		
OAAM [27]	9.95	-	54	50	10.40	5.30	53	50		
SERL [37]	-	4.74	56	48	12.13	5.63	53	49		
AVIG [36]	20.60	4.36	58	40	21.60	4.33	60	41		
PREVALENT[12]	10.19	4.71	58	53	10.51	5.30	54	51		
ORIST	10.90	4.72	57	51	11.31	5.10	57	52		

Table 3. Single run results on the R2R dataset compared against state-of-the-art methods. The top two results are highlighted in red bold and blue italic fonts, respectively.

metric SPL and the remote object grounding metric RGSPL, respectively. As analysed later in the ablation study, this mainly arises from the sequential design for the base BERT model and the direction loss. We also note that its navigation performance on the Val Seen split ranks second, which indicates our ORIST still has room to further learn from the training data.

Results on NDH. NDH is a recently proposed VLN task. We compare against CMN [42] and PREVELANT [12]. We also compare against Seq2Seq [4], the baseline proposed along with the NDH dataset.

Table 2 presents the navigation results using GP as the metric in three settings. The results show that our ORIST method achieves the best performance on both Val Unseen and Test Unseen splits under all settings. Specifically, ORIST achieves about 0.9 meters absolute improvement (about 40% relative improvement) under the "Navigator" setting on the Test Unseen split. As shown later in the ablation study, both our sequential design and new losses contribute to the final success.

Module				NDH Val Seen	NDH Val Unseen	1	REVERIE Val Seen			REVERIE Val Unseen				
	S	$\mathcal{L}^{\mathbf{D}}$	$\mathcal{L}^{\mathbf{R}_n}$	$\mathcal{L}^{\mathbf{R}_{g}}$	GP↑	GP↑	SR↑	OSR↑	SPL↑	TL↓	SR↑	OSR↑	SPL↑	TL↓
#1					4.00	2.42	41.9	47.1	39.5	10.83	9.50	23.9	7.7	12.43
#2	1				3.80	2.86	41.4	44.9	40.1	10.95	15.1	21.7	13.3	11.23
#3	1	1			3.88	3.09	41.5	44.3	39.5	10.64	18.6	24.0	16.9	10.41
#4	1	1	1		4.12	3.15	42.9	46.9	40.3	10.96	17.1	24.3	14.9	12.24
#5	1	1	1	1	4.21	3.30	45.2	49.1	42.2	10.73	16.8	25.0	15.1	10.90

Table 4. Ablation study of our model. Symbols S, $\mathcal{L}^{\mathbf{D}}$, $\mathcal{L}^{\mathbf{R}_n}$, $\mathcal{L}^{\mathbf{R}_g}$ denotes our sequential design, direction loss, room-type loss for the next step, and room-type loss for the goal location, respectively.

Results on R2R. We compare our model against eight state-of-the-art navigation methods, including FAST [16], SelfMonitor [22], EnvDrop [32], AuxRN [41], OAAM [27], SERL [37], AVIG [36], and PREVALENT [12]. PREVA-LENT [12] is pre-trained on a large scale augmented VLN dataset, and hence is not directly equivalent. VLN-BERT [24] is not compared because it needs to obtain a set of candidate paths via extra VLN methods. Table 3 presents the comparison of navigation results. Our ORIST method achieves favorable results on the Test split in terms of the main metric SPL, even taking PREVALENT into consideration.

4.5. Ablation Study

In this section, we evaluate the impact of the key components of our model: the sequential design for BERT and the newly proposed three loss functions. The evaluations are conducted on the NDH and REVERIE datasets. NDH utilises detailed navigation instructions as in the R2R task, while REVERIE utilises much shorter and high level instructions. The results are shown in Table 4, where the symbols S, $\mathcal{L}^{\mathbf{D}}$, $\mathcal{L}^{\mathbf{R}_n}$, and $\mathcal{L}^{\mathbf{R}_g}$ denote the sequential design for BERT, direction loss, room-type loss for the next step, and room-type loss for the goal location, respectively.

Effectiveness of the Sequential Design. To evaluate the effectiveness of the proposed sequential design, we compare against its variant which is a direct application of BERT (*i.e.*, without the connection from step t to step t+1). Both are initialised from UNITER. The results are presented in rows 1 and 2 of Table 4. They show that: (I) For Val Unseen splits, our sequential design achieves significant improvements: 0.44 meters absolute improvement (18% relative improvement) for NDH task and 6% absolute improvement (73% relative improvement) for REVERIE task. (II) For Val Seen splits, a slight performance regression is observed for the NDH task and a slight improvement appears on the REVERIE task. These results show that our sequential design effectively enhances the generalization ability and may reduce over-fitting in seen environments.

Effectiveness of Direction and Room-type Losses. By comparing the results in rows 2 and 3 of Table 4, we observe that the direction loss $\mathcal{L}^{\mathbf{D}}$ brings consistent perfor-

mance improvement on all four splits (in terms of SR for REVERIE), and particularly on the unseen splits. This might be attributed to that the direction information increases the distinguishability of navigation candidates and is easily generalized to unseen scenarios. Regarding the room-type losses (rows 3, 4, and 5), we find that for the NDH task both losses boost the performance (*e.g.*, from 3.88 to 4.21 on Val Seen, and from 3.09 to 3.3 on Val Unseen); for the REVERIE task, both losses improve performance on the Val Seen split (the SPL score increases from 39.5 to 42.2, SR from 41.5 to 45.2) while the next-room type loss $\mathcal{L}^{\mathbf{R}_n}$ harms on Val Unseen split. This might be caused by the fact that there is no next room mentioned in REVERIE. In contrast, its improvement on the Val Seen split could arise from over-fitting to some seen scenarios.

4.6. Qualitative Results

In Figure 3, we visualise one navigation trajectory on the R2R task to give an intuitive example of how our agent navigates. The visualisation is based on the attention visualisation tool BertViz [35]. The left panel shows the process of attention shifting as the agent navigates. Navigable locations at each step are marked using blue cylinders and the perceived views for each navigable location are bounded using an indexed rectangle. The green rectangle denotes the view associated with the next step predicted by our model. Objects in each view are marked with red bounding boxes. As shown in the left panel, the attention on the instruction shifts intuitively as the agent navigates. Concretely, at the first step, the agent focuses on "exit the room" and it goes forward to the door. Next, the agent becomes aware of "turn left" and performs the correct action. Then, at steps $3 \sim 5$, the attention is updated from "the left room" to "proceed through the room", and the agent behaves accordingly. At the last two steps, the attention focuses on the end of the instruction and the agent recognises the bathroom by virtue of the presence of the sink. Hence, it chooses the correct viewpoint and finally decides to stop.

To illustrate that the agent is on the road to "know where", we take the decision progress of step 6 as an example and show details in the right panel. Specifically, the right panel shows two kinds of attention distributions under each navigation candidate: 1) Attention on the instruc-



Figure 3. Visualisation of a navigation trajectory. The left panel depicts the whole trajectory, and the right panel shows attention distributions of a specific step (see Section 4.6 for details). Object tokens are indicated by "[]", and "[ORI]" denotes the orientation token.

tion. As shown by the pink color on the left column words (the darker the larger), the "bathroom" is focused on. 2) Self-attention on all tokens. In particular, the right column shows attention distribution of "bathroom" on all tokens. We can see "bathroom" is strongly connected to the object "[drawer#sink#table]" (a darker line between two tokens denotes a stronger connection). The 12 kinds of colors denote attention computed from 12 attention heads. Additionally, at the bottom of the right panel, we present the progress prediction, the next room and goal room type prediction, and the direction of each navigable viewpoint. The agent successfully infers these clues from the instruction and visual perceptions. All of these information assist the agent making the correct navigation decision.

5. Conclusion

We have proposed a novel unified sequential BERT model for general indoor VLN tasks. The sequential char-

acteristic enables the BERT to be better applied to several VLN tasks. By taking object-level and word-level inputs, our model is able to learn fine-grained relationships across textual and visual modalities for VLN. Moreover, we design a new direction loss and two room-type losses to facilitate the model to predict more accurate navigation actions. Extensive experimental evaluations on three VLN tasks demonstrate the effectiveness and generalisation ability of the proposed method.

As our sequential BERT is able to directly predict navigation actions, it enables us to combine the multi-task VLN pre-training and downstream task fine-tuning. We leave this as a future work.

6. Acknowledgements

This work is supported in part by the ARC DE190100539 and the NSF CAREER Grant #1149783.

References

- Dong An, Yuankai Qi, Yan Huang, Qi Wu, Liang Wang, and Tieniu Tan. Neighbor-view enhanced model for vision and language navigation. In ACM MM, 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 Roshan Zamir. On evaluation of embodied navigation agents. *CoRR*, abs/1807.06757, 2018. 6
- [3] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, pages 6077–6086, 2018. 4
- [4] 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, 5, 6
- [5] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. VQA: visual question answering. In *ICCV*, pages 2425–2433, 2015. 2
- [6] Jize Cao, Zhe Gan, Yu Cheng, Licheng Yu, Yen-Chun Chen, and Jingjing Liu. Behind the scene: Revealing the secrets of pre-trained vision-and-language models. In ECCV, 2020. 4
- [7] 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. 5
- [8] Howard Chen, Alane Suhr, Dipendra Misra, Noah Snavely, and Yoav Artzi. TOUCHDOWN: natural language navigation and spatial reasoning in visual street environments. In *CVPR*, pages 12538–12547, 2019. 1
- [9] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. UNITER: universal image-text representation learning. In ECCV, pages 104–120, 2020. 2, 4, 5
- [10] 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. 2
- [11] Tsu-Jui Fu, Xin Wang, Matthew Peterson, Scott T. Grafton, Miguel P. Eckstein, and William Yang Wang. Counterfactual vision-and-language navigation via adversarial path sampling. *CoRR*, abs/1911.07308, 2019. 2
- [12] Weituo Hao, Chunyuan Li, Xiujun Li, Lawrence Carin, and Jianfeng Gao. Towards learning a generic agent for visionand-language navigation via pre-training. In *CVPR*, pages 13134–13143, 2020. 2, 6, 7
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778. IEEE Computer Society, 2016. 2, 3

- [14] Yicong Hong, Qi Wu, Yuankai Qi, Cristian Rodriguez Opazo, and Stephen Gould. A recurrent vision-and-language BERT for navigation. In CVPR, 2021. 2
- [15] Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara L. Berg. Referitgame: Referring to objects in photographs of natural scenes. In *EMNLP*, pages 787–798, 2014.
- [16] 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. 1, 6, 7
- [17] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 123:32– 73, 2017. 5
- [18] Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, Yejin Choi, and Jianfeng Gao. Oscar: Object-semantics aligned pre-training for vision-language tasks. In ECCV, pages 121–137, 2020. 2
- [19] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In *ECCV*, pages 740–755, 2014. 5
- [20] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019. 5
- [21] 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. 2, 4
- [22] Chih-Yao Ma, Jiasen Lu, Zuxuan Wu, Ghassan AlRegib, Zsolt Kira, Richard Socher, and Caiming Xiong. Selfmonitoring navigation agent via auxiliary progress estimation. In *ICLR*, 2019. 1, 2, 6, 7
- [23] Chih-Yao Ma, Zuxuan Wu, Ghassan AlRegib, Caiming Xiong, and Zsolt Kira. The regretful agent: Heuristic-aided navigation through progress estimation. In *CVPR*, pages 6732–6740, 2019. 1
- [24] Arjun Majumdar, Ayush Shrivastava, Stefan Lee, Peter Anderson, Devi Parikh, and Dhruv Batra. Improving visionand-language navigation with image-text pairs from the web. In *ECCV*, pages 259–274, 2020. 2, 7
- [25] Khanh Nguyen and Hal Daumé III. Help, anna! visual navigation with natural multimodal assistance via retrospective curiosity-encouraging imitation learning. In *EMNLP*, pages 684–695, 2019. 1
- [26] Vicente Ordonez, Girish Kulkarni, and Tamara L. Berg. Im2text: Describing images using 1 million captioned photographs. In *NeurIPS*, pages 1143–1151, 2011. 5
- [27] Yuankai Qi, Zizheng Pan, Shengping Zhang, Anton van den Hengel, and Qi Wu. Object-and-action aware model for visual language navigation. In *ECCV*, pages 303–317, 2020. 1, 2, 6, 7
- [28] 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. 1, 2, 3, 5, 6

- [29] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In ACL, pages 2556–2565, 2018. 5
- [30] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. VL-BERT: pre-training of generic visual-linguistic representations. In *ICLR*, 2020. 2, 4
- [31] Hao Tan and Mohit Bansal. LXMERT: learning crossmodality encoder representations from transformers. In *EMNLP*, pages 5099–5110, 2019. 2
- [32] 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, 5, 6, 7
- [33] Jesse Thomason, Michael Murray, Maya Cakmak, and Luke Zettlemoyer. Vision-and-dialog navigation. In *CoRL*, pages 394–406, 2019. 1, 2, 5, 6
- [34] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, pages 5998–6008, 2017. 4
- [35] Jesse Vig. A multiscale visualization of attention in the transformer model. arXiv preprint arXiv:1906.05714, 2019. 7
- [36] Hanqing Wang, Wenguan Wang, Tianmin Shu, Wei Liang, and Jianbing Shen. Active visual information gathering for vision-language navigation. In *ECCV*, 2020. 2, 6, 7
- [37] Hu Wang, Qi Wu, and Chunhua Shen. Soft expert reward learning for vision-and-language navigation. In ECCV, pages 126–141, 2020. 2, 6, 7
- [38] Xin Wang, Qiuyuan Huang, Asli Çelikyilmaz, Jianfeng Gao, Dinghan Shen, Yuan-Fang Wang, William Yang Wang, and Lei Zhang. Reinforced cross-modal matching and selfsupervised imitation learning for vision-language navigation. In CVPR, pages 6629–6638, 2019. 2, 6
- [39] Licheng Yu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, Mohit Bansal, and Tamara L. Berg. Mattnet: Modular attention network for referring expression comprehension. In *CVPR*, pages 1307–1315, 2018. 6
- [40] Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. From recognition to cognition: Visual commonsense reasoning. In CVPR, pages 6720–6731, 2019. 2
- [41] Fengda Zhu, Yi Zhu, Xiaojun Chang, and Xiaodan Liang. Vision-language navigation with self-supervised auxiliary reasoning tasks. In *CVPR*, pages 10009–10019, 2020. 2, 4, 5, 6, 7
- [42] 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. 6