

CityWalker: Learning Embodied Urban Navigation from Web-Scale Videos

Xinhao Liu* Jintong Li* Yicheng Jiang Niranjana Sujay Zhicheng Yang
Juexiao Zhang John Abanes Jing Zhang Chen Feng[†]
New York University

<https://ai4ce.github.io/CityWalker/>

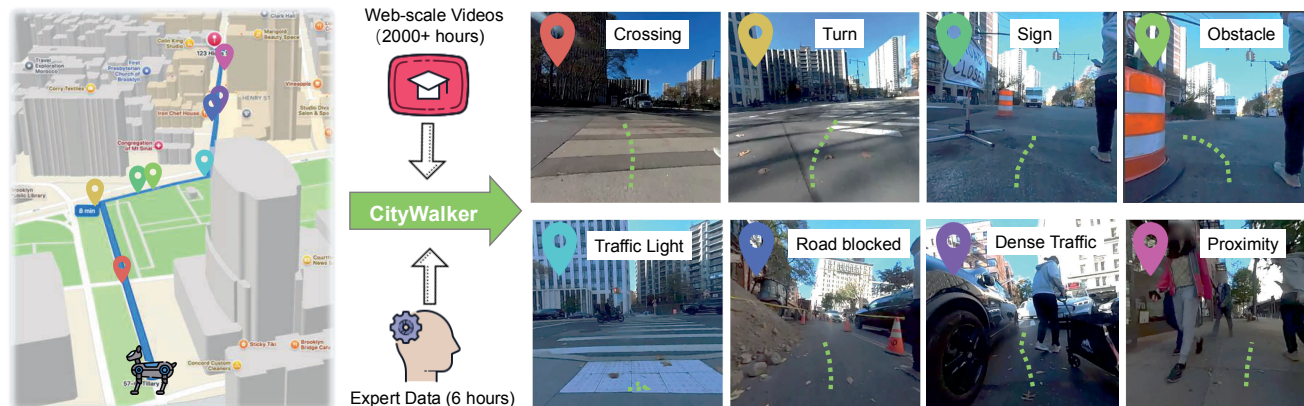


Figure 1. **Embodied Urban Navigation.** Navigating urban spaces is challenging for (especially off-street) mobile agents. The differently colored pins (📍) along the route highlight various critical scenarios unique to complex and dynamic urban landscapes. Thumbnails on the right with corresponding colored pins demonstrate the real-world observation of these challenging cases. Our CityWalker model is trained with over 2000 hours of city walking videos and fine-tuned with a small amount of expert data to address these challenges effectively.

Abstract

Navigating dynamic urban environments presents significant challenges for embodied agents, requiring advanced spatial reasoning and adherence to common-sense norms. Despite progress, existing visual navigation methods struggle in map-free or off-street settings, limiting the deployment of autonomous agents like last-mile delivery robots. To overcome these obstacles, we propose a scalable, data-driven approach for human-like urban navigation by training agents on thousands of hours of in-the-wild city walking and driving videos sourced from the web. We introduce a simple and scalable data processing pipeline that extracts action supervision from these videos, enabling large-scale imitation learning without costly annotations. Our model learns sophisticated navigation policies to handle diverse challenges and critical scenarios. Experimental results show that training on large-scale, diverse datasets significantly enhances navigation performance, surpassing current methods. This work shows the potential of using abundant online video data to develop robust navigation policies for embodied agents in dynamic urban settings.

*Equal contribution ({xinhao.liu, jintong.li}@nyu.edu)

[†]Corresponding author (cfeng@nyu.edu)

1. Introduction

Visual navigation is a crucial capability for mobile agents. In previously unseen environments, humans typically rely on navigation tools (such as Google Maps) to reach a goal location by following a series of waypoints. These tools provide high-level guidance, but the actual navigation between waypoints requires sophisticated spatial awareness and decision-making. In urban settings, this involves understanding and adhering to complex rules, such as common sense regulations and norms, while dynamically responding to obstacles and environmental changes. The diversity and complexity encountered in urban navigation scenarios pose significant challenges for navigation policies but are essential for deploying mobile robots in real-world settings. Applications like delivery robots and robotaxis, especially without high-definition maps, highlight the growing need for embodied agents to navigate efficiently and safely in dynamic urban landscapes. *How can we teach embodied agents to navigate dynamic urban landscapes as effectively as humans do?*

The question has been explored in early works [1–3] before the deep learning era. These approaches combined SLAM-based mapping with modular systems for search-

ing, detection, analysis, and planning. The development of high-performance photorealistic simulators [4–7] significantly advanced research in visual navigation by enabling reinforcement learning (RL) in static indoor environments, achieving near-perfect point-goal navigation [8]. This simulation-RL paradigm was also extended to more complex settings, including real-world panoramic street views [9–12]. Recent works begin to explore imitation learning from expert demonstrations for long-horizon or cross-embodiment navigation tasks [13, 14].

Despite their success in simulators or moderately complex environments, existing methods fall short in addressing the complexity of **embodied urban navigation**. Public urban spaces are inherently dynamic and unpredictable, characterized by “*multifarious terrains, diverse obstacles, and dense pedestrians*” that require real-time adaptability [15]. Moreover, effective navigation in such settings demands adherence to common sense navigation rules and social norms, such as using sidewalks, obeying traffic signals, and maintaining appropriate personal space to avoid conflicts [16–18]. Current reinforcement learning and imitation learning approaches typically excel in static or controlled environments but struggle with real-world urban navigation. This is because these nuanced behaviors and delicate constraints are challenging to incorporate into simulation environments while ensuring reliable and generalizable sim-to-real transfer. Moreover, coverage of these complexities is still lacking in existing teleoperation demonstration datasets. As a result, embodied urban navigation remains an **unsolved problem**.

While a straightforward approach to the problem is to collect expert trajectories through teleoperation, this is limited by data size and lack of diversity, which weakens the agent’s ability to generalize across various urban scenarios. *Is there a data-driven approach that requires minimum annotation but maximizes generalizability?* Inspired by the success of scaling laws in language, vision, and robotics tasks [19–22], we propose a scalable framework, **City-Walker**, which leverages web-scale city walking and driving videos to train models for embodied urban navigation. Our model is trained on over two thousand hours of internet-sourced videos across different geolocations, weather conditions, and times of day.

The subsequent question is *how to effectively obtain action supervision from in-the-wild videos for imitation learning*. Current solutions rely heavily on prompting proprietary VLMs to generate *action* labels from videos [23, 24], which is costly and challenging to control, making it difficult to scale for large datasets. We argue that noisy pseudo labels from off-the-shelf visual odometry (VO) models [25] are sufficiently effective for imitation learning in embodied urban navigation. Moreover, this data processing method generalizes beyond walking data. We demonstrate that

a model trained with driving videos also performs well on quadruped agents, with further improvements observed when training in a cross-domain setting that combines both city walking and driving data.

We summarize our contribution as the following:

- We identify embodied urban navigation as an intriguing yet unsolved problem. To address this problem, we propose a scalable, data-driven solution by leveraging web-scale city walking and driving videos.
- We introduce a simple yet scalable data processing paradigm, making large-scale imitation learning feasible for urban navigation without extensive manual labeling.
- We demonstrate training on web-scale data significantly improves navigation performance in real-world experiments, allowing embodied agents to handle the complexities of urban scenarios. We also show model performance scales well with training data size.

2. Related Work

Navigation in Simulation. The introduction of high-performance photorealistic simulators [4–7] has significantly boosted navigation research in indoor household environments. Based on the form of goal specification, navigation tasks are usually categorized into point-goal [4, 8, 26, 27], image-goal [28–31], object-goal [32–34], and vision-language navigation [35–37]. While tasks like image-goal and object-goal navigation focus on efficient goal-finding or instruction understanding, point-goal navigation has a more straightforward goal specification and emphasizes safe and efficient trajectory planning to the goal. Our problem formulation resembles point-goal navigation in that accurate waypoints (represented as GPS coordinates) are readily available from navigation tools. Recently, point-goal navigation has been considered a “solved” problem due to near-perfect performance achieved in simulators [8]. However, our investigation shows that waypoint-goal navigation in public urban spaces remains an unsolved challenge.

Real-World Navigation. Real-world robot navigation has a long history, predating the deep learning era. Early methods relied on ranging sensors and modular SLAM systems for navigation [1–3]. Witnessing the success in simulators, various methods have been proposed to bridge the sim-to-real gap [38–41]. Another thread of work learns navigation policies from expert trajectories [42, 43]. By incorporating data from different embodiments [14] and carefully designed heuristics [13], these methods have been extended to long-horizon navigation on different robots [44, 45]. While deployed on real-world embodied agents, these methods still operate in relatively simple and static environments such as suburban areas, parking garages, or controlled obstacle courses. They are not guaranteed to generalize to significantly different, dynamic, and complex urban landscapes. Our work follows this line of research by training naviga-

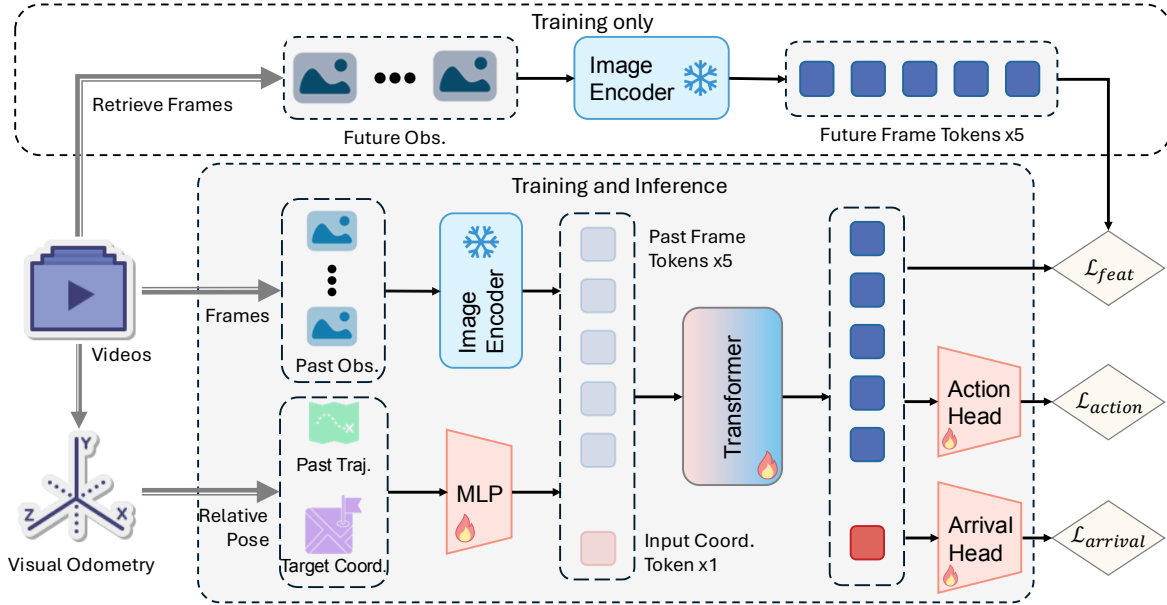


Figure 2. **Overall Illustration of CityWalker.** Our training pipeline starts with internet-sourced videos, using visual odometry to obtain relative poses between frames. At each time step, the model receives past observations, past trajectory, and target location as input. They are encoded via a frozen image encoder and a trainable coordinate encoder. A transformer processes these inputs to generate future tokens. An action head and an arrival head decode these tokens into action and arrival status predictions. During training, future frame tokens from future frames guide the transformer to hallucinate future tokens.

tion models with imitation learning. But we focus on learning the rules and norms essential for urban environments, which is crucial for enabling applications such as delivery robots and robo-taxis.

In-the-Wild Video Learning. Learning from web-scale data has proven successful in language and vision tasks [19–21, 46]. While some works utilize in-the-wild videos to extract intermediate representations, priors, or reward functions [47, 48], a key challenge remains: the lack of action labels. Patel et al. [49] use a reconstruction pipeline to obtain hand-object trajectories as action labels for learning robot manipulation. SelfD [50] argues BEV model is better than VO in pseudo-labeling, while we show VO works well with large-scale data. Recent works also rely on prompting closed VLMs to get action labels [23, 51]. LeLaN [24], a concurrent work most relevant to ours, also relies on VLM prompting and uses pretrained navigation models to generate action labels. In contrast, we process in-the-wild video data differently by using only the noisy pseudolabels from off-the-shelf visual SLAM tools [25]. This approach enables scalable and cost-effective action label generation for imitation learning in urban navigation tasks.

3. Embodied Urban Navigation

3.1. Problem Formulation

In this work, we address the problem of embodied urban navigation for agents in dynamic urban environments. The

agent’s task is to navigate from its current location to a specified target waypoint location, following a series of waypoints provided by a navigation tool. Formally, we define this as a point-goal navigation problem in real-world urban settings. At each time step t , the agent receives an RGB observation o_t , its current GPS location p_t , and a sub-goal waypoint w_t . The agent aims to learn a policy $\pi(a_t | o_{(t-k):t}, p_{(t-k):t}, w_t)$ that maps the past observations and positional information to an action a_t from the action space \mathcal{A} , represented as a series of action waypoints in Euclidean space. Typically, we take $k = 5$ and also predict actions for the next 5 time steps.

As accurate waypoints are easily available from navigation tools (e.g. Google Maps API), we focus our problem on the navigation between two consecutive waypoints. The agent treats the current sub-goal w_t as the immediate target and is responsible for determining when it has reached this sub-goal based on its observations and positional data. Upon reaching w_t , the agent proceeds to the next waypoint in the sequence.

3.2. Evaluation Metrics

To enable analysis and assessment of different methods on the problem, our evaluation metrics are twofold. We use the ground truth trajectory in teleoperation-collected data to provide a thorough analysis of each method’s performance across various scenarios. Below, we discuss the metrics we used in offline data evaluation. For real-world deployment,

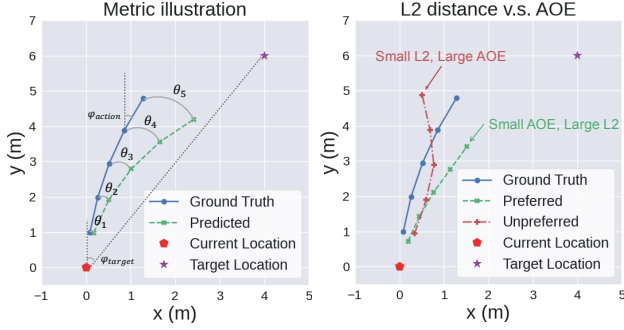


Figure 3. **Evaluation Metrics.** *Left:* The orientation error is defined to be the angle between each predicted and ground truth action pair, labeled by θ_n in the figure. The action angle φ_{action} and target angle φ_{target} are defined with respect to the positive y-axis. *Right:* Both green and red trajectories are predicted actions. The **green** trajectory is the preferred one, having a large L2 distance but a small AOE. Vice versa for the **red** trajectory.

we primarily use navigation success rate as the key metric for performance comparison.

Average Orientation Error. We use the error between each predicted action and the ground truth action as the main evaluation metric on the offline teleoperation data. While L2-distance in Euclidean space is a straightforward metric, we find it does not adequately capture the problem. As shown on the right side of Fig. 3, the red trajectory has a smaller L2 error compared to the green trajectory but moves in a direction that deviates from the target. This observation motivates us to come up with a new evaluation metric that better reflects the quality of the predicted actions. We define *average orientation error* (AOE) as the angle difference between the predicted and ground truth actions, averaged across all samples at each step (left side of Fig. 3):

$$\text{AOE}(k) = \frac{1}{n} \sum_i \theta_{i_k} = \frac{1}{n} \sum_i \arccos \frac{\langle \hat{a}_{i_k}, a_{i_k} \rangle}{\|\hat{a}_{i_k}\| \|a_{i_k}\|}, \quad (1)$$

where k is the index of predicted action and n is the number of data. Additionally, we observe that taking the mean across k predicted future time steps can lead to an underestimation due to steps with very small errors. Hence, we propose using the maximum average orientation error (MAOE) over all predicted actions as an overall estimation:

$$\text{MAOE} = \frac{1}{n} \sum_i \max_k \theta_{i_k}. \quad (2)$$

As orientation error might be noisy when action distance is small, we also use L2-distance as a complimentary evaluation metric.

Critical Scenarios. Not all time steps are equally critical in a trajectory. For example, errors at road crossings or turning points are more likely to result in navigation failure compared to mistakes during forward movement in open areas.




Figure 4. **Data Sample and Visual Odometry (VO) Result.** Our internet-source training data includes both walking and driving videos. These videos cover various scenarios in the urban environment. The VO tool gives noisy trajectories globally, but trustworthy local relative pose within a short time period.

We identify several key scenarios that are most critical during navigation. AOE and MAOE are calculated separately for each scenario, and we use the scenario mean rather than the sample mean for a more comprehensive evaluation. Detailed definitions of each scenario are provided in Sec. 4.

3.3. Learning from In-the-Wild Videos

City Walking Videos. To facilitate large-scale imitation learning for human-like urban navigation, we utilize city walking videos sourced from the internet. Walking in urban environments is intuitive for humans, resulting in an abundance of first-person view videos available online. These videos naturally capture the complexities of navigating dynamic urban landscapes, including interactions with pedestrians, adherence to traffic signals, and maneuvering around various obstacles. Motivated by the extensive availability of such data and its inherent alignment with our navigation objectives, we curate a dataset comprising over 2000 hours of city walking videos, encompassing diverse weather and lighting conditions. These videos exhibit consistent motion patterns characteristic of walking, providing continuous motion cues that can directly serve as action labels.

Action Labels from Videos. To generate actionable labels for imitation learning, we employ VO [25, 52–54] tools to extract trajectory poses from the collected city walking videos. Specifically, we utilize the state-of-the-art DPVO [25] while any other VO method should also work. The resulting VO trajectories are illustrated in Fig. 4. However, using VO poses presents two primary challenges. First, global trajectory accuracy is compromised by accumulation errors inherent in VO systems. Nevertheless, since our model focuses on predicting a limited sequence of actions (5 past and 5 future), we mitigate this issue by relying solely on relative poses within short temporal windows. Second, VO methods suffer from scale ambiguity, leading to inconsistencies in trajectory lengths across different data. For instance, as shown in Fig. 4, a driving video may exhibit a longer actual trajectory compared to a walking video, yet the DPVO output incorrectly reflects a shorter trajectory. To address this, we normalize each action by the average step length within its trajectory, following the approach in

previous work [14]. This normalization resolves both scale ambiguity and inconsistencies in step lengths between different types of embodiments all at once. This enables the learning of policies in a consistent, abstracted action space. During deployment, the predicted actions are denormalized using the specific step length of the target robot, ensuring appropriate movement execution in real-world settings.

Generalizability and Scalability. Our data processing pipeline is both generalizable and scalable. It can be applied to any video sources beyond city walking, as long as the videos are captured with egomotion, such as driving videos, which are also abundantly available online. This flexibility enables the training of more generalizable navigation policies that can be zero-shot applied or fine-tuned for cross-domain and cross-embodiment tasks, such as transitioning from pedestrian to vehicular navigation. Furthermore, our method offers significant scalability advantages. Unlike pipelines that rely on prompting large language models for action label generation [23, 24], our visual odometry-based approach can be easily parallelized. Processing 2000 hours of video data with our pipeline requires negligible wall-clock time, making it both feasible and cost-efficient to scale up the training dataset.

3.4. Pipeline and Training

Pipeline Architecture. Figure 2 shows our pipeline. It builds upon previous works [44, 45]. At the core of our model is a transformer that processes a temporal sequence of input tokens, including k image features and one coordinate embeddings. Specifically, the coordinate input is stacked by k past positions and one single target position, and outputs a sequence of the same length. We use an action head and an arrival prediction head to decode the output of the transformer. We utilize DINOv2 [21] representation and freeze the backbone during training.

Feature Hallucination. Inspired by previous works on feature learning [55–57], we use feature hallucination as an auxiliary loss during training. Specifically, we compute the MSE loss between the output image tokens and that directly extracted from future frames. We hope this will guide the transformer to predict more informative future tokens that best mimic future frames, which in turn helps the two MLP heads predict the final action and arrival status.

Loss Functions. As mentioned in Sec. 3.2, L1 or L2 loss can be less effective than orientation error in evaluating the quality of the predicted actions. We also employ orientation error to supervise the training. We define the orientation loss to be the negative cosine similarity between the predicted actions and the ground truth actions:

$$\mathcal{L}_{\text{ori}} = -\frac{1}{k} \sum_{i=1}^k \frac{\langle \hat{a}_i, a_i \rangle}{\|\hat{a}_i\| \|a_i\|}. \quad (3)$$

Together with the L1 loss for predicted action and the BCE

loss for predicted arriving status, the final loss function is a weight sum of four individual losses:

$$\mathcal{L} = \omega_{l1} \mathcal{L}_{l1} + \omega_{\text{ori}} \mathcal{L}_{\text{ori}} + \omega_{\text{arr}} \mathcal{L}_{\text{arr}} + \omega_{\text{feat}} \mathcal{L}_{\text{feat}}. \quad (4)$$

We choose each weight so that each loss items are within the same magnitude.

4. Experiment

4.1. Setup

Baselines. We compare our model with outdoor navigation models closely related to our setup, including GNM [14], ViNT [44], and NoMaD [45]. Although originally developed for image-goal navigation, we tested them using goal-images from our collected data. Additionally, ViNT as a foundation model for visual navigation supports fine-tuning, which we also evaluated. We acknowledge CoN-VOI [58] as a recent work with a similar setup, but could not test it due to the lack of open source code and prompt. We also experimented with VLM-based approaches similar to CoNVOI and defer the results to the appendix.

Data. We collected expert data via teleoperation for fine-tuning and offline testing. We collect data via a Unitree Go1 quadruped equipped with a Livox Mid-360 LiDAR and a webcam for RGB observations. We employed the LiDAR-SLAM method [59] to obtain the robot’s poses as ground truth actions. Additionally, we retrieved location data from a smartphone using a webpage. We reduce noise from the quadruped’s vibrations by holding the phone [38]. In total, we gathered 15 hours of teleoperation data across various areas in New York City, allocating 6 hours for fine-tuning and 9 hours for testing.

Critical Scenarios As discussed in Sec. 3.2, we identify several critical scenarios in our dataset that we want to emphasize in our evaluation. We analyze the ground truth trajectory and run object detection [60] to aid us in standardizing the definition. We define these scenarios as follows:

- *Turn:* When the ground truth action changes direction significantly. Defined when $\varphi_{\text{action}} > 20^\circ$ (see Fig. 3).
- *Crossing:* When the agent is at a road crossing. Defined if any traffic light is detected with score > 0.5 .
- *Detour:* When the action angle deviates from the target angle. Defined when $|\varphi_{\text{action}} - \varphi_{\text{target}}| > 45^\circ$ (see Fig. 3).
- *Proximity:* When any person is close to the agent. Defined when the largest bounding box of any detected person is larger than 25% of the image area.
- *Crowd:* When the agent is surrounded by a crowd of people. Defined when ≥ 5 people are detected.

We define these scenarios in a way that they are not mutually exclusive. A single data sample can belong to multiple scenarios. Although these scenarios account for less than half of the data, we argue they are the most significant factors contributing to successful urban navigation.

Table 1. **Benchmark on Offline Data.** We evaluate three metrics in each critical scenario for all methods. Percentages under scenarios indicate their data proportions. The “Mean” column shows scenario means averaged over six scenarios; “All” shows sample means over all data samples. We highlight the best performance in **bold** and the second-best performance with underline.

Method	Metric	Mean	Turn 8%	Crossing 12%	Detour 12%	Proximity 6%	Crowd 7%	Other 55%	All 100%
GNM [14] (fine-tuned)	↓ L2 (m)	<u>1.22</u>	2.36	1.36	1.42	0.88	0.76	0.55	<u>0.74</u>
	↓ MAOE (°)	<u>16.2</u>	31.1	<u>14.8</u>	12.5	<u>14.7</u>	<u>12.8</u>	<u>11.0</u>	<u>12.1</u>
	↑ Arrival (%)	68.6	66.4	69.7	66.4	69.0	69.2	70.7	70.0
ViNT [44] (fine-tuned)	↓ L2 (m)	1.30	1.91	1.13	1.14	0.77	0.66	<u>0.57</u>	0.70
	↓ MAOE (°)	16.5	31.1	15.4	<u>12.9</u>	14.8	13.3	11.6	12.6
	↑ Arrival (%)	70.5	71.2	68.7	70.7	73.0	68.6	71.0	70.7
NoMaD [45] (fine-tuned)	↓ L2 (m)	1.39	2.49	1.56	1.55	1.06	0.95	0.76	<u>0.74</u>
	↓ MAOE (°)	19.1	35.1	18.5	15.6	18.1	14.3	12.8	<u>12.1</u>
	↑ Arrival (%)	68.6	66.4	69.7	66.4	69.0	69.3	70.7	70.0
Ours (zero-shot)	↓ L2 (m)	1.34	<u>1.30</u>	<u>1.09</u>	1.33	1.44	1.48	1.39	1.38
	↓ MAOE (°)	16.5	<u>26.8</u>	15.5	16.3	16.3	13.1	11.3	12.7
	↑ Arrival (%)	<u>79.1</u>	<u>69.3</u>	<u>71.4</u>	78.8	<u>84.1</u>	<u>84.5</u>	<u>86.4</u>	<u>84.1</u>
Ours (fine-tuned)	↓ L2 (m)	1.11	1.27	1.00	<u>1.15</u>	1.06	1.12	1.06	1.07
	↓ MAOE (°)	15.2	26.6	14.1	13.9	14.3	12.0	10.4	11.5
	↑ Arrival (%)	81.8	68.9	75.3	<u>78.5</u>	90.6	87.5	90.2	87.8

4.2. Performance Benchmarking

In this subsection, we aim to answer the question **Q1**: Can our CityWalker model navigate successfully in complex urban environments? To answer this, we benchmark our model and the baselines both with offline data and in real-world navigation deployment.

Overall Benchmark. Table 1 shows the benchmark in different critical scenarios. Overall, our fine-tuned model has superior performance on all metrics in all scenarios except “detour”. We think the less competent performance in the detour scenario is due to the small portion of such data in our video training data. It can be verified by the fact that the fine-tuning performance gain is most significant in this scenario. It is worth noting that our zero-shot model has similar or even better performance compared to the fine-tuned baselines. This demonstrates the power of large-scale pretraining, especially considering the domain and modality gap from human walking videos.

Real-world Deployment. We deploy our model and the baseline methods on the same Unitree Go1 quadruped for real-world navigation, in previously unseen urban environments. We adapt the PD controller provided in [44] to provide velocity command to the quadruped. We choose area with relatively constant density of dynamics such as pedestrian and traffic to ensure the consistency among different trials. In each trial, the target location is generally 50-100m far away from the starting location. We categorize the real-world test into *forward*, *left turn*, and *right turn* cases to provide case-by-case analysis. We conduct 8-14 trails for each cases. Each trial is marked success when the model predicts arrival within a 5m distance to the target location.

Table 2. **Real-World Navigation.** The table shows the success rate of real-world experiments in different scenarios. * and † indicate zero-shot inference and fine-tuned models respectively.

Method	All	Forward	Left turn	Right turn
ViNT*[44]	37.7	62.5	0.0	50.0
ViNT†[44]	57.1	100.0	25.0	25.0
NoMad*[45]	42.9	75.0	16.7	28.6
Ours†	77.3	100.0	62.5	66.7

All human interruptions due to potential collision or time-out are treated as failure cases.

Table 2 shows the success rate on all real-world test cases. Our model achieves the highest success rate among all cases with a significant gap to the second-best performing fine-tuned ViNT. As the baselines are trained with navigation data mostly in suburban or off-road environments [14, 61], they are limited in handling complex maneuvers, making them unsuitable for urban navigation tasks, as shown in Fig. 5. Our model’s robust performance in both forward and turns highlights its ability to manage dynamic and varied urban scenarios effectively, allowing it to be more reliably used in real-world environments where frequent and precise directional changes are essential. The substantial improvement not only increases the overall navigation success but also extends the applicability of our model to more demanding and realistic urban settings, positioning it as a superior solution for embodied agents in dynamic environments.

4.3. Power of Data Scaling

Our model is significantly boosted by the large-scale training data. We ask **Q2**: To what extent does increasing the

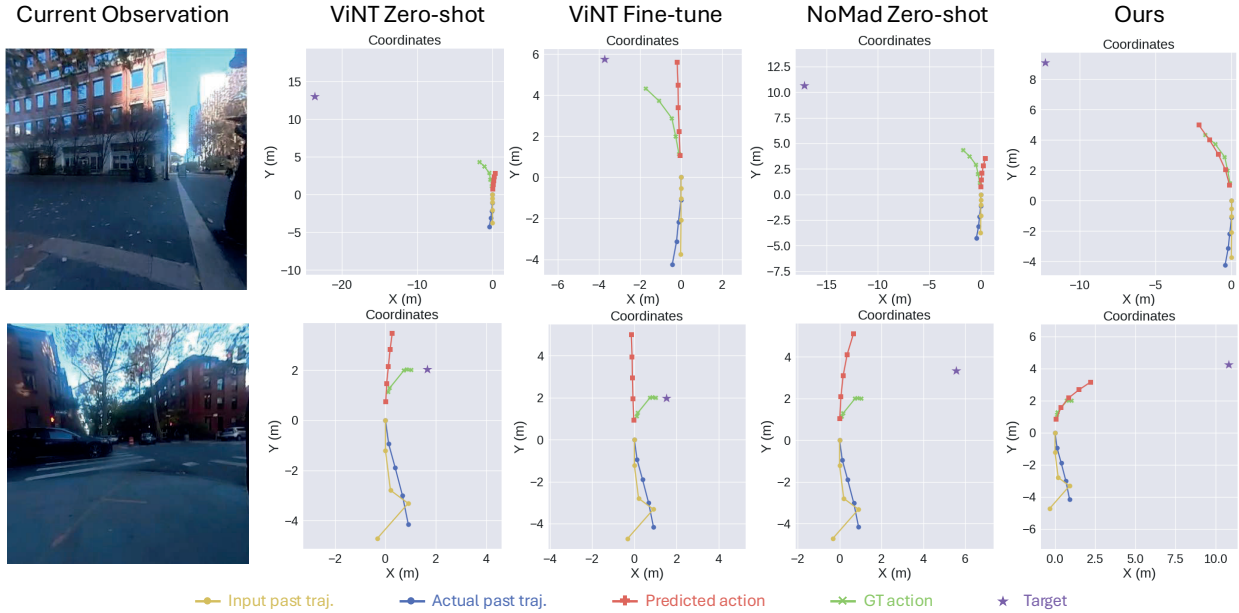


Figure 5. **Qualitative Results.** Left image shows current observations of two samples. Right plots displays input trajectory, ground truth actions, and predicted actions in the current coordinate system with the agent at the origin.




Figure 6. **Performance and Data Size.** We show the model performance evaluated by MAOE with respect to the size of the training data measured by video length in hours. We also show the zero-shot performance of our model trained with only driving videos and mixed driving and walking videos

size of the training dataset improve the performance of our urban navigation model?

Performance Gain and Data Scale. Figure 6 shows the model performance when trained with different data sizes. The plot demonstrates that our model’s zero-shot performance improves significantly as we increase the amount of training data. Notably, when trained with more than 1000 hours of videos, our zero-shot model (red line) achieves a better performance than the fine-tuned ViNT model (\square). This indicates that large-scale training on diverse video data can even surpass the performance of models that are fine-tuned with limited expert data, highlighting the potential of leveraging web-scale datasets for embodied navigation.

Domain and Embodiment Gap. An unexpected observation is that our model *without* the feature hallucination loss outperforms the one *with* it in zero-shot inference. We hypothesize this phenomenon mainly due the domain and embodiment gap between human walking videos and quadruped navigation policies. The feature hallucination loss aims to guide the transformer model to produce output tokens that mimic future observations, essentially encouraging the model to anticipate future visual inputs. This enforcement might inadvertently hinder performance by focusing on predicting future observations based on human motion. The model may learn representations that are less applicable or even misleading for quadruped navigation. Consequently, the feature hallucination loss could introduce a bias that detracts from the model’s ability to make accurate navigation decisions in the robot’s operational domain.

Feature Hallucination Loss. Fortunately, as shown in Tab. 1, this issue no longer exists after fine-tuning. We also show the benefit of feature hallucination loss from another perspective. Figure 7 shows two training loss curve for different data size and training setups with and without feature hallucination loss. We observe that although feature hallucination loss should increase the absolute value of the total loss, the training with feature hallucination loss lead to a lower total loss for 2000-hour training and a lower orientation loss in both cases. The training with feature hallucination loss also shows a steeper decreasing trend, showing a potential of having even lower error with more training steps. Therefore, this is another proof of the effectiveness of our proposed feature hallucination loss.

Driving Videos. As discussed in Sec. 3.3, our data process-




Figure 7. **Improvement from Feature Hallucination.** The plot shows the orientation loss for different amounts of training data and loss setups. For all four setups, we train with $\omega_{11} = 1.0$, $\omega_{arr} = 1.0$, and $\omega_{ori} = 5.0$. The unproportional training steps are due to different batch sizes used for different data sizes. All models are trained for 10 epochs.

ing pipeline also enables the use of abundant online driving videos. We evaluate models trained exclusively on driving videos and on a combination of walking and driving videos, results shown in Fig. 6. We can observe that the model trained only with driving data (+) has a similar performance to the zero-shot baselines. As a result, we can think of driving as a form of navigation from a different domain and embodiment. However, combining driving and walking data yields significant performance gains (x). With just 250 hours of mixed data, the model nearly matches the performance of models trained on 1000 hours of walking data alone. This signifies the generalizability and robustness provided by cross-domain cross-embodiment data.

4.4. Pipeline component analysis

To provide a more in-depth analysis of our model and pipeline, we ask **Q3**: How do the individual components of our pipeline enhance the performance and reliability of our model for real-world urban navigation?

AOE at Different Time Steps. In Tab. 4, we study the AOE at different time steps, i.e. $AOE(k)$ for k from 1 to 5. ViNT exhibits increasing orientation errors as the prediction steps progress, indicating a degradation in accuracy over time. In contrast, our model maintains relatively consistent AOE from step 2 to step 5, which likely contributes to its superior performance in real-world navigation deployments with stable action prediction across multiple future steps. All methods experience higher errors at the first step. We attribute it to the proximity of the initial action to the origin, as minor linear deviations result in substantial angular error. It is worth noting that the MAOE is significantly larger than any individual AOE from steps 1 to 5, suggesting that the peak errors occur at varying time steps across different samples. This variability suggests the potential benefit of taking softmax for orientation errors instead of simple averaging, as it

Table 3. **Ablation Study.** We ablate on the contribution of orientation loss, feature hallucination, and fine-tuning with expert data. We show the MAOE($^\circ$) in scenario-mean. All results are trained with 1000-hour walking data.

Training Components			MAOE
Ori. Loss	Feature Hall.	Fine-tuning	Mean
			17.03
✓			17.00
✓	✓		17.02
		✓	15.23
✓		✓	<u>15.21</u>
✓	✓	✓	15.16

Table 4. **AOE at Different Time Steps.** The table shows MAOE and different AOE of scenario-mean, all in degrees ($^\circ$). * and † indicates zero-shot inference and fine-tuned models respectively.

Method	AOE(1)	AOE(2)	AOE(3)	AOE(4)	AOE(5)	MAOE
ViNT* [44]	10.57	7.93	9.09	10.21	11.45	17.52
ViNT† [44]	11.24	7.96	8.40	8.93	9.87	16.53
Ours*	12.91	8.92	8.64	8.46	8.43	16.54
Ours†	11.44	7.77	7.68	7.75	7.97	15.23

could better capture the distribution of peak errors.

Ablation Study. Table 3 examines the individual contributions of different pipeline components to our model’s performance. A clear improvement in MAOE when fine-tuning with expert data, demonstrating the effectiveness of using in-domain expert information. The improvement from orientation loss and feature hallucination loss is marginal, within error range of our limited testing data. We defer more detailed ablation study to the appendix.

5. Conclusion

Summary. In this work, we addressed embodied urban navigation by training agents with thousands of hours of in-the-wild city walking videos, introducing a scalable pipeline for imitation learning from diverse walking videos. Our experiments show that large-scale, diverse training data significantly enhances navigation performance, surpassing existing methods and demonstrating the benefits of data scaling. Our findings highlight the potential of leveraging abundant online videos to develop robust navigation policies for embodied agents in dynamic urban environments.

Limitation and Future Work. During real-world deployment, our current system implementation is sensitive to potential large location noises read naively from iPhone location services, which could be improved with better GPS hardware. Future work could also enhance the robustness of the model to such location noises.

Acknowledgment

The work was supported by NSF grants 2238968, 2121391, 2322242 and 2345139; and in part through the NYU IT High Performance Computing resources, services, and staff expertise. We also thank Xingyu Liu and Zixuan Hu for their help in data collection.

References

- [1] Quirin Muhlbauer, Stefan Sosnowski, Tingting Xu, Tianguang Zhang, Kolja Kuhnlenz, and Martin Buss. Navigation through urban environments by visual perception and interaction. In *ICRA*, pages 3558–3564. IEEE, 2009. 1, 2
- [2] Rainer Kümmerle, Michael Ruhnke, Bastian Steder, Cyrill Stachniss, and Wolfram Burgard. A navigation system for robots operating in crowded urban environments. In *ICRA*, pages 3225–3232. IEEE, 2013.
- [3] Yoichi Morales, Alexander Carballo, Eiji Takeuchi, Atsushi Aburadani, and Takashi Tsubouchi. Autonomous robot navigation in outdoor cluttered pedestrian walkways. *Journal of Field Robotics*, 26(8):609–635, 2009. 1, 2
- [4] Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, et al. Habitat: A platform for embodied ai research. In *CVPR*, pages 9339–9347, 2019. 2
- [5] Fei Xia, Amir R Zamir, Zhiyang He, Alexander Sax, Jitendra Malik, and Silvio Savarese. Gibson env: Real-world perception for embodied agents. In *CVPR*, pages 9068–9079, 2018.
- [6] Eric Kolve, Roozbeh Mottaghi, Winson Han, Eli VanderBilt, Luca Weihs, Alvaro Herrasti, Matt Deitke, Kiana Ehsani, Daniel Gordon, Yuke Zhu, et al. Ai2-thor: An interactive 3d environment for visual ai. *arXiv preprint arXiv:1712.05474*, 2017.
- [7] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from RGB-D data in indoor environments. *3DV*, 2017. 2
- [8] Erik Wijmans, Abhishek Kadian, Ari S. Morcos, Stefan Lee, Irfan Essa, Devi Parikh, Manolis Savva, and Dhruv Batra. Dd-ppo: Learning near-perfect pointgoal navigators from 2.5 billion frames. In *ICLR*, 2019. 2
- [9] Piotr Mirowski, Matt Grimes, Mateusz Malinowski, Karl Moritz Hermann, Keith Anderson, Denis Teplyashin, Karen Simonyan, Andrew Zisserman, Raia Hadsell, et al. Learning to navigate in cities without a map. *NeurIPS*, 31, 2018. 2
- [10] Ang Li, Huiyi Hu, Piotr Mirowski, and Mehrdad Farajtabar. Cross-view policy learning for street navigation. In *ICCV*, pages 8100–8109, 2019.
- [11] Karl Moritz Hermann, Mateusz Malinowski, Piotr Mirowski, Andras Banki-Horvath, Keith Anderson, and Raia Hadsell. Learning to follow directions in street view. In *AAAI*, volume 34, pages 11773–11781, 2020.
- [12] Piotr Mirowski, Andras Banki-Horvath, Keith Anderson, Denis Teplyashin, Karl Moritz Hermann, Mateusz Malinowski, Matthew Koichi Grimes, Karen Simonyan, Koray Kavukcuoglu, Andrew Zisserman, et al. The streetlearn environment and dataset. *arXiv preprint arXiv:1903.01292*, 2019. 2
- [13] Dhruv Shah and Sergey Levine. ViKiNG: Vision-Based Kilometer-Scale Navigation with Geographic Hints. In *RSS*, 2022. 2
- [14] Dhruv Shah, Ajay Sridhar, Arjun Bhorkar, Noriaki Hirose, and Sergey Levine. Gnm: A general navigation model to drive any robot. In *ICRA*, pages 7226–7233. IEEE, 2023. 2, 5, 6
- [15] Wayne Wu, Honglin He, Yiran Wang, Chenda Duan, Jack He, Zhizheng Liu, Quanyi Li, and Bolei Zhou. Metaurban: A simulation platform for embodied ai in urban spaces. *arXiv preprint arXiv:2407.08725*, 2024. 2
- [16] Adarsh Jagan Sathyamoorthy, Utsav Patel, Moumita Paul, Nithish K Sanjeev Kumar, Yash Savle, and Dinesh Manocha. Comet: Modeling group cohesion for socially compliant robot navigation in crowded scenes. *IEEE Robotics and Automation Letters*, 7(2):1008–1015, 2021. 2
- [17] Muchen Sun, Francesca Baldini, Peter Trautman, and Todd Murphey. Move beyond trajectories: Distribution space coupling for crowd navigation. In *RSS*, 2021.
- [18] Yuya Onozuka, Ryosuke Matsumi, and Motoki Shino. Autonomous mobile robot navigation independent of road boundary using driving recommendation map. In *IROS*, pages 4501–4508. IEEE, 2021. 2
- [19] Tom B Brown. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020. 2, 3
- [20] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *ICCV*, pages 4015–4026, 2023.
- [21] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel HAZIZA, Francisco Massa, Alaaeldin El-Nouby, Mido Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning robust visual features without supervision. *Transactions on Machine Learning Research*, 2024. 3, 5
- [22] Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, et al. Palme: An embodied multimodal language model. *arXiv preprint arXiv:2303.03378*, 2023. 2
- [23] Beichen Wang, Juexiao Zhang, Shuwen Dong, Irving Fang, and Chen Feng. Vlm see, robot do: Human demo video to robot action plan via vision language model. *arXiv preprint arXiv:2410.08792*, 2024. 2, 3, 5
- [24] Noriaki Hirose, Catherine Glossop, Ajay Sridhar, Dhruv Shah, Oier Mees, and Sergey Levine. Lelan: Learning a language-conditioned navigation policy from in-the-wild videos. *arXiv preprint arXiv:2410.03603*, 2024. 2, 3, 5
- [25] Zachary Teed, Lahav Lipson, and Jia Deng. Deep patch visual odometry. *NeurIPS*, 36, 2024. 2, 3, 4

- [26] Max Jaderberg, Volodymyr Mnih, Wojciech Marian Czarnecki, Tom Schaul, Joel Z Leibo, David Silver, and Koray Kavukcuoglu. Reinforcement learning with unsupervised auxiliary tasks. In *ICLR*, 2017. 2
- [27] Saurabh Gupta, James Davidson, Sergey Levine, Rahul Suktanekar, and Jitendra Malik. Cognitive mapping and planning for visual navigation. In *CVPR*, pages 2616–2625, 2017. 2
- [28] Nikolay Savinov, Alexey Dosovitskiy, and Vladlen Koltun. Semi-parametric topological memory for navigation. In *ICLR*, 2018. 2
- [29] Meera Hahn, Devendra Singh Chaplot, Shubham Tulsiani, Mustafa Mukadam, James M Rehg, and Abhinav Gupta. No rl, no simulation: Learning to navigate without navigating. *NeurIPS*, 34:26661–26673, 2021.
- [30] Obin Kwon, Nuri Kim, Yunho Choi, Hwiyeon Yoo, Jeongho Park, and Songhwai Oh. Visual graph memory with unsupervised representation for visual navigation. In *CVPR*, pages 15890–15899, 2021.
- [31] Jacob Krantz, Theophile Gervet, Karmesh Yadav, Austin Wang, Chris Paxton, Roozbeh Mottaghi, Dhruv Batra, Jitendra Malik, Stefan Lee, and Devendra Singh Chaplot. Navigating to objects specified by images. In *CVPR*, pages 10916–10925, 2023. 2
- [32] Theophile Gervet, Soumith Chintala, Dhruv Batra, Jitendra Malik, and Devendra Singh Chaplot. Navigating to objects in the real world. *Science Robotics*, 8(79):eadf6991, 2023. 2
- [33] Devendra Singh Chaplot, Dhiraj Prakashchand Gandhi, Abhinav Gupta, and Russ R Salakhutdinov. Object goal navigation using goal-oriented semantic exploration. *NeurIPS*, 33:4247–4258, 2020.
- [34] Arjun Majumdar, Gunjan Aggarwal, Bhavika Devnani, Judy Hoffman, and Dhruv Batra. Zson: Zero-shot object-goal navigation using multimodal goal embeddings. *NeurIPS*, 35:32340–32352, 2022. 2
- [35] Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. Embodied question answering. In *CVPR*, pages 1–10, 2018. 2
- [36] Hao Tan, Licheng Yu, and Mohit Bansal. Learning to navigate unseen environments: Back translation with environmental dropout. *arXiv preprint arXiv:1904.04195*, 2019.
- [37] Jacob Krantz, Erik Wijmans, Arjun Majumdar, Dhruv Batra, and Stefan Lee. Beyond the nav-graph: Vision-and-language navigation in continuous environments. In *ECCV*, pages 104–120. Springer, 2020. 2
- [38] Maks Sorokin, Jie Tan, C Karen Liu, and Sehoon Ha. Learning to navigate sidewalks in outdoor environments. *IEEE Robotics and Automation Letters*, 7(2):3906–3913, 2022. 2, 5
- [39] Joanne Truong, Max Rudolph, Naoki Harrison Yokoyama, Sonia Chernova, Dhruv Batra, and Akshara Rai. Rethinking sim2real: Lower fidelity simulation leads to higher sim2real transfer in navigation. In *CoRL*, pages 859–870. PMLR, 2023.
- [40] Joanne Truong, Denis Yarats, Tianyu Li, Franziska Meier, Sonia Chernova, Dhruv Batra, and Akshara Rai. Learning navigation skills for legged robots with learned robot embeddings. In *IROS*, pages 484–491. IEEE, 2021.
- [41] Wei Liu, Huihua Zhao, Chenran Li, Joydeep Biswas, Billy Okal, Pulkit Goyal, Yan Chang, and Soha Pouya. X-mobility: End-to-end generalizable navigation via world modeling. *arXiv preprint arXiv:2410.17491*, 2024. 2
- [42] Gregory Kahn, Pieter Abbeel, and Sergey Levine. Badgr: An autonomous self-supervised learning-based navigation system. *IEEE Robotics and Automation Letters*, 6(2):1312–1319, 2021. 2
- [43] Gregory Kahn, Pieter Abbeel, and Sergey Levine. Land: Learning to navigate from disengagements. *IEEE Robotics and Automation Letters*, 6(2):1872–1879, 2021. 2
- [44] Dhruv Shah, Ajay Sridhar, Nitish Dashora, Kyle Stachowicz, Kevin Black, Noriaki Hirose, and Sergey Levine. ViNT: A foundation model for visual navigation. In *CoRL*, 2023. 2, 5, 6, 8
- [45] Ajay Sridhar, Dhruv Shah, Catherine Glossop, and Sergey Levine. Nomad: Goal masked diffusion policies for navigation and exploration. In *ICRA*, pages 63–70, 2024. 2, 5, 6
- [46] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763. PMLR, 2021. 3
- [47] Annie S Chen, Suraj Nair, and Chelsea Finn. Learning generalizable robotic reward functions from” in-the-wild” human videos. *arXiv preprint arXiv:2103.16817*, 2021. 3
- [48] Shikhar Bahl, Abhinav Gupta, and Deepak Pathak. Human-to-robot imitation in the wild. In *RSS*, 2022. 3
- [49] Austin Patel, Andrew Wang, Ilija Radosavovic, and Jitendra Malik. Learning to imitate object interactions from internet videos. *arXiv preprint arXiv:2211.13225*, 2022. 3
- [50] Jimuyang Zhang, Ruizhao Zhu, and Eshed Ohn-Bar. Selfd: Self-learning large-scale driving policies from the web. In *CVPR*, pages 17316–17326, 2022. 3
- [51] Naoki Wake, Atsushi Kanehira, Kazuhiro Sasabuchi, Jun Takamatsu, and Katsushi Ikeuchi. Gpt-4v (ision) for robotics: Multimodal task planning from human demonstration. *IEEE Robotics and Automation Letters*, 2024. 3
- [52] Jakob Engel, Vladlen Koltun, and Daniel Cremers. Direct sparse odometry. *TPAMI*, 40(3):611–625, 2017. 4
- [53] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. Orb-slam: a versatile and accurate monocular slam system. *IEEE transactions on robotics*, 31(5):1147–1163, 2015.
- [54] Zachary Teed and Jia Deng. Droid-slam: Deep visual slam for monocular, stereo, and rgb-d cameras. *NeurIPS*, 34:16558–16569, 2021. 4
- [55] Yuhang He, Irving Fang, Yiming Li, Rushi Bhavesh Shah, and Chen Feng. Metric-Free Exploration for Topological Mapping by Task and Motion Imitation in Feature Space. In *RSS*, July 2023. 5
- [56] Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. Self-supervised learning from images with a joint-embedding predictive architecture. In *CVPR*, pages 15619–15629, 2023.

- [57] Adrien Bardes, Quentin Garrido, Jean Ponce, Xinlei Chen, Michael Rabbat, Yann LeCun, Mahmoud Assran, and Nicolas Ballas. Revisiting feature prediction for learning visual representations from video. *arXiv preprint arXiv:2404.08471*, 2024. [5](#)
- [58] Adarsh Jagan Sathyamoorthy, Kasun Weerakoon, Mohamed Elnoor, Anuj Zore, Brian Ichter, Fei Xia, Jie Tan, Wenhao Yu, and Dinesh Manocha. Convoi: Context-aware navigation using vision language models in outdoor and indoor environments. *arXiv preprint arXiv:2403.15637*, 2024. [5](#)
- [59] Wei Xu and Fu Zhang. Fast-lid: A fast, robust lidar-inertial odometry package by tightly-coupled iterated kalman filter. *IEEE Robotics and Automation Letters*, 6(2):3317–3324, 2021. [5](#)
- [60] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. [5](#)
- [61] Noriaki Hirose, Dhruv Shah, Ajay Sridhar, and Sergey Levine. Sacson: Scalable autonomous control for social navigation. *IEEE Robotics and Automation Letters*, 2023. [6](#)