

Augmented Commonsense Knowledge for Remote Object Grounding

Bahram Mohammadi¹, Yicong Hong², Yuankai Qi³, Qi Wu¹, Shirui Pan⁴, Javen Qinfeng Shi¹

¹Australian Institute for Machine Learning (AIML), University of Adelaide

²Australian National University

³Macquarie University

⁴Griffith University

{bahram.mohammadi, qi.wu01, javen.shi}@adelaide.edu.au,

yicong.hong@anu.edu.au, yuankai.qi@mq.edu.au, s.pan@griffith.edu.au

Abstract

The vision-and-language navigation (VLN) task necessitates an agent to perceive the surroundings, follow natural language instructions, and act in photo-realistic unseen environments. Most of the existing methods employ the entire image or object features to represent navigable viewpoints. However, these representations are insufficient for proper action prediction, especially for the REVERIE task, which uses concise high-level instructions, such as "Bring me the blue cushion in the master bedroom". To address enhancing representation, we propose an augmented commonsense knowledge model (ACK) to leverage commonsense information as a spatio-temporal knowledge graph for improving agent navigation. Specifically, the proposed approach involves constructing a knowledge base by retrieving commonsense information from ConceptNet, followed by a refinement module to remove noisy and irrelevant knowledge. We further present ACK which consists of knowledge graph-aware cross-modal and concept aggregation modules to enhance visual representation and visual-textual data alignment by integrating visible objects, commonsense knowledge, and concept history, which includes object and knowledge temporal information. Moreover, we add a new pipeline for the commonsense-based decision-making process which leads to more accurate local action prediction. Experimental results demonstrate our proposed model noticeably outperforms the baseline and archives the state-of-the-art on the REVERIE benchmark. The source code is available at <https://github.com/Bahram-Mohammadi/ACK>.

Introduction

Navigating an embodied agent through complex and unseen environments by following natural language instructions is a challenging problem in artificial intelligence research. In this regard, vision-and-language navigation (VLN) (Anderson et al. 2018) has drawn the attention of many researchers in recent years (Fried et al. 2018; Tan, Yu, and Bansal 2019; Hao et al. 2020; Zhu et al. 2021; Chen et al. 2022; Li et al. 2023) and a variety of VLN tasks have been introduced in different levels, such as room-to-room (R2R) (Anderson et al. 2018) and remote embodied visual referring expression in real indoor environments (REVERIE) (Qi et al. 2020b). In R2R, the agent aims to reach a pre-specified location from a

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

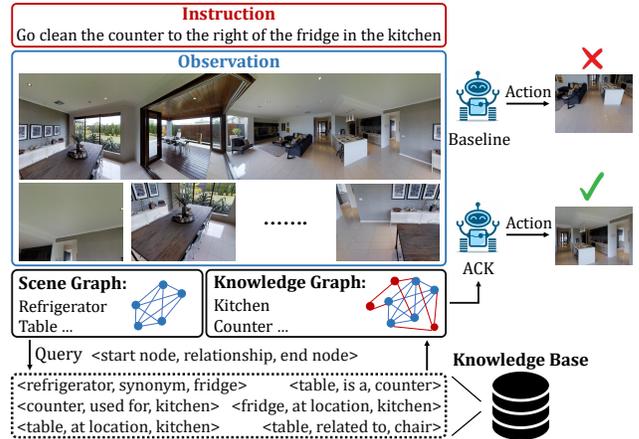


Figure 1: Action prediction by the baseline and our method. Utilizing visible objects alongside commonsense knowledge as a spatio-temporal knowledge graph improves visual representation and action prediction. Best viewed in color.

starting point by following fine-grained instructions. However, To be more practical, REVERIE introduces a goal-oriented task in which the agent needs to explore the environment and localize the target object according to concise instructions, e.g., "Go to the living room and clean the table next to the couch". Therefore, the agent cannot complete the task successfully just by strictly following instructions and requires more information about the environment to predict the correct action at each step.

Many of the proposed methods exploit scene-level features to represent visual perception (Qiao et al. 2022; Guhr et al. 2021; Chen et al. 2021). To provide the agent with richer visual clues, a wide range of previous works (Hong et al. 2020; Qi et al. 2020a, 2021) use object-level features. It is very common to use an object as a landmark in the instructions which means they can also be utilized as visual landmarks. Hence, object localization along with their orientation encoding regarding the heading and elevation angles of the agent is very beneficial to align the detected objects in the scene with the object labels in the instruction. For example, Given the instruction "With the fireplace on your right walk down the walkway and stop at the end before you enter

the next room”, the agent first needs to detect the fireplace and then correctly identify it as the visual landmark on the right to predict the proper action. High-level instructions and the lack of sufficient data in REVERIE disturb the navigation performance. This problem motivated us to provide the agent with additional information that is not present in the scene but can be inferred from the visible content of navigable directions. As Humans need to reason over complementary information to make a more precise decision, we aim to introduce commonsense knowledge into the REVERIE task to improve navigation. Employing external knowledge has shown significant performance in other vision-and-language problems such as image captioning and visual question answering (VQA) (Marino et al. 2019, 2021; Ding et al. 2022). Not only does such information enhance the vision and text alignment, but it also generalizes the action reasoning of the agent. Furthermore, the spatial and temporal connection between objects and knowledge during navigation boosts the exploration and generalization ability of the agent in unseen environments. For instance, *hallway* is followed by *living room* in the training environment, thus learning this sequence helps the agent to take the right action by observing the same concept, *hallway*, in an unseen environment.

In this work, to fulfill the above-mentioned requirements, we incorporate commonsense knowledge into the REVERIE task as a spatio-temporal knowledge graph by constructing a knowledge base followed by our proposed method ACK. We build an internal knowledge base according to the visible contents of the image and ConceptNet (Liu and Singh 2004) as the external knowledge base. Then, the pre-trained contrastive language-image pre-training (CLIP) model (Radford et al. 2021) is utilized to collect and rank the most pertinent knowledge to the scene and detected objects. These object and knowledge features are complementary to the existing visual representation and align the viewpoint images with instructions. Afterward, we present ACK which takes advantage of commonsense knowledge to enhance visual representation and action reasoning. ACK consists of two modules to integrate object and knowledge features along with their historical information. The knowledge graph-aware cross-modal encoder models the relationship between concepts and instructions while the concept history of the previous step is taken into account. Subsequently, the concept aggregation module outputs a single concept feature per each navigable direction which is utilized for visual representation enhancement and local action prediction in the commonsense-based decision-making pipeline. Figure 1 demonstrates the action prediction comparison between DUET (Chen et al. 2022) as the baseline and ACK. As shown in this figure, our model can make the right decision by exploiting object-level and knowledge-level features.

The experiments are conducted on the REVERIE dataset and results show that our proposed approach, ACK, outperforms the state-of-the-art methods. In summary, the main contributions of this work are as follows:

- We integrate object-level features with object-related commonsense knowledge to complement the existing visual representation and enhance vision-text alignment.
- We propose ACK followed by the commonsense-based decision-making pipeline for leveraging object and knowledge features to provide a more informative representation and hold their historical information for making a sequence of correct decisions.
- We conduct experiments on REVERIE to validate the effectiveness and generalization ability of ACK and the results show the superiority of our approach over the state-of-the-art methods and the baseline model.

Related Work

Vision-and-Language Navigation. In the VLN task, an agent is required to find the optimal path toward the target location given the visual and textual input data. In recent years, numerous approaches (Wang et al. 2021; Deng, Narasimhan, and Russakovsky 2020; Lin et al. 2022; Liu et al. 2021) have been proposed to improve the performance of the agent. The first baseline for the VLN is introduced by (Anderson et al. 2018) which designs a multi-modal Seq2Seq model. Then (Fried et al. 2018) proposes the speaker-follower method to augment the data and improve the generalization. EnvDrop (Tan, Yu, and Bansal 2019) extends this work by presenting environmental dropout. (Wang et al. 2019) presents a module for cross-modal grounding which enables the agent to infer the essential parts of the scenes and sub-instructions. (Ma et al. 2019) focuses on progress monitoring regarding the instruction. In light of the successes of vision-and-language pre-training PREVALENT (Hao et al. 2020) pre-trains the navigation model using the self-learning approach and AirBERT (Guhur et al. 2021) improves cross-modality interaction. Many of the recent works in VLN rely on the navigation history to achieve superior performance. Some approaches condense the history into a single vector (Wang et al. 2019; Hong et al. 2021) while others attempt to explicitly store the previous states (Pashevich, Schmid, and Sun 2021; Chen et al. 2022). Most recently, (Zhao, Qi, and Wu 2023) propose to especially mind the passed target location in trajectory histories.

VLN with Commonsense Knowledge. In the context of VLN, incorporating commonsense knowledge is rarely considered in previous works, however, it has recently drawn the attention of researchers in this topic (Gao et al. 2021; Li et al. 2022, 2023). The most widely-used large-scale structured knowledge bases are ConceptNet (Liu and Singh 2004) and DBpedia (Auer et al. 2007) which are created by automatic data extraction and manual annotation, respectively. CKR (Gao et al. 2021) exploits ConceptNet to iteratively perform object- and room-entity reasoning through internal and external graph reasoning during the training. In another line of work, KERM (Li et al. 2023) proposes a knowledge-enhanced reasoning model to take advantage of external knowledge which is retrieved from the visual genome (Krishna et al. 2017) by parsing the region descriptions. In this work, we leverage commonsense knowledge to properly align viewpoints with instructions and improve visual representation as well as local action prediction. The architecture of DUET (Chen et al. 2022) is followed as the main baseline of our proposed method.

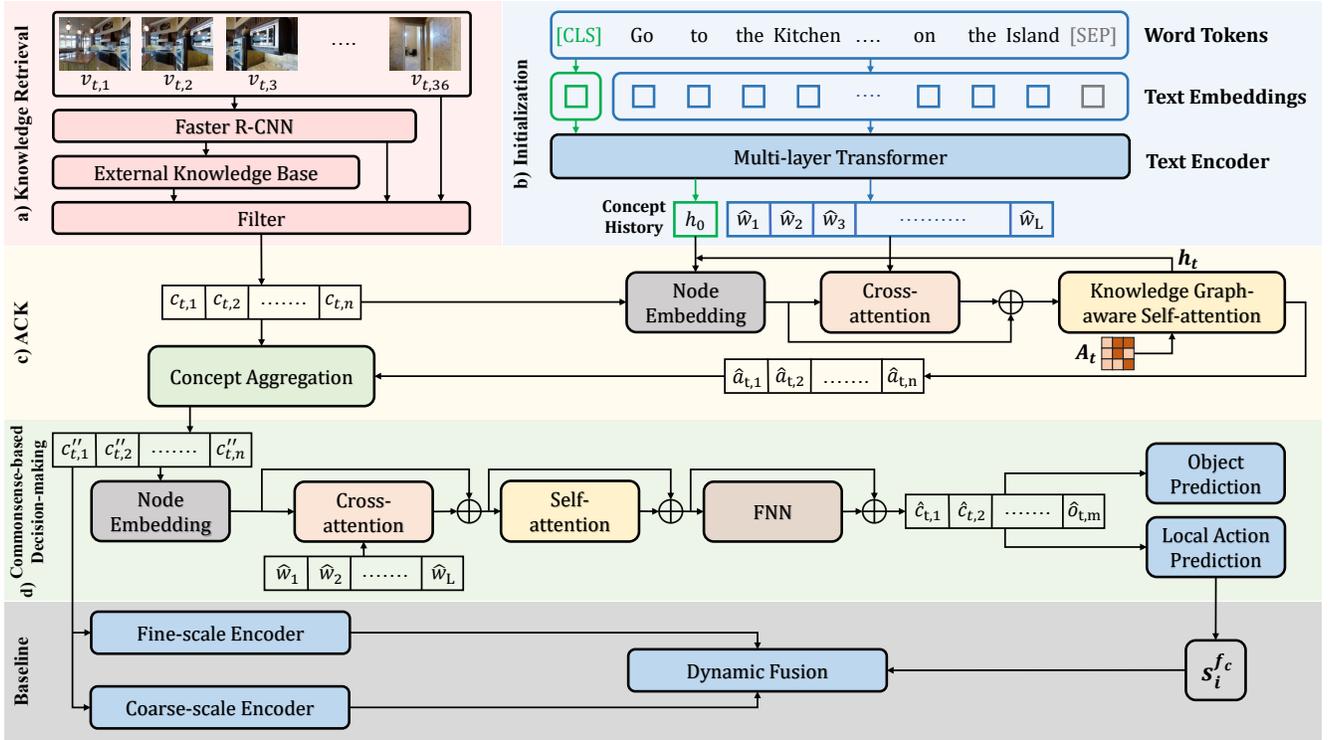


Figure 2: Main architecture of our proposed method. (a) Retrieving and refining the commonsense knowledge. (b) Initializing the concept history which represents the entire instruction and obtaining the text embedding. (c) ACK receives the detected objects, commonsense knowledge, and their temporal information to output weighted raw concept features which are utilized in the commonsense-based decision-making pipeline and the baseline model. (d) Inspired by the baseline agent, we add a new pipeline to produce the local action score and predict the object. Best viewed in color.

Method

In this section, we first explain the problem formulation and the overview of our proposed approach. Then, we present commonsense knowledge retrieval and ACK in detail and describe the commonsense-based decision-making pipeline for local action prediction.

Problem Formulation. In the context of REVERIE (Qi et al. 2020b), given a language instruction as a sequence of words denoted as $I = \{w_i\}_{i=1}^L$, where w_i represents the i^{th} word and L is the length of the sequence, the agent navigates through an undirected connectivity graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, where \mathcal{V} corresponds to viewpoints and \mathcal{E} shows the connection between nodes. At time step t , the agent is located at node \mathcal{V}_t and perceives a panoramic view $\mathcal{V}_t = \{v_{t,i}\}_{i=1}^{36}$. The agent infers an action a_t to transfer from state s_t to state s_{t+1} only based on the navigable directions $\mathcal{N}(\mathcal{V}_t) = \{v_{t,i}\}_{i=1}^K$, where $\mathcal{N}(\mathcal{V}_t) \subseteq \mathcal{V}_t$. Each state includes a triplet $\{v_{t,i}, \theta_{t,i}, \psi_{t,i}\}$, where $v_{t,i}$ is the viewpoint image, and $\{\theta_{t,i}, \psi_{t,i}\}$ are angels of heading and elevation, respectively, to determine the orientation of the image with respect to the agent. The agent is required to walk on the connectivity graph by selecting the next location at each node until it decides to stop or the number of action steps exceeds the threshold. The episode ends when the agent identifies the position of the target object within the panoramic view.

Method Overview. We follow DUET(Chen et al. 2022) architecture as the main baseline. This method includes two modules, topological mapping and global action planning. The former module is responsible for the gradual construction of a map over time by adding new observed locations during the path and updating the representation of each node. Afterward, the action, including the next location or stop action, is predicted by the latter module. DUET dynamically fuses action prediction of two scales: a fine-scale representation of the current location and a coarse-scale representation of the topological map to balance fine-grained language grounding against reasoning over the graphs.

In this work, we aim to improve the visual representation, which affects both fine- and coarse-scale encoders, and local action prediction. To do so, we incorporate commonsense knowledge into the baseline agent by leveraging the visible entities in viewpoint images and adding a new pipeline for local action reasoning. Through this paper, we use concepts and objects/knowledge interchangeably. To achieve more accurate alignment between instructions and candidate directions, we utilize object labels rather than their bounding boxes. As shown in Figure 2(a) we first build a knowledge base using detected objects and an external knowledge base followed by a filtering module to refine it by removing noisy and irrelevant data according to the entire image and objects.

In the initialization phase Figure 2(b), we obtain the text embedding and initialize concept history as a representation of the entire instruction. Then, we exploit ACK, Figure 2(c), to generate weighted raw concept features. Finally, we add a new pipeline to reason over the concept features and output the scores for object and local action prediction as illustrated in Figure 2(d). These scores are used as inputs for dynamic fusion to help the agent predict the correct local action.

Commonsense Knowledge Retrieval

Inspired by the human decision-making process which is based on background knowledge, we aim to provide the agent with complementary information to scene-level and object-level data. Commonsense knowledge not only helps the agent to understand the surroundings comprehensively, but it also facilitates data matching between the images and instructions. To obtain appropriate external knowledge, we first construct a knowledge base and then filter out the irrelevant information regarding visible content in each direction.

Object Detection. For $v_{t,i}$ which is the i^{th} view at time step t , the Faster R-CNN model (Ren et al. 2015) pre-trained on visual genome (VG) (Krishna et al. 2017) is utilized to obtain the object-set $\mathcal{O}_{v_{t,i}}$. To avoid overlooking helpful information we use all of the objects not just the most salient ones. During navigation, the adopted object detector is capable of differentiating 1600 categories $\{O_i\}_{i=1}^{1600}$, including those that have been annotated in the REVERIE dataset.

Knowledge Base Construction. To build the internal knowledge base, we employ ConceptNet (Liu and Singh 2004) as the external source of information. Each query we send to ConceptNet contains three parameters, start node, end node, and relationship type. The response from the ConceptNet is represented by a tuple $f_{i,j} = (s_i, r_{i,j}, e_j)$, which indicates that the start node s_i is connected to the end node e_j through the relationship $r_{i,j}$. More than 30 different relationships are available in ConceptNet, but we only use 8 most relevant ones based on their descriptions¹. For each object set $\mathcal{O}_{v_{t,i}}$, all related data is extracted from the knowledge base. For example, if $\mathcal{O}_{v_{t,i}}$ contains *bed*, then $\mathcal{K}_{v_{t,i}}$ includes *bedroom* according to the triplet (*bed*, *AtLocation*, *bedroom*) and the agent can infer it faces the *bedroom*.

Knowledge Selection. The extracted knowledge may be irrelevant or noisy, which could potentially impact the accuracy of the downstream task. To address this issue, we consider a refinement module to select the top- k pertinent supporting facts for viewpoints $\{k_{t,i}\}_{i=1}^K$. To do so, the pre-trained CLIP model (Radford et al. 2021) is used which consists of CLIP-I and CLIP-T encoders for encoding image and text, respectively, into a joint embedding space. We employ CLIP-I to encode entire images while object and knowledge labels are encoded by CLIP-T. Afterward, we calculate the similarity score for each fact according to its average cosine similarity with the whole image and its objects. The higher score means the corresponding knowledge is more suitable to be utilized. Finally, we select top- k commonsense knowledge. In this case, even for the same object sets we may retrieve different knowledge sets.

¹<https://github.com/commonsense/conceptnet5/wiki/Relations>

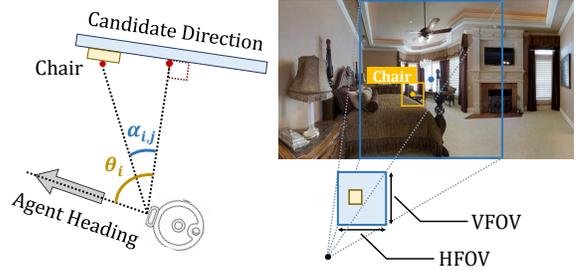


Figure 3: Encoding the relative position of objects with respect to the heading and elevation angles of the agent.

Augmented Commonsense Knowledge Model

At time step t , the object features \mathcal{O}_t , knowledge features \mathcal{K}_t , concept history features h_{t-1} , and instruction features \mathcal{I}_t are fed into our proposed model to either improve visual representation or local action prediction.

Knowledge Graph This study intends to treat objects, knowledge, and their history as a spatio-temporal knowledge graph. We first generate a fully-connected scene graph using detected objects and then expand it to construct the knowledge graph by adding extracted knowledge. To provide the agent with temporal insights about concepts during navigation, we consider an extra node that is connected to all the others. In the lines below, we elucidate the node representation and adjacency matrix generation.

Node Embedding. The acquired knowledge exists in textual format, thus, in order to maintain consistency between objects and knowledge and also to provide a better vision-text alignment, object labels are used instead of their visual features. To encode the object and knowledge labels CLIP-T is used. Node type encoding and directional encoding are added to each node feature. The node type encoding embeds 0 for history, 1 for objects, and 2 for knowledge to distinguish between different nodes. The directional encoding embeds the relative position of objects with respect to the heading and elevation angles of the agent. We aim to use objects as visual landmarks to improve alignment between textual and visual data. As shown in Figure 3, the relative orientation of objects is encoded as follows:

$$\{\mathcal{D}_{t,\mathcal{O}_{i,j}}\}_{j=1}^N = (\sin(\theta_i \pm \alpha_{i,j}), \cos(\theta_i \pm \beta_{i,j}), \sin(\psi_i \pm \alpha'_{i,j}), \cos(\psi_i \pm \beta'_{i,j})), \quad (1)$$

where $\mathcal{D}_{t,\mathcal{O}_{i,j}}$ is the directional encoding of j^{th} object in the viewpoint v_i at time step t and N is the number of objects. The directional encoding for knowledge is set to zero $\{\mathcal{D}_{t,\mathcal{K}_{i,j}}\}_{j=1}^K = \vec{0}$, where K is the number of knowledge.

Adjacency Matrix. At time step t , we generate a fully-connected scene graph for each navigable direction $v_{t,i}$ using the detected object set $\mathcal{O}_{v_{t,i}}$. We then add knowledge nodes and establish the connections between concepts if they are correlated in the knowledge base. Specifically, we connect two nodes that represent either objects or knowledge if their relationship is defined in the knowledge base

as the triplet $(c_i, r_{i,j}, c_j)$, where c_i and c_j denote concept nodes and $r_{i,j}$ denotes the relationship between them. The connections between nodes are considered undirected.

History Initialization It is not necessary to encode instructions during the navigation $t > 0$. Therefore, to encode the entire instructions at the initialization step $t = 0$, they are fed to a multi-layer transformer as the text encoder. The embedded [CLS] token is defined as the initial representation of concept history h_0 that represents the entire instruction.

$$h_0, \hat{\mathcal{I}} = \text{TextEncoder}([\text{CLS}], I, [\text{SEP}]) \quad (2)$$

where [CLS] and [SEP] are pre-defined tokens in BERT model (Devlin et al. 2019) and I is the language instruction.

Knowledge Graph-aware Cross-modal Encoder Having obtained concept features $[\mathcal{C}_t]$, including object and knowledge features $[\mathcal{O}_t; \mathcal{K}_t]$, and the history feature of previous time step h_{t-1} , we use a multi-layer cross-modal transformer to model the interaction between concepts, history, and instruction. Each layer of this transformer consists of a cross-attention layer and a self-attention layer. In the cross-attention layer, the concatenation of concept and history features $[h_{t-1}; \mathcal{C}_t]$ serves as queries while instruction embedding $\hat{\mathcal{I}}$ is fed to the transformer as keys and values. The cross-attention between $[h_{t-1}; \mathcal{C}_t]$ and $\hat{\mathcal{I}}$ is calculated as:

$$[h'_{t-1}; \mathcal{C}'_t] = \text{CrossAttention}([h_{t-1}; \mathcal{C}_t], \hat{\mathcal{I}}) \quad (3)$$

In general, transformers do not consider the structure of the input features. To address this problem, we introduce a knowledge graph-aware self-attention that slightly differs from the standard attention mechanism in standard transformers. In this case, the constructed knowledge graph structure is exploited to compute the attention as follows:

$$[h_t; \mathcal{A}_{t,C}] = \text{KGS}([h'_{t-1}; \mathcal{C}'_t])$$

$$\text{KGS}(X) = \text{Softmax}\left(\frac{XW_q(XW_k)^T}{\sqrt{d}} + M\right)XW_v \quad (4)$$

$$M = A_tW_a + b_a$$

where h_t is the concept history of the current step, $\mathcal{A}_{t,C}$ is concept self-attention scores, X is the node representation $[h_{t-1}; \mathcal{C}_t]$, A is the adjacency matrix, and W_a, b_a are two learnable parameters.

Concept Aggregation At time step t , \mathcal{C}_t^k is the concept tokens at head k and $\mathcal{A}_{t,C}^k$ is the attention scores over the concept tokens. Then, we average the score over all the attention heads ($K = 12$) and apply a Softmax function to get the overall concept attention weights as:

$$\hat{\mathcal{A}}_{t,C} = \text{Softmax}(\bar{\mathcal{A}}_{t,C}) = \text{Softmax}\left(\frac{1}{K} \sum_{k=1}^K \mathcal{A}_{t,C}^k\right) \quad (5)$$

Now, we perform a weighted sum over the input concept tokens to retrieve the weighted raw concept features as:

$$\mathcal{C}_t'' = \hat{\mathcal{A}}_{t,C} \mathcal{C}_t \quad (6)$$

where \mathcal{C}_t'' is the aggregated concept features since we need one concept feature per each viewpoint image.

Commonsense-based Decision-making Pipeline Inspired by the baseline agent, we add a new commonsense-based decision-making pipeline to predict a navigation score s^{fc} and an object score $s^{\mathcal{O}}$ to enhance local action reasoning. Afterward, these scores are used in the dynamic fusion module of the baseline model to calculate the final score and predict the next action.

Training and Inference We pre-train the baseline model on single-step action prediction (SAP) (Krantz et al. 2020), masked language modeling (MLM) (Devlin et al. 2019), masked region classification (MRC) (Lu et al. 2019), and object grounding (OG) (Lin, Li, and Yu 2021) tasks. However, our proposed model is only incorporated into policy learning. Analogous to the baseline, in addition to SAP loss L_{SAP} and OG loss L_{OG} , fine-tuning is guided by supervision provided by a pseudo-interactive demonstrator instead of behavioral cloning. In this case, the agent selects the next location with the overall shortest distance to the destination. We also use the loss calculated through the commonsense-based decision-making pipeline L_{CD} .

For inference, an action is predicted by the agent in each time step. If the agent exceeds the maximum number of action steps or the predicted action is the stop action, it stops at the current node. Otherwise, the agent moves to the predicted state. Eventually, when the agent stops at the final location, the object with the highest score is selected as the designated target object.

Experiments

Implementation Details

The ACK is not incorporated into pre-training tasks of DUET (Chen et al. 2022) and we only fine-tune the proposed model for 20k iterations on a single NVIDIA 3090 GPU. We use AdamW optimizer (Loshchilov and Hutter 2018) and the learning rate is 10^{-5} during the training. Similar to the baseline, viewpoint images and object bounding boxes, which have been provided by REVERIE, are encoded by ViT-B/16 (Dosovitskiy et al. 2020) pre-trained on ImageNet (Russakovsky et al. 2015). For the object detection task, we use the Faster R-CNN model (Ren et al. 2015) pre-trained on VG (Krishna et al. 2017). ConceptNet (Liu and Singh 2004) is utilized as the external commonsense knowledge base. To encode object and knowledge labels we employ the pre-trained CLIP model (Radford et al. 2021).

Dataset and Evaluation Metrics

REVERIE is a goal-oriented task with concise and high-level instructions that integrates R2R navigation with referring expression grounding. The agent navigates through the environment to identify the referred object that is not visible in the first view. In this dataset, the average length of instructions is 18 words. Also, there are more than 4,000 target objects falling into 489 categories. We utilize the widely-used and standard metrics for performance evaluation. trajectory length (TL), oracle success rate (OSR), success rate (SR),

Methods	Validation Unseen						Test Unseen					
	Navigation				Grounding		Navigation				Grounding	
	TL	OSR↑	SR↑	SPL↑	RGS↑	RGSPL↑	TL	OSR↑	SR↑	SPL↑	RGS↑	RGSPL↑
Human	-	-	-	-	-	-	21.18	86.83	81.51	53.66	77.84	51.44
Seq2Seq (Anderson et al. 2018)	11.07	8.07	4.20	2.84	2.16	1.63	10.89	6.88	3.99	3.09	2.00	1.58
SMNA (Ma et al. 2019)	9.07	11.28	8.15	6.44	4.54	3.61	9.23	8.39	5.80	4.53	3.10	2.39
RCM (Wang et al. 2019)	11.98	14.23	9.29	6.97	4.89	3.89	10.60	11.68	7.84	6.67	3.67	3.14
FAST-MATTN (Qi et al. 2020b)	45.28	28.20	14.40	7.19	7.84	4.67	39.05	30.63	19.88	11.61	11.28	6.08
CKR (Gao et al. 2021)	26.26	31.44	19.14	11.84	11.45	-	22.46	30.40	22.00	14.25	11.60	-
SIA (Lin, Li, and Yu 2021)	41.53	44.67	31.53	16.28	22.41	11.56	48.61	44.56	30.80	14.85	19.02	9.20
ORIST (Qi et al. 2021)	10.90	25.02	16.84	15.14	8.52	7.58	11.38	29.20	22.19	18.97	10.68	9.28
Airbert (Guhur et al. 2021)	18.71	34.51	27.89	21.88	18.23	14.18	17.91	34.20	30.28	23.61	16.83	13.28
RecBERT (Hong et al. 2021)	16.78	35.02	30.67	24.90	18.77	15.27	15.86	32.91	29.61	23.99	16.50	13.51
HOP (Qiao et al. 2022)	16.46	36.24	31.78	26.11	18.85	15.73	16.38	33.06	30.17	24.34	17.69	14.34
HAMT (Chen et al. 2021)	14.08	36.84	32.95	30.20	18.92	17.28	13.62	33.41	30.40	26.67	14.88	13.08
TD-STP (Zhao et al. 2022)	-	39.48	34.88	27.32	21.16	16.56	-	40.26	35.89	27.51	19.88	15.40
AZHP (Gao et al. 2023)	22.32	53.65	48.31	36.63	34.00	25.79	21.84	55.31	51.57	35.85	32.25	22.44
KERM (Li et al. 2023)	22.47	53.65	49.02	34.83	33.97	24.14	18.38	57.44	52.26	37.46	32.69	23.15
DUET (Chen et al. 2022)	22.11	51.07	46.98	33.73	32.15	23.03	21.30	56.91	52.51	36.06	31.88	22.06
ACK (Ours)	22.86	52.77	47.49	34.44	32.66	23.92	20.65	59.01	53.97	37.89	32.77	23.15

Table 1: Comparison of the agent performance with state-of-the-art methods on REVERIE dataset in the single-run setting.

and success rate penalized by path length (SPL) indicate the navigation performance. Moreover, remote grounding success rate (RGS) and remote grounding success rate weighted by path length (RGSPL) relate to object grounding task.

Comparison with State-of-the-Arts

Table 1 compares the single-run performance of ACK with state-of-the-art methods on the REVERIE benchmark. Our method achieves state-of-the-art performance and improves all the metrics on test unseen split which demonstrates the effectiveness and generalization ability of our proposed method in unseen environments. According to Table 1, our model remarkably outperforms DUET (Chen et al. 2022) and consistently enhances all metrics on both validation unseen and test unseen splits. In particular, compared to the baseline model on test unseen split, for navigation metrics, OSR, SR, and SPL improved by 2.10%, 1.46%, and 1.83%, respectively, and for grounding metrics, RGS and RGSPL are improved by 0.89% and 1.09%, respectively. Note that ACK is just incorporated into the fine-tuning stage of the baseline. Hence, to have a fair comparison, we mentioned the results of KERM (Li et al. 2023) while its pipeline is only composed of the fine-tuning phase. We also evaluate the performance of ACK on the R2R benchmark, however, no significant improvement is achieved. R2R contains fine-grained and detailed instructions which means the agent can strictly follow instructions, and exploring the environment is not essential. Thus, incorporating commonsense knowledge as complementary data is not very helpful for this task.

Ablation Study

The contribution of each element is assessed through comprehensive experiments which are shown in Table 2 and Table 3. The ACK is merely ablated on the validation unseen split of REVERIE.

KGS	CH	CD	OSR↑	SR↑	SPL↑	RGS↑	RGSPL↑
×	×	×	51.01	45.92	33.77	31.30	23.31
✓	×	×	51.36	46.18	33.81	32.04	23.33
✓	✓	×	52.23	48.08	34.02	32.25	23.39
✓	✓	✓	52.77	47.49	34.44	32.66	23.92

Table 2: Ablation of knowledge graph-aware self-attention, concept history, and commonsense-based decision-making pipeline on REVERIE validation unseen split. The performance is continuously enhanced as the suggested modules are gradually incorporated.

top- k	OSR↑	SR↑	SPL↑	RGS↑	RGSPL↑
0	51.49	46.72	33.91	32.42	23.56
10	52.77	47.49	34.44	32.66	23.92
20	51.66	47.25	33.98	32.49	23.67

Table 3: Ablation of utilizing top- k external knowledge in the training stage on REVERIE validation unseen split.

Knowledge Graph-aware Self-attention. As mentioned in Eq. 4 we ablate the transformer with and without knowledge graph structure in Table 2. According to this table, if the model is aware of the relationship between concepts, navigation performance is improved.

Concept History. Table 2 ablates the impact of incorporating temporal information of concepts into the proposed model. In REVERIE, it is essential to explore the environment more efficiently due to the lack of step-by-step instructions. Hence, As results in Table 2 suggest, holding a memory of visited objects in the path alongside the extracted



Figure 4: Visualization example of navigation performance for comparing ACK and the baseline method. We can see that our method predicts the correct action while DUET selects the wrong candidate direction. The concepts, including detected objects and retrieved commonsense knowledge, with the highest weights are used as landmarks in the instruction. Therefore, taking advantage of these concepts leads to visual representation enhancement and more accurate alignment between visual and textual information. Best viewed in color

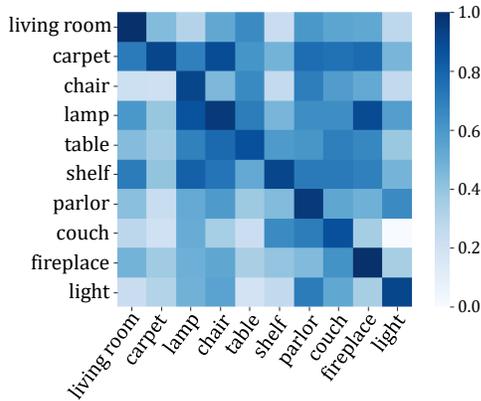


Figure 5: Learned concept-to-concept correlation matrix

relevant commonsense knowledge helps the agent act more properly in unseen environments.

Commonsense-base Decision-making. To show the effectiveness of incorporating commonsense knowledge in the decision-making process, we evaluate the ACK with and without utilizing the new pipeline for action prediction. Table 2 demonstrates that incorporating the commonsense-based decision-making pipeline into the dynamic fusion module of the baseline makes the agent learn from the knowledge graph and improves local action reasoning.

External Knowledge Capacity. To show the impact of commonsense knowledge capacity on agent performance, we evaluate our method by different numbers of knowledge as illustrated in Table 3. Note that, $k = 0$ means that we only use the fully-connected scene graph constructed by visible objects in viewpoint images. Increasing the number of knowledge, $k = 10$, results in better performance, however, by continuously increasing the number of extracted knowledge, $k = 20$, the performance of the agent is deteriorated. It shows extra information may cause noise in the output.

Qualitative Analysis

To visualize our proposed method, we use an example in the validation unseen split of REVERIE. According to the instruction, the target object is the white basket. Figure 4 illustrates an example where ACK selects the correct candidate direction by detecting objects and retrieving commonsense knowledge, while the baseline model makes the wrong decision without the supporting facts. In this figure, a heat map is used to visualize the weight distribution over the top 10 concepts of the selected viewpoint, including objects and knowledge, after the knowledge graph-aware cross-modal encoder. Regarding the instruction, *couch* and *table* are used as landmarks. We can see that annotated concepts in red, e.g., *living room*, *couch*, and *table*, can be leveraged as visual landmarks as well as improve the alignment between textual and visual data. Furthermore, the other retrieved concepts such as *chair* and *fireplace*, with higher weights can also be useful for navigation in this example. Also, the learned concept-to-concept correlations for this example are visualized in Figure 5.

Conclusion

In this paper, we propose ACK to enhance visual representation and local action prediction by incorporating commonsense knowledge into the REVERIE task as a spatio-temporal knowledge graph. At first, a knowledge base is constructed and then refined to output purified commonsense information. We further design ACK to leverage the refined commonsense knowledge, which consists of two modules, the knowledge graph-aware cross-modal encoder, and the concept aggregator. The absence of step-by-step instructions in the REVERIE motivated us to hold a concept history during the navigation for more efficient exploration. The experimental results demonstrate the superiority of ACK over the state-of-the-art methods on REVERIE which shows that taking advantage of commonsense knowledge is a promising direction for the REVERIE task. For future work, we aim to exploit other external repositories as well as different formats of commonsense knowledge.

References

- Anderson, P.; Wu, Q.; Teney, D.; Bruce, J.; Johnson, M.; Sünderhauf, N.; Reid, I.; Gould, S.; and van den Hengel, A. 2018. Vision-and-Language Navigation: Interpreting Visually-Grounded Navigation Instructions in Real Environments. In *CVPR*.
- Auer, S.; Bizer, C.; Kobilarov, G.; Lehmann, J.; Cyganiak, R.; and Ives, Z. 2007. Dbpedia: A nucleus for a web of open data. In *The Semantic Web: 6th International Semantic Web Conference, 2nd Asian Semantic Web Conference, ISWC 2007+ ASWC 2007, Busan, Korea, November 11-15, 2007. Proceedings*, 722–735. Springer.
- Chen, S.; Guhur, P.-L.; Schmid, C.; and Laptev, I. 2021. History Aware Multimodal Transformer for Vision-and-Language Navigation. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 5834–5847. Curran Associates, Inc.
- Chen, S.; Guhur, P.-L.; Tapaswi, M.; Schmid, C.; and Laptev, I. 2022. Think Global, Act Local: Dual-Scale Graph Transformer for Vision-and-Language Navigation. In *CVPR*, 16537–16547.
- Deng, Z.; Narasimhan, K.; and Russakovsky, O. 2020. Evolving Graphical Planner: Contextual Global Planning for Vision-and-Language Navigation. In *NeurIPS*, volume 33, 20660–20672.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, volume 1, 2.
- Ding, Y.; Yu, J.; Liu, B.; Hu, Y.; Cui, M.; and Wu, Q. 2022. MuKEA: Multimodal Knowledge Extraction and Accumulation for Knowledge-Based Visual Question Answering. In *CVPR*, 5089–5098.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*.
- Fried, D.; Hu, R.; Cirik, V.; Rohrbach, A.; Andreas, J.; Morency, L.-P.; Berg-Kirkpatrick, T.; Saenko, K.; Klein, D.; and Darrell, T. 2018. Speaker-Follower Models for Vision-and-Language Navigation. In *NeurIPS*, volume 31.
- Gao, C.; Chen, J.; Liu, S.; Wang, L.; Zhang, Q.; and Wu, Q. 2021. Room-and-Object Aware Knowledge Reasoning for Remote Embodied Referring Expression. In *CVPR*, 3064–3073.
- Gao, C.; Peng, X.; Yan, M.; Wang, H.; Yang, L.; Ren, H.; Li, H.; and Liu, S. 2023. Adaptive Zone-Aware Hierarchical Planner for Vision-Language Navigation. In *CVPR*, 14911–14920.
- Guhur, P.-L.; Tapaswi, M.; Chen, S.; Laptev, I.; and Schmid, C. 2021. Airbert: In-Domain Pretraining for Vision-and-Language Navigation. In *ICCV*, 1634–1643.
- Hao, W.; Li, C.; Li, X.; Carin, L.; and Gao, J. 2020. Towards Learning a Generic Agent for Vision-and-Language Navigation via Pre-Training. In *CVPR*.
- Hong, Y.; Rodriguez, C.; Qi, Y.; Wu, Q.; and Gould, S. 2020. Language and Visual Entity Relationship Graph for Agent Navigation. In *NeurIPS*, volume 33, 7685–7696.
- Hong, Y.; Wu, Q.; Qi, Y.; Rodriguez-Opazo, C.; and Gould, S. 2021. VLN BERT: A Recurrent Vision-and-Language BERT for Navigation. In *CVPR*, 1643–1653.
- Krantz, J.; Wijmans, E.; Majumdar, A.; Batra, D.; and Lee, S. 2020. Beyond the nav-graph: Vision-and-language navigation in continuous environments. In *ECCV*, 104–120.
- Krishna, R.; Zhu, Y.; Groth, O.; Johnson, J.; Hata, K.; Kravitz, J.; Chen, S.; Kalantidis, Y.; Li, L.-J.; Shamma, D. A.; et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 123: 32–73.
- Li, X.; Wang, Z.; Yang, J.; Wang, Y.; and Jiang, S. 2023. KERM: Knowledge Enhanced Reasoning for Vision-and-Language Navigation. In *CVPR*, 2583–2592.
- Li, X.; Zhang, Y.; Yuan, W.; and Luo, J. 2022. Incorporating External Knowledge Reasoning for Vision-and-Language Navigation with Assistant’s Help. *Applied Sciences*, 12(14): 7053.
- Lin, B.; Zhu, Y.; Chen, Z.; Liang, X.; Liu, J.; and Liang, X. 2022. ADAPT: Vision-Language Navigation With Modality-Aligned Action Prompts. In *CVPR*, 15396–15406.
- Lin, X.; Li, G.; and Yu, Y. 2021. Scene-Intuitive Agent for Remote Embodied Visual Grounding. In *CVPR*, 7036–7045.
- Liu, C.; Zhu, F.; Chang, X.; Liang, X.; Ge, Z.; and Shen, Y.-D. 2021. Vision-Language Navigation With Random Environmental Mixup. In *ICCV*, 1644–1654.
- Liu, H.; and Singh, P. 2004. ConceptNet—a practical commonsense reasoning tool-kit. *BT technology journal*, 22(4): 211–226.
- Loshchilov, I.; and Hutter, F. 2018. Decoupled weight decay regularization. *ICLR*.
- Lu, J.; Batra, D.; Parikh, D.; and Lee, S. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *NeurIPS*, 32.
- Ma, C.-Y.; Lu, J.; Wu, Z.; AlRegib, G.; Kira, Z.; Socher, R.; and Xiong, C. 2019. Self-monitoring navigation agent via auxiliary progress estimation. *ICLR*.
- Marino, K.; Chen, X.; Parikh, D.; Gupta, A.; and Rohrbach, M. 2021. KRISP: Integrating Implicit and Symbolic Knowledge for Open-Domain Knowledge-Based VQA. In *CVPR*, 14111–14121.
- Marino, K.; Rastegari, M.; Farhadi, A.; and Mottaghi, R. 2019. OK-VQA: A Visual Question Answering Benchmark Requiring External Knowledge. In *CVPR*.
- Pashevich, A.; Schmid, C.; and Sun, C. 2021. Episodic Transformer for Vision-and-Language Navigation. In *ICCV*, 15942–15952.
- Qi, Y.; Pan, Z.; Hong, Y.; Yang, M.-H.; van den Hengel, A.; and Wu, Q. 2021. The Road To Know-Where: An Object-and-Room Informed Sequential BERT for Indoor Vision-Language Navigation. In *ICCV*, 1655–1664.

- Qi, Y.; Pan, Z.; Zhang, S.; van den Hengel, A.; and Wu, Q. 2020a. Object-and-Action Aware Model for Visual Language Navigation. In *ECCV*, 303–317.
- Qi, Y.; Wu, Q.; Anderson, P.; Wang, X.; Wang, W. Y.; Shen, C.; and Hengel, A. v. d. 2020b. REVERIE: Remote Embodied Visual Referring Expression in Real Indoor Environments. In *CVPR*.
- Qiao, Y.; Qi, Y.; Hong, Y.; Yu, Z.; Wang, P.; and Wu, Q. 2022. HOP: History-and-Order Aware Pre-Training for Vision-and-Language Navigation. In *CVPR*, 15418–15427.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. In Meila, M.; and Zhang, T., eds., *icml*, volume 139, 8748–8763.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. 2015. Imagenet large scale visual recognition challenge. *IJCV*, 115: 211–252.
- Tan, H.; Yu, L.; and Bansal, M. 2019. Learning to navigate unseen environments: Back translation with environmental dropout. In *NAACL*, 2610–2621.
- Wang, H.; Wang, W.; Liang, W.; Xiong, C.; and Shen, J. 2021. Structured Scene Memory for Vision-Language Navigation. In *CVPR*, 8455–8464.
- Wang, X.; Huang, Q.; Celikyilmaz, A.; Gao, J.; Shen, D.; Wang, Y.-F.; Wang, W. Y.; and Zhang, L. 2019. Reinforced Cross-Modal Matching and Self-Supervised Imitation Learning for Vision-Language Navigation. In *CVPR*.
- Zhao, C.; Qi, Y.; and Wu, Q. 2023. Mind the Gap: Improving Success Rate of Vision-and-Language Navigation by Revisiting Oracle Success Routes. In *Proceedings of the 31st ACM International Conference on Multimedia*, 4349–4358.
- Zhao, Y.; Chen, J.; Gao, C.; Wang, W.; Yang, L.; Ren, H.; Xia, H.; and Liu, S. 2022. Target-driven structured transformer planner for vision-language navigation. In *Proceedings of the 30th ACM International Conference on Multimedia*, 4194–4203.
- Zhu, F.; Liang, X.; Zhu, Y.; Yu, Q.; Chang, X.; and Liang, X. 2021. SOON: Scenario Oriented Object Navigation With Graph-Based Exploration. In *CVPR*, 12689–12699.