

# UNeMo: Collaborative Visual-Language Reasoning and Navigation via a Multimodal World Model

Changxin Huang<sup>1,2</sup>, Lv Tang<sup>1,2</sup>, Zhaohuan Zhan<sup>3</sup>, Lisha Yu<sup>4</sup>, Runhao Zeng<sup>5</sup>,  
Zun Liu<sup>1,2</sup>, Zhengjie Wang<sup>6</sup>, Jianqiang Li<sup>1,2\*</sup>

<sup>1</sup> School of Artificial Intelligence, Shenzhen University

<sup>2</sup> National Engineering Laboratory for Big Data System Computing Technology, Shenzhen University

<sup>3</sup> Department of Engineering, Shenzhen MSU-BIT University

<sup>4</sup> School of Computer Science and Engineering, Sun Yat-sen University

<sup>5</sup> Artificial Intelligence Research Institute, Shenzhen MSU-BIT University

<sup>6</sup> School of Mechatronic Engineering, Beijing Institute of Technology

huangchx@szu.edu.cn, lijq@szu.edu.cn

## Abstract

Vision-and-Language Navigation (VLN) requires agents to autonomously navigate complex environments via visual images and natural language instructions—remains highly challenging. Recent research on enhancing language-guided navigation reasoning using pre-trained large language models (LLMs) has shown promising prospects. However, the reasoning of such methods is limited to the linguistic modality, lacking visual reasoning capabilities. Moreover, existing reasoning modules are optimized separately from navigation policies, leading to incompatibility and potential conflicts in optimization objectives. To tackle these challenges, we introduce UNeMo, a novel framework designed for the collaborative optimization of visual state reasoning and navigational decision-making. It introduces a Multimodal World Model (MWM) that takes visual features, language instructions, and navigational actions as inputs to jointly predict subsequent visual states, enabling cross-modal reasoning. Via a Hierarchical Prediction-Feedback (HPN) mechanism, MWM collaborates with navigation policies: the first layer generates actions using current vision-and-language features; MWM then infers post-action visual states to guide the second layer’s fine-grained decisions. This forms a dynamic bidirectional promotion mechanism where MWM reasoning optimizes navigation policies, while policy decisions feedback to improve MWM’s reasoning accuracy. Experiments on R2R and REVERIE datasets show UNeMo outperforms state-of-the-art methods by 2.1% and 0.7% in navigation accuracy for unseen scenes, validating its effectiveness.

## Introduction

Vision-and-Language Navigation (VLN) requires an embodied agent to achieve autonomous navigation from the current to the target position in an unknown environment based on visual observations and language instructions (Anderson et al. 2018; Krantz et al. 2020a; Qi et al. 2020). This capability constitutes the core foundation of an embodied intelligent system. This task not only needs to address the problem of efficient fusion of cross-modal information (Gu

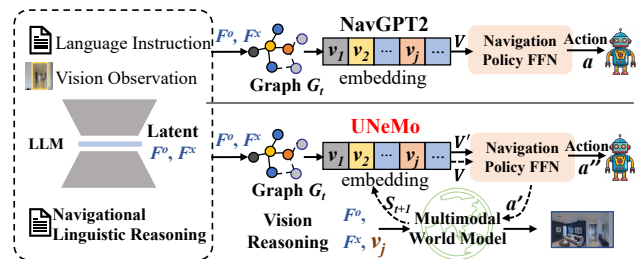


Figure 1: Comparison of main differences between UNeMo and NavGPT2. UNeMo introduces MWM for visual state reasoning and achieves joint optimization with the navigation policy via a hierarchical predictive-feedback navigator.

et al. 2022; Wu et al. 2024b) but also poses new challenges to the construction of a perception-reasoning collaborative mechanism in dynamic environments (Zhang et al. 2024b).

Traditional VLN methods use tailored end-to-end deep learning frameworks, such as sequence-to-sequence models (Fried et al. 2018), attention-based cross-modal fusion (Hong et al. 2021a), and graph-based representation learning (Krantz et al. 2020b). These jointly optimize vision-language alignment and action policy learning to map multimodal inputs to navigation decisions (An et al. 2024; Yue et al. 2024). While performing well on large task-specific datasets, they show poor policy generalization in out-of-distribution environments or with novel instructions.

Large language models (LLMs) have recently shown strong abilities in knowledge generalization and complex reasoning, prompting research into their use for enhancing VLN with promising results. Researchers have attempted to use LLMs to assist VLN in tasks such as navigation instruction parsing (Yang et al. 2024), enhancing scene descriptions (Wu et al. 2024a), and extracting task knowledge (Shah et al. 2023; Schumann et al. 2024; Wang et al. 2023a). NavGPT further leverages the decision-making capabilities of LLMs to directly predict navigation actions (Zhou, Hong, and Wu 2024). However, due to the lack of training on navigation data, the performance of such methods is inferior to tradi-

\*Corresponding author

tional VLN-specialized approaches.

Accordingly, NavGPT2 employs navigation data to conduct knowledge distillation on LLMs, thereby constructing a navigation-specialized LLM (Zhou et al. 2024), as illustrated in Fig. 1. Equipped with a navigational reasoning LLM, this approach achieves performance on par with VLN-specialized methods. However, NavGPT2 is trained exclusively via language-guided reasoning modules, rendering it incapable of visual state reasoning. Furthermore, its navigational reasoning module is frozen during the training phase of navigation policies, creating a disconnect between the reasoning process and navigation optimization. This hinders the dynamic refinement of navigational reasoning capabilities in response to policy training progress, thereby imposing inherent limitations on navigation performance.

To address the aforementioned issues, we propose an Unlock Next Moment (UNeMo) framework, which enables the collaborative optimization of cross-modal reasoning and navigation policies. Specifically, the framework introduces a Multimodal World Model (MWM) to implement vision state reasoning: it takes the agent’s current observed visual features, language instruction features, and navigation action prediction results as inputs, and predicts the next visual state through joint modeling, as shown at the bottom of Fig. 1. The MWM and navigation policies adopt a synchronous optimization mechanism, endowing the navigation policies with vision-language cross-modal reasoning capabilities. To enhance the collaborative effectiveness between the world model and navigation policies, this study designs a Hierarchical Prediction-Feedback Navigator (HPFN). This navigator employs a hierarchical decision-making architecture: first, it performs coarse-grained navigation action prediction based on current vision-language features; subsequently, it uses the MWM to predict the visual state of the future node that the agent will reach after executing the action, and integrates this visual feedback feature into the decision-making process of the navigation policy to guide fine-grained navigation action prediction.

Through the synchronous optimization of multimodal reasoning and navigation decision-making, UNeMo constructs a dynamic bidirectional promotion mechanism: the decisions of the navigation policy feed back into the state prediction of the world model, while the reasoning results of the world model continuously optimize the navigation policy.

The contributions can be summarized as follows:

- We propose MWM, a multimodal world model that jointly captures visual observations, language instructions, and navigation-action predictions to enable dynamic visual state reasoning.
- We design a hierarchical prediction-feedback navigator, employing a closed-loop interaction paradigm of coarse-grained action prediction  $\rightarrow$  world-model visual state feedback  $\rightarrow$  fine-grained decision refinement, thereby strengthening the synergy between navigational reasoning and decision making.
- Extensive experiments on the R2R and REVERIE benchmarks demonstrate that, in unseen environments, UN-

eMo outperforms state-of-the-art methods by 2.1% and 0.7% in success rate, respectively, underscoring its efficacy in enhancing generalization and navigation accuracy in complex VLN scenarios.

## Related Work

Traditional research in Vision-Language Navigation (VLN) has centered on state representation learning, including recurrent neural network (RNN)-based historical state modeling (Hong et al. 2021b), sequential encoding for temporal information processing (Chen et al. 2021), topological map construction for spatial connectivity (An et al. 2024; Chen et al. 2022), and grid map development for precise localization (Liu et al. 2023). Additionally, frameworks such as SUSA (Zhang et al. 2024a) have introduced hybrid semantic-spatial representations to enhance vision-language alignment. Recent advances have shifted toward leveraging Large Language Models (LLMs) to enhance navigation capabilities.

**Navigation with LLM.** LLMs, with their strong semantic understanding and reasoning abilities, have been increasingly applied to navigation agents (Huang et al. 2022; Yao et al. 2023; Shah et al. 2023; Schumann et al. 2024; Wang et al. 2023a). Early approaches decoupled LLMs from core navigation systems—using GPT-2 for subtask hints (Qiao et al. 2023) or GPT-4 prompts to translate instructions into actions (Song et al. 2025)—limiting holistic integration. Recent work embeds LLMs directly, such as zero-shot NavGPT (Zhou, Hong, and Wu 2024), MapGPT (Chen et al. 2024), MC-GPT (Zhan et al. 2024), and fine-tuned LLaMA-7B variants like LangNav (Pan et al. 2023) and NavCoT (Lin et al. 2025). Despite improvements, these agents still lag behind specialized VLN models in accuracy and transparency. NavGPT-2, which freezes the LLM as a vision-language encoder with imitation learning (Zhou et al. 2024), lacks dynamic scene adaptation. In contrast, our work integrates a future-state prediction module to dynamically forecast environmental changes for enhanced decision-making.

**World Model for VLN.** World models (Ha and Schmidhuber 2018) compress high-dimensional sensory data into dynamic state representations, supporting embodied decision-making. Prior studies explored diverse architectures for environmental feature encoding (Chang, Gupta, and Gupta 2020; Purushwalkam et al. 2020), while sampling-based planners used these representations for self-supervised policy learning (Sekar et al. 2020; Yen-Chen, Bauza, and Isola 2020). Pathdreamer (Koh et al. 2021) and Dreamwalker (Wang et al. 2023b) predict future observations (pixel-level or abstract) to extend planning horizons but focus solely on forward exploration. LookAhead (Wang et al. 2024) employs a hierarchical neural radiance field to generate multi-scale future embeddings. However, these methods separate state prediction from policy learning. Our work introduces a joint framework for simultaneous future state prediction and proposes a hierarchical prediction-feedback mechanism that co-optimizes state inference and action selection, fully integrating predicted future states into navigation decisions for more precise navigation.

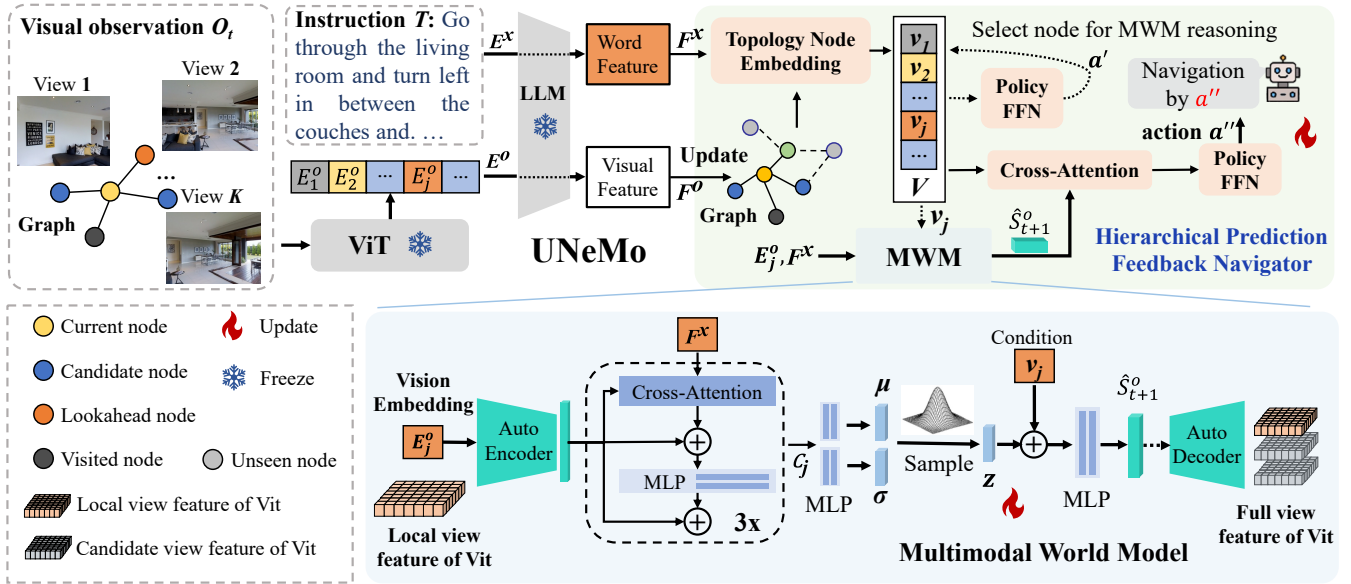


Figure 2: Overview of the UNeMo framework. Similar to NavGPT2, the input visual observation  $O_t$  and instruction  $T$  are processed through a pre-trained LLM encoder to obtain visual and word features, respectively. UNeMo introduces two key components: (1) The Multimodal State Reasoning module receives partial-view representations of the highest-scoring action node and predicts its complete post-execution visual state; (2) The Decision Making module fuses these state-reasoning results with linguistic features for final navigation action selection.

## Preliminary

### Problem setup

In the VLN task, agents need to parse multimodal information to achieve goal-oriented path planning in discrete topological spaces. Specifically, an agent must understand the global natural language instruction  $I$  and select a navigational action  $a_t$  based on the current visual observation  $v_t$ . This process is repeated until the target location is reached or the number of steps exceeds a threshold.

The action prediction in VLN can be formulated as a node-selection problem. Early methods confined this space to the local neighborhood of the agent’s current node; more recent approaches have extended it to the global scale by constructing an explicit topological map (Chen et al. 2022; Zhou et al. 2024). The agent builds a topological map  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  based on the nodes it has visited and the nodes within its visual range, where  $\mathcal{V} = \{v_1, v_2, \dots, v_N\}$  denotes the set of  $N$  navigable nodes, and  $\mathcal{E}$  represents the set of connecting edges between nodes. At each time step  $t$ , the agent acquires observations of adjacent navigable nodes at the current position and updates the topological map  $\mathcal{G}_t$ . For visited nodes, the node embedding is the observation at that node. For unvisited but partially observable nodes, the node embedding aggregates the average of all partial-view embeddings observed from other visited nodes. The final action prediction  $a_t$  corresponds to selecting a node  $v \in \mathcal{V}$ .

### LLM-based navigator

We adopt NavGPT2 (Zhou et al. 2024), one of the most advanced LLM-based navigation approaches to date, as our

base model. NavGPT2 incorporates three core components: node embedding, decision making, and loss function. For node embedding, it employs a Topology Node Embedding (TNE) encoder. Specifically, let  $E^o$  and  $E^x$  denote the visual (observation) and textual (instruction) features, respectively. These are jointly encoded by a pretrained LLM to produce the corresponding features  $F^o$  and  $F^x$ . The TNE encoder then projects these embeddings into a shared topological space to form the node embeddings used for navigation:

$$\mathcal{V}_t = \text{TNE}(F^o, F^x). \quad (1)$$

For decision making, a two-layer feedforward network is used to predict the score of each navigational action, and the node with the highest score is selected for action execution. The decision-making process is as follows:

$$a_t = \max(\text{Softmax}(\text{FFN}(\mathcal{V}_t))). \quad (2)$$

For the loss function, NavGPT2 integrates behavioral cloning and DAGger (Ross, Gordon, and Bagnell 2011) objectives into a single training criterion:

$$\mathcal{L}_{\text{base}} = \lambda \mathcal{L}_{\text{BC}} + \mathcal{L}_{\text{DAG}}, \quad (3)$$

where  $\mathcal{L}_{\text{BC}} = -\sum_{t=1}^T \log \pi(v_t^* | \mathcal{I}, \mathcal{G}_t)$  optimizes for expert demonstrations  $v_t^*$ , while  $\mathcal{L}_{\text{DAG}} = -\sum_{t=1}^T \log \pi(\tilde{v}_t^* | \mathcal{I}, \tilde{\mathcal{G}}_t)$  leverages pseudo-labels  $\tilde{v}_t^*$  derived from shortest-path analysis on the agent’s topological map  $\tilde{\mathcal{G}}_t$ , with hyperparameter  $\lambda$  balancing their contributions.

## Method

### Overview

Our framework is illustrated in Figure 2. Similar to NavGPT2, input visual observations  $O_t$  and instructions  $T$  are processed via a pre-trained LLM encoder to generate visual and word features, respectively. To endow the agent with visual reasoning capabilities, we introduce a Multimodal World Model (MWM) for visual state prediction: this module takes the current visual state encoding, instruction encoding, and navigation action as inputs, and predicts the next visual state the agent will encounter. This predicted visual state updates the embedding of each navigation node, enriching the agent’s representation with forward-looking information to support more informed navigation decisions. To mitigate objective conflicts arising from separate optimization of the navigation policy and reasoning module, we further propose a Hierarchical Prediction-Feedback Navigator. This mechanism enables bidirectional interaction: navigation policy decisions feed back to refine the world model’s state predictions, while the world model’s reasoning outputs continuously optimize the navigation policy. Together, these components facilitate the joint optimization of navigation reasoning and decision-making processes.

### Multimodal World Model

While topological maps encode historical states—encompassing visual observations and node connectivity—they lack information about future states. Given that future state prediction enhances navigation performance (Li and Bansal 2023; Wang et al. 2023b, 2024), we introduce a Multimodal World Model (MWM) to predict future visual states, leveraging the topological map’s representation of past states. We employ Conditional Variational Autoencoders (CVAE) (Sohn, Lee, and Yan 2015) to construct the MWM, which comprises an encoder for aligning and fusing multimodal inputs, and a decoder for reasoning about the next-step visual state.

The encoder uses cross-attention to align visual observations with language instructions. At time  $t$ , the navigator generates a node-scoring distribution  $a_t$  over the topological map, selecting the highest-scoring node  $j$ . The partially observed view embedding  $E_j^o \in \mathbb{R}^{1 \times 768}$  (visual observation of neighbor  $j$ ) is combined with instruction feature  $F^x \in \mathbb{R}^{N \times 768}$ ; both feed into cross-attention for fusion, yielding  $c_j \in \mathbb{R}^{1 \times 768}$ , formally:

$$c_j = \text{Softmax} \left( \frac{\mathcal{W}_q E_j^o (\mathcal{W}_k F^x)^\top}{\sqrt{d}} \right) \mathcal{W}_v F^x, \quad (4)$$

$\mathcal{W}_*$  denotes learnable projection matrices, with subscripts  $q, k, v$  indicating mappings to attention query, key, and value vectors;  $d$  is the scaling factor. Specifically, observed-view embedding  $E_j^o$  projects to the query, and instruction embedding  $F^x$  projects to key/value. Stacking  $L = 3$  cross-attention layers yields the final cross-modal embedding  $c_j \in \mathbb{R}^{1 \times 768}$  (matching query dimensionality). This fused feature is fed through two MLPs to estimate future-state distribution parameters: mean  $\mu_t$  and variance  $\sigma_t^2$ .

In the decoder, the mean  $\mu_t$  and variance  $\sigma_t^2$  are first reparameterized to sample the latent variable  $z_t = \mu + \sigma \odot \varepsilon, \varepsilon \sim \mathcal{N}(0, 1)$ . In the decoder,  $\mu_t$  and  $\sigma_t^2$  are reparameterized to sample latent  $z_t = \mu + \sigma \odot \varepsilon$  ( $\varepsilon \sim \mathcal{N}(0, 1)$ ). Node  $j$ ’s intrinsic embedding  $v_j^t$  is concatenated with  $z_t$  as a conditional feature; an MLP then uses this vector to predict future visual-state embedding  $\hat{S}_{t+1}^o = \text{MLP}(z_t; v_j^t)$ . Therefore, The entire visual state reasoning process of MWM can be formulated as:

$$\hat{S}_{t+1}^o = \text{CVAE}(E_j^o, F^x | v_j), \quad (5)$$

this reasoning visual-state is integrated into the subsequent Hierarchical Prediction-Feedback Navigator to enhance navigation performance.

### Hierarchical Prediction-Feedback Navigator

Although the MWM enables prediction of future visual states, integrating these predictions into the navigation policy remains an open challenge. Here we employ the prediction of MWM to update node embeddings and guide decision making.

To update the node embeddings, we first pass topological-map node embeddings  $\mathcal{V}_t = \{v_{t,i}\}_{i=1}^K$  through a two-layer feedforward network to predict the look ahead action  $a_t'$ . Then we employ the MWM to predict the visual-state embedding after executing the lookahead exploration  $\hat{S}_{t+1}^o$  as described in Eq. 5. This prediction  $\hat{S}_{t+1}^o$  embodies the expected observation once the agent arrives at node  $j$ . We introduce a cross-attention module to enable deep interaction between the topological-map node features and the MWM’s predicted state. Specifically, we treat the node embeddings  $\mathcal{V}_t$  as query vectors, and the predicted next-state embedding  $\hat{S}_{t+1}^o$  as key/value vectors. Between cross-attention layers, we use the residual connection to preserve the original node information. Through successive cross-attention layers, the node embeddings are progressively enriched with lookahead information, yielding updated representations:

$$\mathcal{V}_t' = \text{Softmax} \left( \frac{\mathcal{W}_q \mathcal{V}_t (\mathcal{W}_k \hat{S}_{t+1}^o)^\top}{\sqrt{d}} \right) \mathcal{W}_v \hat{S}_{t+1}^o, \quad (6)$$

where  $\mathcal{W}_*$  denotes a learnable weight matrix, and the subscripts  $q, k$  and  $v$  refer, respectively, to the query, key, and value projections in the attention mechanism.  $d$  is the scaling factor.

To incorporate this lookahead prediction into the navigation policy, the updated node embeddings  $\mathcal{V}_t'$  are fed into the policy head to compute the action-probability distribution for the next time step:

$$a_t'' = \max(\text{Softmax}(\text{FFN}(\mathcal{V}_t'))). \quad (7)$$

## Experiments

### Experiment Settings

**Datasets** Our work primarily focus on the path instruction following task on the Room-to-Room (R2R) benchmark dataset (Anderson et al. 2018). R2R requires an agent to follow step-by-step natural language instructions, averaging

32 words in length. To rigorously evaluate the generalization capabilities of our approach, we additionally conduct evaluations on the REVERIE benchmark (Qi et al. 2020). REVERIE presents more challenging high-level instructions (averaging 21 words), such as "Find the cell phone on the nightstand in the bedroom", and requires the agent to navigate to the target location and localize a distant target object within a pre-defined ground-truth bounding box. Notably, the expert demonstration path lengths in REVERIE range from 4 to 7 steps, compared to an average of 6 steps in R2R.

**Evaluation Metrics** We adopt the widely recognized evaluation framework in vision-and-language navigation. For core navigation performance assessment, key metrics include: Success Rate (SR), measuring the proportion of paths where the agent’s final position is within a preset threshold (typically 3 meters) from the goal; Success weighted by Path Length (SPL), which jointly considers task success and path efficiency; Trajectory Length (TL) representing the average path length in meters; Navigation Error (NE) indicating the average distance between the agent’s final position and the target; and Oracle Success Rate (OSR) showing success rate under ideal stopping policy. For object referral evaluation, following REVERIE (Qi et al. 2020), we employ Remote Grounding Success (RGS) to assess the proportion of successfully reaching the correct area and accurately locating the target object, along with its path-length-weighted variant (RGSPL). Among all metrics, except TL and NE where lower values indicate better performance, higher values represent superior results for other indicators.

**Implementation Details** To demonstrate that UNeMo serves as a generalizable framework applicable to diverse baseline architectures, we validate its effectiveness on two distinct topology-based navigation policies: NavGPT2 (Zhou et al. 2024) and DUET (Chen et al. 2022). Both policies share a common limitation: while capable of processing vision-language inputs and constructing topological representations via imitation learning, they lack prospective state reasoning capabilities. To address this gap, we integrate explicit state reasoning tasks into both systems using the UNeMo framework. Our experiments retain all baseline pre-training results and introduce UNeMo exclusively during the fine-tuning phase. An important implementation distinction lies in their modality processing: DUET employs conventional vision-language pretraining for feature extraction, whereas NavGPT2 acquires representations via LLM intermediate layer distillation. This difference provides a controlled setting for evaluating UNeMo’s adaptability to diverse multimodal inputs. All experiments were conducted on Ubuntu 22.04 using NVIDIA RTX 4090 GPUs (single-GPU configuration), with complete environment specifications to be provided in our open-source code repository later.

## Comparison with State of the Art

**R2R Result** For the R2R dataset experiments, we compared our method with state-of-the-art (SOTA) approaches in LLM-based vision-and-language navigation, including

NavCoT (Lin et al. 2025), LangNav(Pan et al. 2023), Nav-iLLM(Zheng et al. 2024), and NavGPT2(Zhou et al. 2024). These methods formulate navigation as a language modeling problem, employing either prompt engineering or simple imitation learning paradigms for end-to-end navigation decision-making.

As shown in Table 1, the results demonstrate that our method outperforms several competitive methods on both val unseen and test unseen splits. Specifically, on the val unseen split, our approach achieves improvements of 2.1% in SR and SPL compared to baseline methods. For the more challenging test unseen split, these key metrics show gains of 1.5% and 1.3%, respectively. Notably, UNeMo exhibits superior efficiency: with only a FlanT5-1.5B backbone (30% parameters of NavGPT2’s FlanT5-5B), it requires less than half the GPU memory (12GB vs 27GB) while delivering better navigation performance. Furthermore, UNeMo demonstrates particular robustness in long-horizon navigation: on the R2R val-seen split, it exhibits significantly larger SR gains on long paths than on short paths. For long paths with a length of  $\geq 7$ , its SR increases from 64.2% to 69.8%, a 5.6% improvement, whereas for short paths with a length of  $< 7$ , the SR only rises slightly from 71.1% to 72.3%, with a gain of 1.2%. This indicates its enhanced capability in handling complex, long-trajectory navigation scenarios.

These results confirm two core advantages of our method: (1) high efficiency via a compact model that reduces GPU memory footprint while retaining strong performance; (2) superior long-horizon robustness from the multimodal framework’s enhanced reasoning. This provides a practical path for developing efficient, scalable LLM-based navigation systems.

**REVERIE Results** We validate UNeMo’s effectiveness on goal-oriented navigation via the REVERIE dataset, with comparative experimental results presented in Table 2. DUET denotes the original results from the method’s paper, while DUET\* represents our reproduced results—and UNeMo outperforms DUET\* across all metrics. Specifically, UNeMo achieves a navigation SR of 49.36% on Val Unseen and 53.21% on Test Unseen, surpassing both DUET and DUET\*. Its remote RGS reaches 34.28% on Val and 32.15% on Test, outperforming DUET’s 32.15% and 31.88% as well as DUET\*. Notably, UNeMo shows marginal drops in path efficiency metrics SPL and RGSPL. This stems from REVERIE’s natural language instructions providing only coarse-grained object descriptions, unlike R2R’s detailed route guidance, which makes REVERIE’s task emphasize agents’ exploration and path correction capabilities. Unlike DUET’s sole reliance on imitation learning, UNeMo’s hierarchical prediction-feedback mechanism incorporates prospective exploration before navigation decisions, enabling course correction for excessive deviations. This may require extra exploratory actions but ensures target arrival. The trade-off between path efficiency and task success aligns with practical goal-oriented navigation needs, confirming UNeMo prioritizes core navigation requirements—success rates—over marginal efficiency losses.

Methods	Freeze LLM	Val Unseen					Test Unseen				
		TL	NE↓	OSR↑	SR↑	SPL↑	TL	NE↓	OSR↑	SR↑	SPL↑
NavCoT(LLaMa2-7B)	No	9.9	6.3	48	40	37	-	-	-	-	-
LangNav(LLaMa2-7B)	No	-	-	-	46	-	-	-	-	-	-
NaviLLM(Vicuna-7B)	No	12.8	3.5	-	67	59	13.2	3.7	-	68	<u>60</u>
NavGPT2(FlanT5-1.5B)	Yes	12.8	3.3	79	70	59	13.5	3.4	77	<u>71</u>	<u>60</u>
NavGPT2(FlanT5-5B)	Yes	13.6	<u>3.1</u>	<b>81</b>	<u>72</u>	<u>61</u>	14.6	<u>3.4</u>	<u>78</u>	<u>71</u>	<u>60</u>
UNeMo(FlanT5-1.5B)	Yes	13.1	<b>3.0</b>	<u>79.6</u>	<b>72.1</b>	<b>61.1</b>	13.9	<b>3.1</b>	<b>78.4</b>	<b>72.5</b>	<b>61.3</b>

Table 1: Performance comparison with SOTA LLM-based VLN methods on R2R.

Methods	Val Unseen						Test Unseen					
	TL	OSR↑	SR↑	SPL↑	RGS↑	RGSPL↑	TL	OSR↑	SR↑	SPL↑	RGS↑	RGSPL↑
RecBERT	16.78	35.02	30.67	24.9	18.77	15.27	15.86	32.91	29.61	23.99	16.5	13.51
AirBert	18.71	34.51	27.89	21.88	18.23	14.18	17.91	34.2	30.28	23.61	16.83	13.28
HAMT	14.08	36.84	32.95	30.2	18.92	17.28	13.26	33.41	30.40	26.67	14.88	13.08
DUET*	26.17	<u>51.95</u>	46.66	31.03	<u>32.32</u>	21.41	24.25	<u>57.91</u>	<u>52.61</u>	34.82	31.12	20.71
DUET	22.11	51.07	<u>46.98</u>	<b>33.73</b>	32.15	<b>23.03</b>	21.30	56.91	52.51	<b>36.06</b>	<u>31.88</u>	<b>22.06</b>
UNeMo	24.46	<b>53.99</b>	<b>49.36</b>	<u>32.84</u>	<b>34.28</b>	<u>22.90</u>	23.90	<b>58.82</b>	<b>53.21</b>	<u>35.17</u>	<b>32.15</b>	<u>20.90</u>

Table 2: Performance Comparison on REVERIE Dataset. \*indicates reproduced results.

## Ablation Study

This section presents systematic ablation experiments to validate the effectiveness of different state reasoning methods and hierarchical prediction feedback mechanisms.

### Comparison of Different State Reasoning Methods

Since state reasoning’s effectiveness hinges on the chosen approach, we develop several variants to identify the most suitable method. We propose a state deduction framework that enhances multi-modal reasoning through future state prediction, implemented via two paradigms: topology-map-based state prediction and vision-based state decoding.

In the topology-map-based approach, the agent dynamically constructs a graph-structured topological map  $\mathcal{G}_t$  where node features are aggregated via average pooling to form the current state representation  $\mathcal{S}_t$ . This representation is transformed through an MLP, then updated to  $\mathcal{S}_{t+1}$  after expert action execution. As shown in Row 2 (TopoState) of Table 3, this method achieves 2.2% and 1.7% improvements in SR and SPL, respectively. However, it suffers from label acquisition constraints—requiring additional navigation actions and map updates to obtain state labels.

We provide a vision-based alternative: directly predicting candidate view features of expert-action-targeted nodes as future state representations. Row 3 (Cond2Vis) of Table 3 demonstrates 2.4% SR and 2.7% SPL gains on R2R val-unseen data. Notably, predicted states do not directly influence decisions but implicitly optimize node encoding features  $\mathcal{V}_t$  through gradient backpropagation. Both methods significantly improve SR and SPL metrics, proving that this design enables node embedding spaces to autonomously embed latent environmental state dynamics, thereby enhancing navigation performance.

We also conduct comprehensive ablation studies on the

R2R dataset’s val unseen split to evaluate our UNeMo framework’s Multi-modal World Model. As demonstrated in Table 3, the experimental results reveal that utilizing local observations of future states from current nodes as input, coupled with decoding compressed visual features through the joint navigation module’s condition features (Cond2Vis), significantly enhances the agent’s navigation capability (SR: 72.4% vs. baseline 70%). However, the world model-based visual feature decoding approach (Vis-WM), which predicts future states solely from partial observations, shows marginally inferior performance (SR: 71.8%) compared to Cond2Vis. To address this limitation, we incorporate natural language instructions containing action prompts and scene descriptions, implementing cross-attention fusion with local observations of visual features (UNeMo). The final results achieve state-of-the-art (SOTA) performance, exhibiting 2.9% (SR: 72.9%) and 2.7% (SPL: 61.7%) improvements over the baseline.

In conclusion, our ablation experiments substantiate the efficacy of state reasoning methods for LLM-based Vision-and-Language Navigation (VLN) systems. Considering the practical challenges in acquiring topological map state labels, we propose a novel approach that leverages Multi-modal World Model for vision-language data fusion to infer future visual states, while systematically validating the utility of each modular component.

### Ablation Analysis of Hierarchical Prediction-Feedback Navigator

To rigorously validate the pivotal role of the hierarchical prediction-feedback mechanism within the UNeMo framework, this study constructs a dedicated baseline model that explicitly excludes the Hierarchical Prediction-Feedback Navigator (HPFN) for systematic comparative analysis. As demonstrated in Table 4, when the frame-

Methods	Modality		Val Unseen				
	Vision	Language	TL	NE↓	OSR↑	SR↑	SPL↑
NavGPT2	×	×	12.8	3.33	79	70	59
TopoState	×	×	13.2	3.08	80	72.2	60.7
Cond2Vis	✓	×	12.6	3.09	79.7	72.4	61.7
Vis-WM	✓	×	13.4	3.18	<b>80.6</b>	71.8	60.7
<b>UNeMo</b>	✓	✓	13.3	<b>3</b>	<b>80.6</b>	<b>72.9</b>	<b>61.7</b>

Table 3: Comparison of Different State Reasoning Methods. “TopoState” denotes the topology-map-based approach, “Cond2Vis” denotes vision-based approach, “Vis-WM” denotes the vision-based world model, “UNeMo” denotes our multi-modal-based world model.

Methods	Val Unseen				
	TL	NE↓	OSR↑	SR↑	SPL↑
MWM only	13.1	3.3	79.7	71.5	61.0
$a'$ only	13.1	3.0	79.6	72.1	61.1
$a''$ only	13.3	<b>3.0</b>	<b>80.6</b>	<b>72.9</b>	<b>61.7</b>
$a'$ and $a''$	14.1	3.2	80.3	72.1	60.3

Table 4: Ablation Study on Dual-Phase Action Learning policies in Hierarchical Prediction-Feedback Navigator.

work solely employs the MWM as an auxiliary training branch without the integration of HPFN, the navigation metrics—including key indicators such as success rate, path length, and normalized dynamic time warping—on the validation Unseen set exhibit a substantial decline, rendering them significantly inferior to those of the optimized solution that incorporates HPFN. This compelling empirical evidence directly substantiates the critical function of the hierarchical decision-making mechanism in enhancing overall navigation performance, as it bridges the gap between visual state reasoning and actionable navigation policies.

Furthermore, to gain deeper insights into the optimal configuration of HPFN, we designed comparative experiments to evaluate different learning policies for HPFN’s dual-node scoring mechanism, systematically examining three distinct approaches: learning solely from action1 scoring( $a'$  only), learning solely from action2 scoring( $a''$  only), and joint learning from both scoring processes( $a'$  and  $a''$ ). Our experimental results demonstrate that the learning policy focusing on the second scoring output - which integrates state inference results from MWM - optimally balances navigation accuracy and path efficiency in HPFN, thereby establishing the most effective decision-making framework for the UNeMo method.

## Conclusion

This paper addresses the key challenges in Vision-and-Language Navigation (VLN), where existing LLM-based methods are limited by insufficient visual reasoning capabilities and disjoint optimization between reasoning modules and navigation policies. We propose UNeMo, a novel frame-

work that enables collaborative optimization of visual state reasoning and navigational decision-making through its Multimodal World Model (MWM) and hierarchical mechanism. MWM facilitates cross-modal reasoning by jointly predicting subsequent visual states from visual features, language instructions, and navigational actions. The hierarchical interaction between MWM and navigation policies establishes a dynamic bidirectional promotion mechanism, where MWM reasoning enhances policy optimization while policy decisions refine MWM’s reasoning accuracy. Experimental results on R2R and REVERIE datasets demonstrate that UNeMo outperforms state-of-the-art methods, validating its effectiveness in improving VLN performance. This work highlights the potential of integrating multimodal world modeling with hierarchical decision-making mechanism, offering a new direction for developing more robust and adaptive VLN agents. Furthermore, benefiting from Matterport3D’s real-world data, UNeMo narrows the visual gap between simulation and reality. Our future work will extend UNeMo to more complex settings and validate its real-world navigation performance via physical robot deployment.

## Acknowledgements

This work is supported in part by the National Natural Science Funds for Distinguished Young Scholar under Grant 62325307, in part by the National Natural Science Foundation of China under Grants (62403325, 62073225, 62203134, 62373258), in part by the Natural Science Foundation of Guangdong Province under Grants 2023B1515120038, in part by Shenzhen Science and Technology Innovation Commission (20231122104038002, 20220809141216003, KJZD20230923113801004, KJZD20230923115215032, JCYJ20240813141628038), in part by the Guangdong “Pearl River Talent Recruitment Program” under Grant 2019ZT08X603, in part by the Guangdong “Pearl River Talent Plan” under Grant 2019JC01X235, in part by the Scientific Instrument Developing Project of Shenzhen University under Grant 2023YQ019, in part by the Foundation for Distinguished Young Talents in Higher Education of Guangdong under Grant 2025KQNCX112.

## References

- An, D.; Wang, H.; Wang, W.; Wang, Z.; Huang, Y.; He, K.; and Wang, L. 2024. Etpnav: Evolving topological planning for vision-language navigation in continuous environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3674–3683.
- Chang, M.; Gupta, A.; and Gupta, S. 2020. Semantic Visual Navigation by Watching YouTube Videos. *arXiv: Computer Vision and Pattern Recognition, arXiv: Computer Vision and Pattern Recognition*.
- Chen, J.; Lin, B.; Xu, R.; Chai, Z.; Liang, X.; and Wong, K.-Y. K. 2024. Mapgpt: Map-guided prompting with adaptive path planning for vision-and-language navigation. *arXiv preprint arXiv:2401.07314*.
- Chen, S.; Guhur, P.-L.; Schmid, C.; and Laptev, I. 2021. History aware multimodal transformer for vision-and-language navigation. *Advances in neural information processing systems*, 34: 5834–5847.
- 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16537–16547.
- Fried, D.; Hu, R.; Cirik, V.; Rohrbach, A.; Andreas, J.; Morency, L.; Berg-Kirkpatrick, T.; Saenko, K.; Klein, D.; and Darrell, T. 2018. Speaker-Follower Models for Vision-and-Language Navigation. In Bengio, S.; Wallach, H. M.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, 3318–3329.
- Gu, J.; Stefani, E.; Wu, Q.; Thomason, J.; and Wang, X. E. 2022. Vision-and-language navigation: A survey of tasks, methods, and future directions. *arXiv preprint arXiv:2203.12667*.
- Ha, D.; and Schmidhuber, J. 2018. World Models. *Cornell University - arXiv, Cornell University - arXiv*.
- Hong, Y.; Wu, Q.; Qi, Y.; Opazo, C. R.; and Gould, S. 2021a. VLN BERT: A Recurrent Vision-and-Language BERT for Navigation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, 1643–1653. Computer Vision Foundation / IEEE.
- Hong, Y.; Wu, Q.; Qi, Y.; Rodriguez-Opazo, C.; and Gould, S. 2021b. Vln bert: A recurrent vision-and-language bert for navigation. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, 1643–1653.
- Huang, W.; Xia, F.; Xiao, T.; Chan, H.; Liang, J.; Florence, P.; Zeng, A.; Tompson, J.; Mordatch, I.; Chebotar, Y.; et al. 2022. Inner monologue: Embodied reasoning through planning with language models. *arXiv preprint arXiv:2207.05608*.
- Koh, J. Y.; Lee, H.; Yang, Y.; Baldrige, J.; and Anderson, P. 2021. Pathdreamer: A World Model for Indoor Navigation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 14738–14748.
- Krantz, J.; Wijmans, E.; Majumdar, A.; Batra, D.; and Lee, S. 2020a. Beyond the nav-graph: Vision-and-language navigation in continuous environments. In *European Conference on Computer Vision*, 104–120. Springer.
- Krantz, J.; Wijmans, E.; Majumdar, A.; Batra, D.; and Lee, S. 2020b. Beyond the Nav-Graph: Vision-and-Language Navigation in Continuous Environments. In Vedaldi, A.; Bischof, H.; Brox, T.; and Frahm, J., eds., *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XXVIII*, volume 12373 of *Lecture Notes in Computer Science*, 104–120. Springer.
- Li, J.; and Bansal, M. 2023. Improving vision-and-language navigation by generating future-view image semantics. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10803–10812.
- Lin, B.; Nie, Y.; Wei, Z.; Chen, J.; Ma, S.; Han, J.; Xu, H.; Chang, X.; and Liang, X. 2025. Navcot: Boosting llm-based vision-and-language navigation via learning disentangled reasoning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Liu, R.; Wang, X.; Wang, W.; and Yang, Y. 2023. Bird’s-eye-view scene graph for vision-language navigation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10968–10980.
- Pan, B.; Panda, R.; Jin, S.; Feris, R.; Oliva, A.; Isola, P.; and Kim, Y. 2023. Langnav: Language as a perceptual representation for navigation. *arXiv preprint arXiv:2310.07889*.
- Purushwalkam, S.; Gari, S.; Ithapu, V.; Schissler, C.; Robinson, P.; Gupta, A.; and Grauman, K. 2020. Audio-Visual Floorplan Reconstruction. *International Conference on Computer Vision, International Conference on Computer Vision*.
- Qi, Y.; Wu, Q.; Anderson, P.; Wang, X.; Wang, W. Y.; Shen, C.; and Hengel, A. v. d. 2020. Reverie: Remote embodied visual referring expression in real indoor environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9982–9991.
- Qiao, Y.; Qi, Y.; Yu, Z.; Liu, J.; and Wu, Q. 2023. March in chat: Interactive prompting for remote embodied referring expression. In *Proceedings of the IEEE/CVF international conference on computer vision*, 15758–15767.
- Ross, S.; Gordon, G.; and Bagnell, D. 2011. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, 627–635. JMLR Workshop and Conference Proceedings.
- Schumann, R.; Zhu, W.; Feng, W.; Fu, T.-J.; Riezler, S.; and Wang, W. Y. 2024. Velma: Verbalization embodiment of llm agents for vision and language navigation in street view. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 18924–18933.

- Sekar, R.; Rybkin, O.; Daniilidis, K.; Abbeel, P.; Hafner, D.; and Pathak, D. 2020. Planning to Explore via Self-Supervised World Models. *Cornell University - arXiv, Cornell University - arXiv*.
- Shah, D.; Osiniński, B.; Levine, S.; et al. 2023. Lm-nav: Robotic navigation with large pre-trained models of language, vision, and action. In *Conference on robot learning*, 492–504. PMLR.
- Sohn, K.; Lee, H.; and Yan, X. 2015. Learning Structured Output Representation using Deep Conditional Generative Models. In Cortes, C.; Lawrence, N.; Lee, D.; Sugiyama, M.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.
- Song, X.; Chen, W.; Liu, Y.; Chen, W.; Li, G.; and Lin, L. 2025. Towards long-horizon vision-language navigation: Platform, benchmark and method. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 12078–12088.
- Wang, G.; Xie, Y.; Jiang, Y.; Mandlkar, A.; Xiao, C.; Zhu, Y.; Fan, L.; and Anandkumar, A. 2023a. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*.
- Wang, H.; Liang, W.; Van Gool, L.; and Wang, W. 2023b. Dreamwalker: Mental planning for continuous vision-language navigation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 10873–10883.
- Wang, Z.; Li, X.; Yang, J.; Liu, Y.; Hu, J.; Jiang, M.; and Jiang, S. 2024. Lookahead exploration with neural radiance representation for continuous vision-language navigation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 13753–13762.
- Wu, P.; Mu, Y.; Wu, B.; Hou, Y.; Ma, J.; Zhang, S.; and Liu, C. 2024a. Voronav: Voronoi-based zero-shot object navigation with large language model. *arXiv preprint arXiv:2401.02695*.
- Wu, W.; Chang, T.; Li, X.; Yin, Q.; and Hu, Y. 2024b. Vision-language navigation: a survey and taxonomy. *Neural Computing and Applications*, 36(7): 3291–3316.
- Yang, J.; Chen, X.; Qian, S.; Madaan, N.; Iyengar, M.; Fouhey, D. F.; and Chai, J. 2024. Llm-grounder: Open-vocabulary 3d visual grounding with large language model as an agent. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 7694–7701. IEEE.
- Yao, S.; Zhao, J.; Yu, D.; Du, N.; Shafran, I.; Narasimhan, K.; and Cao, Y. 2023. React: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations (ICLR)*.
- Yen-Chen, L.; Bauza, M.; and Isola, P. 2020. Experience-Embedded Visual Foresight. In Kaelbling, L. P.; Kragic, D.; and Sugiura, K., eds., *Proceedings of the Conference on Robot Learning*, volume 100 of *Proceedings of Machine Learning Research*, 1015–1024. PMLR.
- Yue, L.; Zhou, D.; Xie, L.; Zhang, F.; Yan, Y.; and Yin, E. 2024. Safe-vln: Collision avoidance for vision-and-language navigation of autonomous robots operating in continuous environments. *IEEE Robotics and Automation Letters*, 9(6): 4918–4925.
- Zhan, Z.; Yu, L.; Yu, S.; and Tan, G. 2024. Mc-gpt: Empowering vision-and-language navigation with memory map and reasoning chains. *arXiv preprint arXiv:2405.10620*.
- Zhang, X.; Xu, Y.; Li, J.; Hu, Z.; and Hong, R. 2024a. Agent journey beyond rgb: Unveiling hybrid semantic-spatial environmental representations for vision-and-language navigation. *arXiv preprint arXiv:2412.06465*.
- Zhang, Y.; Ma, Z.; Li, J.; Qiao, Y.; Wang, Z.; Chai, J.; Wu, Q.; Bansal, M.; and Kordjamshidi, P. 2024b. Vision-and-language navigation today and tomorrow: A survey in the era of foundation models. *arXiv preprint arXiv:2407.07035*.
- Zheng, D.; Huang, S.; Zhao, L.; Zhong, Y.; and Wang, L. 2024. Towards Learning a Generalist Model for Embodied Navigation. *arXiv:2312.02010*.
- Zhou, G.; Hong, Y.; Wang, Z.; Wang, X. E.; and Wu, Q. 2024. Navgpt-2: Unleashing navigational reasoning capability for large vision-language models. In *European Conference on Computer Vision*, 260–278. Springer.
- Zhou, G.; Hong, Y.; and Wu, Q. 2024. NavGPT: Explicit Reasoning in Vision-and-Language Navigation with Large Language Models. In Wooldridge, M. J.; Dy, J. G.; and Natarajan, S., eds., *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2014, February 20-27, 2024, Vancouver, Canada, 7641–7649*. AAAI Press.