

# ImagiNav: Scalable Embodied Navigation via Generative Visual Prediction and Inverse Dynamics

Jie Chen<sup>1,2</sup>, Yuxin Cai<sup>2,3</sup>, Yizhuo Wang<sup>1</sup>, Ruofei Bai<sup>2,3</sup>, Yuhong Cao<sup>1</sup>,  
Jun Li<sup>2</sup>, Yau Wei Yun<sup>2</sup>, and Guillaume Sartoretti<sup>1</sup>

**Abstract**—Enabling robots to navigate open-world environments via natural language is critical for general-purpose autonomy. Yet, Vision-Language Navigation has relied on end-to-end policies trained on expensive, embodiment-specific robot data. While recent foundation models trained on vast simulation data show promise, the challenge of scaling and generalizing persists due to the limited scene diversity and visual fidelity in simulation. To address this gap, we propose ImagiNav, a novel hierarchical paradigm that formulates navigation in visual space. Instead of predicting actions or metric waypoints, ImagiNav synthesizes a future egocentric video conditioned on language instructions, effectively serving as a high-level visual plan. This plan is subsequently interpreted by an inverse dynamics module, extracting metric trajectories for low-level execution. By decoupling visual planning from robot actuation, this paradigm enables the direct utilization of diverse, in-the-wild human navigation videos. To support this, we develop a scalable auto-labeling data collection pipeline that enhances motion annotation accuracy. Ultimately, ImagiNav demonstrates strong zero-shot transfer to robot navigation without requiring robot demonstrations, paving the way for generalist robots that learn directly from open-world data. The project page is available at: <https://j1dan.github.io/ImagiNav>

## I. INTRODUCTION

The long-standing goal of Embodied Navigation [1] is to develop general-purpose robots capable of following natural language instructions to navigate diverse environments. Mastering this capability is essential for deploying autonomous agents in complex real-world applications. By relying on natural language for intuitive human-robot interaction and egocentric vision for rich semantic understanding, robots can navigate unstructured spaces without depending on precise metric localization, which is notoriously fragile in dynamic environments. To achieve this, Vision-Language Navigation (VLN) agents must possess two critical capabilities: robust high-level reasoning to interpret instructions and precise low-level control to execute actions. While recent advances in end-to-end learning [2]–[4] have shown promise, they predominantly rely on policies trained in simulation. Despite significant efforts to develop continuous and physically realistic simulation platforms [5], [6], a fundamental bottleneck persists: the limited scene diversity and visual fidelity inherent to simulation fail to capture the true complexity of

the real world, severely restricting deployment in open-world scenarios. Conversely, collecting large-scale real-world robot demonstrations is prohibitively expensive. A vastly more abundant resource to overcome this data bottleneck is in-the-wild egocentric navigation video [7], [8]. Whether easily collected or sourced from the web, these videos naturally encode rich navigational intents and capture vast environmental diversity, making them an ideal medium for learning open-world navigation. However, utilizing such data for robotic learning presents a unique challenge: videos lack the action labels required to train conventional policies. To address this, we propose a paradigm shift from *action prediction* to *visual imagination*, treating video as the universal interface between reasoning and control.


In this work, we introduce **ImagiNav**, a scalable generative visual planning framework that formulates navigation directly in visual space. Departing from traditional action-prediction architectures, ImagiNav synthesizes the future visual observations of the navigation process. This generated video effectively serves as a predictive visual plan, from which an inverse dynamics module later extracts the pose trajectories to be tracked by a low-level controller. The key insight is that while robot kinematics vary, the visual outcome of what the agent expects to see while progressing toward a goal is largely embodiment-invariant. By treating videos as plans, we leverage the rich semantic and physical priors inherent in large-scale pretrained video models, while remaining independent of specific control parameterizations. Crucially, this formulation unlocks the vast potential of unlabeled, in-the-wild egocentric videos for scalable robot learning.

To initiate this process, we utilize a pretrained Vision-Language Model (VLM) as a semantic reasoner to decompose the high-level task into textual subgoals. Conditioned on the subgoal and the current egocentric observation, a generative video model synthesizes the visual plan. By finetuning this model on a dataset of handheld smartphone videos captured in everyday environments, we enable the model to effectively ground the reasoner’s abstract intent into physically plausible visual dynamics. To mitigate spatial ambiguity (e.g., confusing “left” with “right”), a common failure mode in video models [9]–[11], we implement an Action-Conditioned Mixture-of-Experts (AC-MoE) strategy that routes the generation task to specialized experts, yielding precise motion dynamics. To translate imagined visual trajectories into executable actions, ImagiNav incorporates an inverse dynamics module that serves as a geometric

<sup>1</sup>Jie Chen, Yizhuo Wang, Yuhong Cao, and Guillaume Sartoretti are with the Department of Mechanical Engineering, National University of Singapore, Singapore.

<sup>2</sup>Jie Chen, Yuxin Cai, Ruofei Bai, Jun Li, and Yau Wei Yun are with the Institute for Infocomm Research (I<sup>2</sup>R), Agency for Science, Technology and Research (A\*STAR), Singapore.

<sup>3</sup>Yuxin Cai and Ruofei Bai are with Nanyang Technological University, Singapore.



**Fig. 1: Overview of the ImagiNav Framework.** Top: Navigation is formulated in visual space. A generative visual planner synthesizes future egocentric observations, which are geometrically grounded into pose trajectories via inverse dynamics and executed by a tracking controller. Bottom: Geometry-first scalable training pipeline from in-the-wild egocentric videos.

grounding layer. This module decodes relative pose changes directly from the synthesized video into metric waypoints. Operating within a temporal hierarchy, the generative model acts as a low-frequency visual planner, while a standard high-frequency controller simply tracks the decoded waypoints. This visual-to-geometric interface provides an embodiment-agnostic planning-control abstraction while rendering the agent’s intent inherently interpretable.

Building on this representation, we further develop a geometry-first data collection pipeline that enables training from in-the-wild human egocentric navigation videos without requiring precise state estimation. By extracting motion primitives through geometric inverse dynamics prior to semantic annotation, we mitigate spatial hallucination errors common in purely VLM labeling. This paradigm eliminates the need for robot teleoperation, precise localization systems, and metric calibration during data collection, making the approach inherently scalable. Our results demonstrate that ImagiNav successfully transfers navigation priors learned from in-the-wild human navigation videos to robots. This successful zero-shot transfer, combined with video generation evaluations confirming the superior out-of-distribution generalization of human data over simulation, underscores the potential of massive egocentric navigation videos as a scalable pathway to generalist embodied intelligence that effectively bypasses the bottleneck of expensive, platform-specific robot demonstrations.

## II. RELATED WORK

### A. Vision-Language Navigation

Vision-Language Navigation (VLN) tasks agents with following natural language instructions in unseen 3D envi-

ronments [1]. Early works primarily relied on reconstructed indoor scenes [12], [13] or synthetic environments [14]–[17] within simulators [5], [18], often simplifying navigation into discrete graph traversal or assuming idealized robot dynamics. Recent advances have incorporated large pretrained language models and transformer-based architectures to enhance high-level reasoning [2]–[4], [19]. While effective in controlled settings, these methods rely heavily on robot data, predominantly from simulation, leading to generalization gaps in diverse real-world environments.

Efforts to reduce domain gaps include more realistic simulation platforms such as VLN-PE [6], [20] that introduce realistic challenges like physical shaking. Other approaches focus on leveraging web-scale video data [7], [8]. However, these approaches typically require either precise state estimation for action labeling or retain action-space supervision, limiting their scalability. In contrast, we reformulate navigation planning in visual space, separating semantic intent from control parameterization.

### B. Generative Models

Generative modeling has revolutionized content creation, expanding from static image synthesis [21], [22] and dynamic video generation [23]–[25] to complex robotic planning. In the domain of robotics, diffusion models effectively represent multi-modal action distributions [26]–[28], while Vision-Language-Action models [29], [30] advance this by leveraging large backbones to inherit rich semantic priors for action diffusion.

Beyond direct policy learning, video generation has emerged as a world model, capable of simulating interactive environments [31], [32] or serving as a trajectory scorer for navigation agents [33], [34]. Distinct from methods that use

these models primarily for simulation, we directly leverage video generation as a policy, functioning as a language-conditioned world model that imagines the future trajectory.

### C. Inverse Dynamics from Visual Observations

Visual Odometry (VO) and Structure-from-Motion (SfM) techniques have long been used for estimating camera motion from image sequences [35], [36]. Recent geometric foundation models further scale pose estimation by leveraging large video corpora and transformer architectures [37], [38], enabling robust relative pose prediction in diverse environments. Rather than using geometric estimation purely for localization, we reinterpret inverse dynamics as a geometric grounding layer that translates visual trajectories into physically executable motion commands. This design enables us to bridge generative visual planning with embodiment-agnostic control without requiring explicit action supervision.

## III. IMAGINAV

The ImagiNav framework is structured as a hierarchical planning architecture comprising three interconnected modules: Semantic Reasoning, Visual Imagination, and Geometry Inverse Dynamics. The overall architecture is illustrated in Figure 1.

### A. Semantic Reasoning

We leverage a pretrained VLM to act as the hierarchical task decomposer. Operating on the current egocentric frame and the global navigation instruction, this module formulates the navigational intent into a natural language prompt. This textual output acts as the semantic condition for the video generation module, translating high-level reasoning into a specific description of the expected future visual trajectory.

### B. Visual Imagination

The core planning component is a generative video model  $\mathcal{G}$  that operates in visual space. Given the current observation  $I_t$  and a textual instruction  $L_t$ , the model predicts a sequence of future egocentric frames:

$$V_{pred} = \mathcal{G}(I_t, L_t), \quad (1)$$

where  $V_{pred} = \{I_{t+1}, \dots, I_{t+H}\}$  represents an imagined future visual trajectory over a horizon of  $H$  steps. Unlike conventional policies that output control signals or waypoints,  $V_{pred}$  encodes the navigation plan implicitly through expected visual progression.

Specifically, we instantiate  $\mathcal{G}$  using the pre-trained LTX-2B architecture, which employs a Diffusion Transformer backbone [39] operating within a compressed spatiotemporal latent space. The generation process follows a rectified flow matching objective [40], [41], where a latent representation  $z_\tau$  is evolved from a noise prior  $z_1 \sim \mathcal{N}(0, \mathbf{I})$  to a clean latent  $z_0$ :

$$\frac{dz_\tau}{d\tau} = v_\theta(z_\tau, \tau, c), \quad (2)$$

where  $v_\theta$  denotes the transformer-based velocity field prediction parameterized by  $\theta$ , conditioned on the encoded

start frame and text instruction  $c = \{\mathcal{E}(I_t), \mathcal{E}(L_t)\}$ . The final video is reconstructed via a causal Video-VAE decoder that performs the final denoising step during latent-to-pixel decoding. To ensure domain relevance, we finetune  $\theta$  on our scalable, human-collected video dataset (detailed in Section IV), aligning the generative flow with physically plausible ego-motion dynamics. To implement the AC-MoE strategy for mitigating the spatial ambiguity common in video generation, we bifurcate our training data to finetune two specialized expert models. Both experts handle forward motions but specialize in left and right turns, respectively. During inference, the semantic reasoner dynamically routes the generation task to the appropriate expert based on the intended subgoal.

The generative model grounds high-level textual subgoals into synthesized visual trajectories, explicitly encoding the robot’s navigational intent and affording highly modular integration into downstream control pipelines. For instance, key frames could be extracted to serve as target observations for a learned image-goal policy. In our proposed architecture, we bypass intermediate visual policies by directly decoding the trajectory from the synthesized video via an IDM.

### C. Geometry Inverse Dynamics

To bridge visual planning and physical execution, we introduce a geometric grounding layer to extract the trajectory from the video plan, implemented via an Inverse Dynamics Model (IDM) denoted as  $\Phi$ :

$$\mathbf{T}_{plan} = \{(\Delta x_i, \Delta y_i, \Delta \theta_i)\}_{i=1}^N = \Phi(V_{pred}). \quad (3)$$

The IDM decodes the synthesized video  $V_{pred}$  into a sequence of  $N$  relative ego-motion waypoints  $\mathbf{T}_{plan}$ , which is passed to a low-level kinematic controller, converting the spatial displacements into executable motor commands for the robot’s hardware interface.

## IV. SCALABLE DATA COLLECTION PIPELINE

A key advantage of planning in visual space is that supervision no longer requires embodiment-specific action labels or pose annotations. Because the planner operates purely in visual space, learning reduces to modeling physically plausible visual progression conditioned on text. This observation enables a shift from metric-annotated data collection to egocentric video curation.

Instead of collecting expensive robot demonstrations, we leverage in-the-wild human walking videos captured using commodity cameras (e.g., smartphones and GoPros). These videos naturally contain rich ego-motion dynamics and diverse environmental layouts, without requiring robotic hardware, localization systems, or camera calibration.

### A. Auto-Labeling via IDM-Guided Annotation

Raw human videos lack textual navigation instructions. Directly relying on VLMs for annotation introduces spatial hallucination errors, particularly in directional reasoning (e.g., confusing left and right turns). To address this, we adopt a *geometry-first, semantics-second* supervision

pipeline. Geometric motion primitives are extracted prior to semantic labeling, ensuring that textual descriptions remain physically grounded.

1) *Geometric Motion Primitive Extraction*: We first segment raw videos into fixed-length clips and apply a pre-trained IDM to recover relative ego-motion trajectories. Based on these poses, a heuristic classifier then categorizes the motion patterns into distinct primitives, including forward translation, rotational turns, and compound movements. This step ensures motion labels are grounded in physical geometry rather than visual interpretation. Importantly, this step does not require explicit camera calibration or robot-specific kinematic modeling. All motion signals are derived directly from image-space geometry.

2) *Motion-Primitive-Guided Captioning*: Given the extracted motion primitives, we prompt a pretrained large language model to generate dense scene descriptions that are consistent with the recovered motion. The geometric motion prior constrains the semantic annotation process, mitigating spatial hallucination errors. To better align with the distribution of video model pretraining data, we map navigation verbs to camera movement terminology (e.g., replacing “move” with “dolly” and “turn” with “pan”). This results in a scalable, low-cost dataset  $\mathcal{D} = \{(V_{clip}, T_{desc})\}$  where the visual trajectory is physically consistent with the text, enabling the video model to learn robust navigation dynamics without direct robot supervision.

## V. EXPERIMENTS

We evaluate ImagiNav from two complementary perspectives. First, we assess whether the visual planner can leverage its learned priors to achieve zero-shot generalization from real-world human videos to unseen simulated environments, driving a physically simulated robot in a high-fidelity embodied benchmark. Second, we study how the source of finetuning data (simulation vs. human videos) affects the video planner’s visual quality and kinematic realism.

### A. Simulation Experimental Setup

We evaluate ImagiNav in the VLN-PE benchmark [6], [20], built on NVIDIA Isaac Sim [42], providing high-fidelity rendering and physically realistic dynamics. All methods are evaluated on the same simulated Unitree H1 humanoid embodiment. The low-level locomotion policy is trained via reinforcement learning and takes linear and angular velocity commands to execute bipedal walking. We report VLN-PE standard navigation and stability metrics on 100 episodes from the InteriorNav dataset, which utilizes the high-quality InteriorAgent scenes [16].

### B. Implementation Details

**Real-World Data Collection.** To validate the proposed pipeline, we collect a pilot-scale dataset using a hand-held iPhone in diverse environments (e.g., shopping malls and university campuses). Video clips are sampled from the raw videos of approximately 33 minutes, and labeled using the IDM-guided procedure in Section IV-A. To reduce class

imbalance, we apply geometric mirroring augmentation, yielding a more balanced coverage of motion primitives (Fig. 2), resulting in 2,558 clips of 121 frames each.

**System Instantiation.** We instantiate the abstract modules of ImagiNav with foundation models:

- **Reasoner**: We use Gemini 3.0 Flash [43] for high-level task decomposition.
- **Video Generation Model ( $\mathcal{G}$ )**: We employ the LTX-Video-2B architecture [24], a latent diffusion model parameterized by a Diffusion Transformer backbone. To adapt the model to the domain of ego-centric navigation, we finetune the pretrained model on our collected dataset using Low-Rank Adaptation (LoRA) [44]. The model is configured to synthesize video sequences with a spatial resolution of  $480 \times 256$  (width  $\times$  height) and a temporal depth of 121 frames, operating at a frame rate of 24 Hz.
- **Inverse Dynamics Model ( $\Phi$ )**: We employ VGGT [38], a geometric foundation model, to extract relative camera poses from the generated videos, where the frames are sampled at an interval.

During inference, we utilize a 70-frame functional video generation horizon, where the IDM operates on frames subsampled at a stride of 4.

**Low-Level Controller.** While ImagiNav is compatible with various path-tracking algorithms, we employ a simple heuristic proportional controller for the experiments. To map the IDM-predicted displacements  $(\Delta x, \Delta y, \Delta \theta)$  to the robot’s velocity command  $(v, \omega)$ , we use:

$$v = \begin{cases} \text{clip}(k_v \sqrt{\Delta x^2 + \Delta y^2}, 0, v_{max}) & \text{if } |\Delta \theta| \leq \frac{\pi}{4} \\ 0 & \text{if } |\Delta \theta| > \frac{\pi}{4} \end{cases}, \quad (4)$$

$$\omega = \text{clip}(k_\theta \cdot \text{wrap}(\Delta \theta), -\omega_{max}, \omega_{max}), \quad (5)$$

where  $k_v$  and  $k_\theta$  are proportional gains.


### C. Baselines

To rigorously assess ImagiNav’s performance, we benchmark it against zero-shot modular baselines, state-of-the-art domain-trained policies, and a simulation-trained variant of our own method to isolate data efficacy.


- **NavDP [28]**: We compare against NavDP, a state-of-the-art diffusion-based policy trained on massive simulation data. It processes RGB-D observations alongside a goal coordinate  $(x_{goal}, y_{goal})$  to output a trajectory of future waypoints  $\mathbf{T} = \{(\Delta x_i, \Delta y_i, \Delta \theta_i)\}_{i=1}^H$ :

$$\mathbf{T} = \pi_{\text{NavDP}}(I_{\text{RGB}}, I_{\text{D}}, x_{\text{goal}}, y_{\text{goal}}). \quad (6)$$

- **InternVLA-N1**: We evaluate InternVLA-N1, the first open dual-system navigation foundation model. It employs a finetuned 7B VLM as a system-2 planner and a diffusion-based navigation policy as the system-1 local planner. Trained on the massive InternData-N1 dataset, it achieves state-of-the-art performance on various benchmarks.
- **ImagiNav-Sim**: To investigate the role of real-world video data versus simulation data, we finetune this



(a) Scene Distribution



(b) Motion Primitive Distribution

**Fig. 2: Dataset Statistics.** (a) Distribution of raw videos across different real-world environments. (b) Distribution of motion primitives in the final augmented dataset, showing balanced coverage.

variant exclusively on simulation data from a subset of the InternData-N1 dataset. We carefully curate samples from *3D-Front*, *R2R*, and *HSSD* scenes to strictly match the total size and motion primitive distribution of the real-world dataset.

For all zero-shot evaluations, we adopt an open-loop local execution strategy: the agent synthesizes a trajectory based on the current observation and executes the derived sequence of motor commands. The robot acquires a fresh observation and initiates re-planning after the current action buffer is fully consumed.


#### D. Simulation Results

We employ a suite of metrics to evaluate navigation performance. We report **Success Rate (SR)** for stops within a goal threshold, and Success weighted by Path Length (**SPL**), measuring efficiency relative to the optimal path, defined as  $SPL = \frac{1}{N} \sum_{i=1}^N S_i \frac{\ell_i}{\max(p_i, \ell_i)}$ , where  $S_i$  is a binary success indicator,  $\ell_i$  is the shortest path distance, and  $p_i$  is the agent’s path length. We also report **Oracle Success Rate (OS)**, which assesses if any point on the path fell within that threshold, **Navigation Error (NE)** (distance to goal in meters), and **Trajectory Length (TL)**.

Our quantitative results are presented in Table I. As expected, InternVLA-N1, benefiting from massive in-domain supervision ( $\sim 1000$  hours), establishes the performance upper bound. In the zero-shot transfer regime, the results highlight distinct performance patterns:

**1. Modality Gap in Modular Baselines.** Although NavDP is a state-of-the-art visual navigation policy, it slightly underperforms ImagiNav in the VLN setting where the point goal is given by the semantic reasoner. We attribute this performance gap to the inherent deficiency of VLMs in grounding semantic understanding into precise metric coordinates. In the NavDP pipeline, the VLM is tasked with abstracting semantic subgoals into geometric coordinates, which directly exposes the language model’s lack of fine-grained spatial reasoning capabilities. In contrast, ImagiNav bridges this gap via visual imagination: by generating the future video explicitly, it grounds the VLM’s textual reasoning into a dense, pixel-perfect trajectory.

**2. Strong Zero-Shot Generalization.** Despite being



**Fig. 3: Qualitative Visualization of the Imagination-Action Cycle.** (Left) The *imagined* future frame generated by the model. (Middle) The actual robot observation after executing the extracted plan. (Right) Top-down trajectory map, where the **red** line indicates the executed path and the **green** line represents the reference ground truth. The close alignment confirms the physical consistency of the generated plans.

trained on out-of-domain human videos with negligible exposure to residential indoor scenes, **ImagiNav-Real** demonstrates robust transferability to the VLN-PE benchmark. This confirms that our planner leverages strong semantic priors from the pretrained video foundation model to generalize the concept of navigable space without overfitting to specific textures. In particular, Figure 3 presents a qualitative visualization of the imagination-action cycle. The high visual correspondence between the imagined and actual frames demonstrates the model’s ability to synthesize physically feasible plans that are realizable by the low-level controller.

As expected, **ImagiNav-Sim**, benefiting from direct exposure to in-domain indoor layouts, slightly outperforms ImagiNav-Real. However, the performance gap is remarkably

**TABLE I:** Quantitative Results on InteriorNav Dataset. We report performance relative to input modality (Depth), domain alignment (Sim→Sim), and data scale.

| Method                      | Depth | Sim → Sim | Data Size | TL ↓        | NE ↓        | OS ↑        | SR ↑        | SPL ↑       |
|-----------------------------|-------|-----------|-----------|-------------|-------------|-------------|-------------|-------------|
| <i>In-Domain Supervised</i> |       |           |           |             |             |             |             |             |
| InternVLA-N1 [4]            | ✓     | ✓         | ~1000h    | <b>6.42</b> | <b>3.48</b> | <b>0.63</b> | <b>0.53</b> | <b>0.44</b> |
| <i>Zero-Shot Transfer</i>   |       |           |           |             |             |             |             |             |
| NavDP [28]                  | ✓     | ✓         | ~101h     | <b>3.48</b> | 4.19        | 0.37        | 0.32        | 0.31        |
| ImagiNav-Sim                | ×     | ✓         | ~0.5h     | 3.92        | <b>3.94</b> | <b>0.45</b> | <b>0.39</b> | <b>0.37</b> |
| ImagiNav-Real               | ×     | ×         | ~0.5h     | 4.03        | 4.13        | 0.41        | 0.36        | 0.35        |

*Note:* TL: Trajectory Length, NE: Navigation Error (m), OS: Oracle Success Rate, SR: Success Rate, SPL: Success weighted by Path Length. Best results for each category are **bolded**.

**TABLE II:** Quantitative Evaluation of LTX-2B Video Model with Sim/Real Data

| Model Variant             | FVD ↓        | LPIPS ↓      | PSNR ↑       | SSIM ↑      | Mot. Fid. ↑ | RPE-T (m) ↓  | RPE-R (°) ↓ |
|---------------------------|--------------|--------------|--------------|-------------|-------------|--------------|-------------|
| Base LTX-2B (Zero-shot)   | 391.14       | 0.520        | 12.61        | 0.37        | 0.40        | 0.036        | 1.39        |
| Real-Finetuned w/o AC-MoE | 73.08        | 0.486        | <b>13.36</b> | <b>0.40</b> | 0.60        | 0.028        | 1.31        |
| Sim-Finetuned             | 72.72        | 0.514        | 13.03        | 0.39        | 0.67        | 0.051        | 1.73        |
| Real-Finetuned            | <b>65.39</b> | <b>0.481</b> | 13.27        | <b>0.40</b> | <b>0.72</b> | <b>0.024</b> | <b>1.18</b> |

*Note:* Mot. Fid.: Motion Fidelity, RPE-T: Translational Relative Pose Error, RPE-R: Rotational Relative Pose Error. Best results are **bolded**.

narrow. We attribute this to the kinematic artifacts inherent to simulation datasets. Trajectories generated by deterministic planners like A\* on discrete graphs often exhibit unnatural, oscillatory zig-zag patterns and instantaneous directional changes, while our real-world dataset possesses natural, continuous momentum. This suggests that, while in-domain simulation data provides better visual layout priors, real-world human data offers superior motion dynamics, partially offsetting this domain mismatch.

### E. Evaluation of Video Generation

We assess the generalization capabilities of our video generation model using a real-world test set, distinct from the training distribution. This dataset encompasses diverse environmental scenes and varying lighting conditions, constituting a rigorous out-of-distribution evaluation.

The evaluation methodology mirrors the setup described in Section V. Our test corpus comprises 710 video clips, each containing 121 frames, extracted from approximately 12.3 minutes of raw videos. We evaluate the synthesized videos through both comprehensive quantitative metrics and qualitative visual analysis.

1) *Quantitative Results:* To quantify visual quality, we report standard video generation metrics: Fréchet Video Distance (FVD) [45], measuring distributional divergence of spatiotemporal dynamics, Learned Perceptual Image Patch Similarity (LPIPS) for perceptual consistency, and Peak Signal-to-Noise Ratio (PSNR) alongside Structural Similarity Index Measure (SSIM) for low-level reconstruction fidelity.

While the aforementioned metrics assess visual realism, they do not explicitly measure the kinematic accuracy required for downstream navigation tasks. For our visual planner, it is imperative that the synthesized video accurately reflects the intended motion dynamics rather than just visual

plausibility. Consequently, we introduce a specific evaluation of kinematic fidelity based on the Relative Pose Error (RPE) and *Motion Fidelity*. The ground-truth camera extrinsics are estimated via VGGT. Let  $\mathbf{p}_{ref}^{(t)} = (x_{ref}^{(t)}, y_{ref}^{(t)}, \theta_{ref}^{(t)})$  and  $\mathbf{p}_{gen}^{(t)} = (x_{gen}^{(t)}, y_{gen}^{(t)}, \theta_{gen}^{(t)})$  denote the reference and generated poses at time step  $t$ , respectively. We define the Translational RPE ( $RPE_{trans}$ ) and Rotational RPE ( $RPE_{rot}$ ) as the mean error over the trajectory of length  $N$ :

$$RPE_{trans} = \frac{1}{N} \sum_{t=1}^N \sqrt{(x_{ref}^{(t)} - x_{gen}^{(t)})^2 + (y_{ref}^{(t)} - y_{gen}^{(t)})^2}, \quad (7)$$


$$RPE_{rot} = \frac{1}{N} \sum_{t=1}^N \left| \text{wrap}(\theta_{ref}^{(t)} - \theta_{gen}^{(t)}) \right|, \quad (8)$$

where  $\text{wrap}(\cdot)$  normalizes the angular difference to the range  $[-180^\circ, 180^\circ]$ .  $RPE_{trans}$  measures the Euclidean distance error in meters, while  $RPE_{rot}$  captures the absolute yaw deviation in degrees.


While RPE captures precise geometric deviations, it is sensitive to domain-specific kinematic shifts (e.g., velocity mismatches between simulation and reality), often penalizing semantically correct actions. To evaluate high-level behavioral alignment, we introduce Motion Fidelity, in which we discretize the continuous relative trajectories into a sequence of semantic action primitives  $L = \{a_t\}_{t=1}^N$ . Each action  $a_t$  is assigned to one of four categories based on empirically determined thresholds: Forward (F), Turn Left (TL), Turn Right (TR), or Static (S). Motion Fidelity quantifies the alignment between the reference sequence  $L_{ref}$  and the generated sequence  $L_{gen}$  using the normalized Levenshtein distance:

$$\text{Motion Fidelity} = 1 - \frac{\text{Levenshtein}(L_{ref}, L_{gen})}{\max(|L_{ref}|, |L_{gen}|)}, \quad (9)$$

where  $\text{Levenshtein}(\cdot)$  computes the minimum edit distance. Because  $L_{ref}$  and  $L_{gen}$  are of equal length in our formulation, this distance simplifies to point-wise semantic substitutions.




(a) **Scene A: Navigable Affordance Perception.** Generation instruction: *Move forward along the walkway towards the building.*



(b) **Scene B: Dynamic Actor Modeling.** Generation instruction: *Move forward and turn right towards the silver trailer.*

**Fig. 4: Qualitative Comparison of Simulation vs. Real-World Training.** We visualize the generated future trajectories across Ground Truth (Top Row), ImagiNav-Real (Middle Row), and ImagiNav-Sim (Bottom Row). Real-world data demonstrates superior understanding of static geometry (Scene A, the navigable walkway is circled in green), and dynamic behavior (Scene B).



**Fig. 5: Controllability and Geometric Grounding.** Given a single start frame, the model generates diverse trajectories conditioned on different language instructions (from left to right: “Move forward and turn left”, “Move forward”, and “Move forward and turn right”).

Quantitative results (Table II) reveal a distinct performance hierarchy. Because our objective is reliable visual planning rather than mere frame-level synthesis, we prioritize spatiotemporal consistency (FVD) and behavioral alignment (Motion Fidelity) as the primary indicators of success. The Base LTX-2B exhibits poor performance across all metrics, confirming that generic video models lack the fine-grained motion priors required for navigation. Our AC-MoE strategy efficiently establishes these priors, outperforming a single-model baseline in FVD and Motion Fidelity. While Sim-Finetuning improves on *frame-level visual metrics* (PSNR, SSIM, LPIPS), it lags significantly in *spatiotemporal consistency* (FVD 72.72 vs. 65.39) and *kinematic realism* (Motion Fidelity 0.67 vs. 0.72). This gap reflects the dual limitations of simulation: the inherent lack of environmental diversity and the presence of synthetic kinematic artifacts. In contrast, the Real-Finetuned model dominates across all categories. By leveraging diverse human data, it captures natural kinematic flow and generalizes effectively to unseen environments.

2) *Qualitative Results:* We provide a qualitative comparison between models trained on simulation versus real-world data in Figure 4. As shown in the bottom row (Scene A), the ImagiNav-Sim model exhibits brittle behavior in unstructured environments, hallucinating unsafe trajectories that collide with the guardrail. In contrast, ImagiNav-Real (middle row) accurately perceives the walkable affordance,

generating a safe path that strictly adheres to the walkway.

Furthermore, Scene B highlights the critical role of real-world data in modeling dynamic scenes. While the simulation-trained model (bottom row) treats pedestrians as static objects due to the lack of dynamic actors in its training distribution, ImagiNav-Real (middle row) successfully anticipates human motion, generating physically consistent walking trajectories for the pedestrians in the scene.

Figure 5 demonstrates the controllability and geometric grounding of our real-finetuned generation model. Given an initial observation, the model synthesizes diverse future trajectories aligned with distinct natural language instructions (“Move forward”, “Move forward and turn left/right”). Notably, the generation is not merely a rigid execution of motion primitives; it exhibits context-aware adaptation. For instance, under the “Move forward and turn left” instruction, the model first steers slightly right to navigate around a stone bench before executing the turn. This demonstrates that the model does not blindly follow the text, but rather grounds the instruction in the physical geometry of the scene for collision avoidance.

## VI. CONCLUSION

In this work, we introduced **ImagiNav**, a modular and hierarchical framework that redefines VLN by treating video generation as an embodiment-agnostic planner. By formulating navigation as a video prediction task, our approach leverages the rich semantic priors inherent in large-scale video models while enabling a scalable data pipeline that transforms readily available in-the-wild videos into actionable navigation policies. Our analysis demonstrates that the superior environmental diversity and kinematic naturalness of in-the-wild human data enable robust zero-shot generalization to physical robot navigation, paving the way for generalist robots that learn directly from human visual experience.

**Limitations and Future Work.** The primary limitation of our pipeline is inference latency, where the high computational cost of video generation currently restricts the system to a low-frequency control regime, limiting reactivity to sudden dynamic hazards. Additionally, operating purely in RGB space can lead to geometric hallucinations, where the model synthesizes visually plausible but physically unsafe trajectories; these may be mitigated through a visual plan

scorer or traditional safety fallbacks. Finally, while our AC-MoE strategy successfully mitigates spatial ambiguity, a notoriously difficult and persistent challenge for generative video models, it remains a preliminary solution that introduces additional overhead by requiring the management of multiple experts. Future work will focus on: (1) distilling the heavy generative model into a lightweight, real-time policy to enable high-frequency closed-loop control; (2) integrating explicit depth modalities into the generation process to enforce stricter geometric consistency and safety in complex 3D environments; and (3) overcoming the fundamental limits of spatial reasoning by distillation of specialized experts into a single, unified model that maintains directional accuracy.

## ACKNOWLEDGMENT

The authors acknowledge the use of AI tools (Gemini) for linguistic refinement and assistance with figure and table formatting. All suggestions were refined and verified by the authors, who take full responsibility for the paper’s content, data, and conclusions.

## REFERENCES

- [1] P. Anderson *et al.*, “Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 3674–3683.
- [2] J. Zhang *et al.*, “Uni-navid: A video-based vision-language-action model for unifying embodied navigation tasks,” *arXiv preprint arXiv:2412.06224*, 2024.
- [3] M. Wei *et al.*, “Streamvln: Streaming vision-and-language navigation via slowfast context modeling,” *arXiv preprint arXiv:2507.05240*, 2025.
- [4] M. Wei, C. Wan, J. Peng, X. Yu *et al.*, “Ground slow, move fast: A dual-system foundation model for generalizable vision-and-language navigation,” *arXiv preprint arXiv:2512.08186*, 2025.
- [5] J. Krantz, E. Wijmans, A. Majumdar, D. Batra, and S. Lee, “Beyond the nav-graph: Vision-and-language navigation in continuous environments,” in *European Conference on Computer Vision*. Springer, 2020, pp. 104–120.
- [6] L. Wang *et al.*, “Rethinking the embodied gap in vision-and-language navigation: A holistic study of physical and visual disparities,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025, pp. 9455–9465.
- [7] K. Lin, P. Chen, D. Huang, T. H. Li, M. Tan, and C. Gan, “Learning vision-and-language navigation from youtube videos,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 8317–8326.
- [8] X. Liu *et al.*, “Citywalker: Learning embodied urban navigation from web-scale videos,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 6875–6885.
- [9] Z. Huang *et al.*, “Vbench: Comprehensive benchmark suite for video generative models,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 21 807–21 818.
- [10] D. Zheng *et al.*, “Vbench-2.0: Advancing video generation benchmark suite for intrinsic faithfulness,” *arXiv preprint arXiv:2503.21755*, 2025.
- [11] Y. Luo *et al.*, “V-reasonbench: Toward unified reasoning benchmark suite for video generation models,” *arXiv preprint arXiv:2511.16668*, 2025.
- [12] A. Chang *et al.*, “Matterport3d: Learning from rgb-d data in indoor environments,” *arXiv preprint arXiv:1709.06158*, 2017.
- [13] S. K. Ramakrishnan *et al.*, “Habitat-matterport 3d dataset (hm3d): 1000 large-scale 3d environments for embodied ai,” *arXiv preprint arXiv:2109.08238*, 2021.
- [14] H. Fu *et al.*, “3d-front: 3d furnished rooms with layouts and semantics,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 10 933–10 942.
- [15] M. Khanna *et al.*, “Habitat synthetic scenes dataset (hssd-200): An analysis of 3d scene scale and realism tradeoffs for objectgoal navigation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 16 384–16 393.
- [16] M. T. I. SpatialVerse Research Team, “Interioragent: Interactive usd interior scenes for isaac sim-based simulation,” <https://huggingface.co/datasets/spatialverse/InteriorAgent>, 2025.
- [17] M. Deitke *et al.*, “Proctor: Large-scale embodied ai using procedural generation,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 5982–5994, 2022.
- [18] M. Savva *et al.*, “Habitat: A platform for embodied ai research,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 9339–9347.
- [19] Y. Cai, X. He, M. Wang, H. Guo, W.-Y. Yau, and C. Lv, “Cl-cotnav: Closed-loop hierarchical chain-of-thought for zero-shot object-goal navigation with vision-language models,” *arXiv preprint arXiv:2504.09000*, 2025.
- [20] I. Contributors, “InternNav: InternRobotics’ open platform for building generalized navigation foundation models,” <https://github.com/InternRobotics/InternNav>, 2025.
- [21] J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.
- [22] C. Saharia *et al.*, “Photorealistic text-to-image diffusion models with deep language understanding,” *Advances in neural information processing systems*, vol. 35, pp. 36 479–36 494, 2022.
- [23] A. Blattmann *et al.*, “Stable video diffusion: Scaling latent video diffusion models to large datasets,” *arXiv preprint arXiv:2311.15127*, 2023.
- [24] Y. HaCohen *et al.*, “Ltx-video: Realtime video latent diffusion,” *arXiv preprint arXiv:2501.00103*, 2024.
- [25] T. Wan *et al.*, “Wan: Open and advanced large-scale video generative models,” *arXiv preprint arXiv:2503.20314*, 2025.
- [26] C. Chi *et al.*, “Diffusion policy: Visuomotor policy learning via action diffusion,” *The International Journal of Robotics Research*, vol. 44, no. 10–11, pp. 1684–1704, 2025.
- [27] A. Sridhar, D. Shah, C. Glossop, and S. Levine, “Nomad: Goal masked diffusion policies for navigation and exploration,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 63–70.
- [28] W. Cai *et al.*, “Navdp: Learning sim-to-real navigation diffusion policy with privileged information guidance,” *arXiv preprint arXiv:2505.08712*, 2025.
- [29] K. Black *et al.*, “ $\pi_0$ : A vision-language-action flow model for general robot control,” *arXiv preprint arXiv:2410.24164*, 2024.
- [30] P. Intelligence *et al.*, “ $\pi_{0.5}$ : a vision-language-action model with open-world generalization,” *arXiv preprint arXiv:2504.16054*, 2025.
- [31] J. Parker-Holder *et al.*, “Genie 2: A large-scale foundation world model,” URL: <https://deepmind.google/discover/blog/genie-2-a-large-scale-foundation-world-model>, vol. 2, 2024.
- [32] M. Yang, Y. Du, K. Ghasemipour, J. Tompson, D. Schuurmans, and P. Abbeel, “Learning interactive real-world simulators,” *arXiv preprint arXiv:2310.06114*, 2023.
- [33] J. Y. Koh, H. Lee, Y. Yang, J. Baldridge, and P. Anderson, “Pathdreamer: A world model for indoor navigation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 14 738–14 748.
- [34] A. Bar, G. Zhou, D. Tran, T. Darrell, and Y. LeCun, “Navigation world models,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 15 791–15 801.
- [35] C. Campos, R. Elvira, J. J. G. Rodriguez, J. M. M. Montiel, and J. D. Tardos, “Orb-slam3: An accurate open-source library for visual, visual-inertial, and multimap slam,” *IEEE Transactions on Robotics*, vol. 37, no. 6, p. 1874–1890, Dec. 2021.
- [36] J. L. Schönberger and J.-M. Frahm, “Structure-from-motion revisited,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 4104–4113.
- [37] N. Keetha *et al.*, “Mapanything: Universal feed-forward metric 3d reconstruction,” *arXiv preprint arXiv:2509.13414*, 2025.
- [38] J. Wang, M. Chen, N. Karaev, A. Vedaldi, C. Rupprecht, and D. Novotny, “Vggt: Visual geometry grounded transformer,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 5294–5306.
- [39] W. Peebles and S. Xie, “Scalable diffusion models with transformers,”

- in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 4195–4205.
- [40] Y. Lipman, R. T. Chen, H. Ben-Hamu, M. Nickel, and M. Le, “Flow matching for generative modeling,” *arXiv preprint arXiv:2210.02747*, 2022.
  - [41] X. Liu, C. Gong, and Q. Liu, “Flow straight and fast: Learning to generate and transfer data with rectified flow,” *arXiv preprint arXiv:2209.03003*, 2022.
  - [42] J. Liang, V. Makoviychuk, A. Handa, N. Chentanez, M. Macklin, and D. Fox, “Gpu-accelerated robotic simulation for distributed reinforcement learning,” in *Conference on Robot Learning*. PMLR, 2018, pp. 270–282.
  - [43] Gemini Team and Google, “Gemini 3.0: Multimodal reasoning at scale,” Google DeepMind, Tech. Rep., 2025. [Online]. Available: <https://deepmind.google/technologies/gemini/>
  - [44] E. J. Hu *et al.*, “Lora: Low-rank adaptation of large language models.” *Iclr*, vol. 1, no. 2, p. 3, 2022.
  - [45] S. Ge, A. Mahapatra, G. Parmar, J.-Y. Zhu, and J.-B. Huang, “On the content bias in fréchet video distance,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.