

NavGSim: High-Fidelity Gaussian Splatting Simulator for Large-Scale Navigation

Jiahang Liu^{1,2,*} Yuanxing Duan^{2,*} Jiazhao Zhang^{1,2,*}
 Minghan Li² Shaoan Wang¹ Zhizheng Zhang^{2,3,†} He Wang^{1,2,3,†}
¹Peking University ²Galbot ³BAAI




Fig. 1: NavGSim’s agent-simulator architecture. Users interact with the system through an upper-level python API interface, which provides access to the core functionalities of NavGSim.

Abstract—Simulating realistic environments for robots is widely recognized as a critical challenge in robot learning, particularly in terms of rendering and physical simulation. This challenge becomes even more pronounced in navigation tasks, where trajectories often extend across multiple rooms or even entire floors. In this work, we present NavGSim, a Gaussian Splatting-based simulator designed to generate high-fidelity, large-scale navigation environments. Built upon a hierarchical 3D Gaussian Splatting framework, NavGSim enables photorealistic rendering in expansive scenes spanning hundreds of square meters. To simulate navigation collisions, we introduce a Gaussian Splatting-based slice technique that directly extracts navigable areas from reconstructed Gaussians. Additionally, for ease of use, we provide comprehensive NavGSim APIs supporting multi-GPU development, including tools for custom scene reconstruction, robot configuration, policy training, and evaluation. To evaluate NavGSim’s effectiveness, we train a Vision-Language-Action (VLA) model using trajectories collected from the NavGSim and assess its performance in both simulated and real-world environments. Our results demonstrate that NavGSim significantly enhances the VLA model’s scene understanding, enabling the policy to handle diverse navigation queries effectively. *NavGSim is publicly available at: <https://github.com/2003jiahang/NavGSim>*

I. INTRODUCTION

Simulating realistic environments is widely regarded as a key component of robot learning[1], [2], providing a playground for robots to acquire diverse skills. Existing

robotics simulators cover a broad range of applications, including locomotion[3], [4], manipulation[5], [6], [7], and navigation[8], [9]. Notably, unlike locomotion and manipulation tasks, navigation tasks are distinguished by their requirement for high-fidelity, large-scale scene environments—often spanning multiple rooms or even an entire floor. To simulate such large-scale environments, existing navigation simulators (e.g., Habitat [10] and AI2-THOR [11]) rely on reconstructed or synthetic meshes. However, these approaches often result in unrealistic visual rendering. While artists can construct highly realistic scenes in game engines like Unreal or Unity, creating large-scale environments still requires significant effort.

In recent years, significant advancements have been made in achieving realistic rendering in navigation simulators through the application of 3D Gaussian splatting (3DGS) [12]. 3DGS, which provides high-quality visual rendering, has gained traction in various domains, including navigation, due to its ability to easily reconstruct real-world environments by moving a camera and LiDAR around the environment. Recent works [13], [14], [15], [16], [17], [18] have leveraged 3DGS to enhance the fidelity and scalability of navigation tasks.

Building on these innovations, we proposed NavGSim, a 3D Gaussian Splatting-based simulator designed for high-fidelity, large-scale navigation tasks. Our method extends hierarchical 3D Gaussian Splatting techniques[19] and in-

* Equal Contribution, † Equal Advising

corporates Gaussian-based collision simulation. Specifically, NavGSim organizes 3D Gaussians into a multi-level hierarchical structure, enabling adaptive level-of-detail (LOD) rendering that dynamically balances visual fidelity and computational efficiency during navigation. To support realistic collision detection, Gaussian ellipsoids are sliced by multiple horizontal planes, and the resulting intersection regions define the object’s occupancy representation on each plane. Consequently, NavGSim converts reconstructed scenes into navigable areas, enabling efficient collision checking and path planning.

Building on NavGSim, we introduce NavGSim-API, a Python library for custom scene reconstruction, robot configuration, policy training, and evaluation. Our API enables researchers to reconstruct scenes from their own posed RGB images. The resulting 3D Gaussian Splatting scenes are fully compatible with NavGSim and support core navigation simulator functionalities, as demonstrated in prior work [10], [11].

To evaluate the advantages of NavGSim over existing mesh-based navigation simulators, we fine-tune a vision-language-action (VLA) model [20] using navigation samples collected directly from NavGSim. We gather 4.5K navigation samples to reach 9 annotated targets. Our results show that the fine-tuned VLA model achieves superior navigation performance in NavGSim and demonstrates strong generalizability in real-world environments by effectively interpreting environmental cues to complete navigation tasks. We hope NavGSim will benefit the research community by enabling the construction of customized navigation simulators for diverse applications.

II. RELATED WORKS

Navigation Simulator. Navigation simulators[21], [11], [10], [22], [23], [24], [25], [26], [27] play a crucial role in embodied AI and robot learning by offering controllable environments for navigation tasks such as Vision-and-language navigation[28], [29], Object Goal Navigation[30], [31], Embodied Question Answering[32], [33] and more. Classic simulators such as Gazebo [21], AI2-THOR [11] and Habitat [10] provide various levels of realism and interactivity. Note that, the scene environments of navigation simulators often span across multi rooms or one floor. To balance the rendering quality and efficiency, existing navigation simulators adopt conventional rendering techniques[34], [35] with reconstructed meshes[36], [37], [38] or synthetic meshes[11], [39], which may limit visual fidelity. In contrast, NavGSim leverages 3D Gaussian Splatting, enabling more realistic and scalable simulations.

Gaussian Splatting for Robots. 3D Gaussian Splatting (3DGS) [12] has recently emerged as a compelling alternative to traditional methods like NeRF[40]. It models a scene with a set of 3D Gaussian kernels that jointly encode density, color, and opacity. As a significant advancement in 3D reconstruction, 3DGS has high-quality reconstruction results while keeping real-time rendering performance, offering a powerful tool for bridging real-world data to simulated

environments, i.e., Real2Sim[41], [42]. For navigation tasks, the required reconstructed environments often span multiple rooms or even entire floors. Although numerous previous works [43], [44], [45], [46] have addressed large-scale scene reconstruction, they cannot be directly applied to navigation tasks, which involve complex layouts with diverse geometries and heavy occlusions. To address this gap, Hierarchical 3D Gaussian Splatting (H3DGS) [19] proposes a hierarchical region-wise reconstruction framework that spatially decomposes the environment and renders it at multiple levels of detail. This design supports efficient and high-fidelity simulation over large, complex spaces.

```
class NavGSim:
    def __init__(self, handler):
        self.handler = handler
        self.handler.launch()
        self.info = {}

    def reset(self):
        self.handler.reset()
        self._update_info()

    def step(self, action):
        self.handler.step(action)
        self._update_info()

    def render(self, camera):
        return self.handler.render(camera)

    def path_planner(self, goal):
        return self.handler.path_planner(goal)

    def collision_check(self, states):
        return self.handler.collision_check(states)

    def close(self):
        self.handler.close()

    def _update_info(self):
        states = self.handler.get_states()
        self.info["obs"] = self.handler.get_obs()
        self.info["states"] = states
        self.info["success"] = self._get_success(states)
        self.info["done"] = self._get_done(states)
```

Fig. 2: NavGSim’s sample code

III. NAVGSIM PLATFORM

A. NavGSim Architecture

Fig. 1 illustrates NavGSim, a modular and scalable simulator for large-scale navigation tasks with photo-realistic rendering path planning, and collision detection. These capabilities are accessible through a simple Python API, supporting the entire pipeline from scene reconstruction to embodied agent training. The usage of APIs is illustrated in Code 2. Together, these components form a unified architecture that balances visual realism, navigational reasoning, and ease of use.

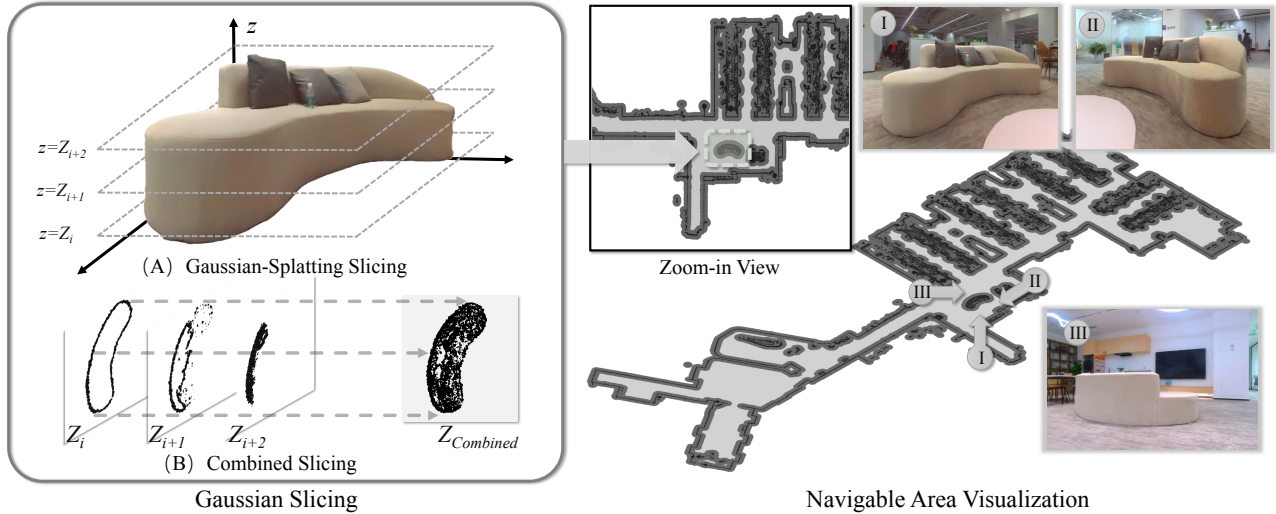


Fig. 3: **Illustration of Gaussian Slicing for Navigable Areas.** The left image shows the 3D Sofa model projected onto multiple horizontal Z-planes (Z_i, Z_{i+1}, Z_{i+2}) at fixed height intervals. Each plane contains a 2D Gaussian slice that represents the projection of the Sofa at that particular height. These slices are combined to form a 2D collision map, with the right image showing the entire scene’s 2-D occupancy map that serves as geometric prior information for downstream tasks such as navigation cost-map construction or collision checking.

NavGSim leverages a Hierarchical 3D Gaussian Splatting [19] (H3DGS) engine Sec. III-B that models indoor scenes as a hierarchy of anisotropic 3D Gaussians. This enables high-quality, real-time rendering from arbitrary viewpoints using only camera intrinsics and extrinsics. To support navigation, we further introduce a planar GS-slice technique Sec. III-C that extracts floor-level traversable regions directly from the 3D Gaussian field. This slice provides a dense and accurate approximation of collision volumes, enabling fast and efficient collision check and path planning during simulation. Lastly, to facilitate scene annotation, we developed an interactive visualization interface, allowing users to browse the reconstructed scene and label specific objects during their exploration.

B. Hierarchical 3D Gaussian Splatting scenes

NavGSim utilizes an H3DGS representation for scene modeling and real-time rendering, which is built around a tree-based structure where each node corresponds to a spatial region. Leaf nodes encode fine-grained visual details using dense sets of anisotropic Gaussians, while internal nodes aggregate and abstract these details into coarser levels of representation. This hierarchical structure supports Level-of-Detail (LOD) rendering by dynamically selecting Gaussians from different levels and smoothly interpolating transition regions. To handle large scenes, such as entire floors or multi-room layouts, we adopt a divide-and-conquer strategy. The input image set is divided into spatial chunks, each optimized independently. After local models are trained, a global hierarchy is built by merging and optimizing the chunks, including node cleanup and alignment. We show our reconstruction results in Fig. 5. This approach enables NavSim to process scenes spanning hundreds to thousands of

square meters using multiple GPUs in parallel. Ensuring both rendering speed and quality, this 3DGS-based representation also provides the foundation for collision detection in the navigation simulator.

C. GS-based Navigation collision simulation

To support efficient and realistic collision detection, NavGSim uses a multi-slice Gaussian projection method, which leverages the continuous and probabilistic nature of 3D Gaussian fields, while avoiding the problem of directly projecting 3D ellipses to 2D. Inspired by 4D-Rotor-GS [47], which slices 4DGS by different time planes, we define a set of horizontal slicing planes at fixed height intervals (e.g., every 10 cm in the robot’s height range) parallel to the floor. Each plane intersects the 3D Gaussian field, and for each Gaussian component, we analytically compute its sliced 2D Gaussian based on its position and 3D covariance. This slicing process is illustrated in Fig. 3.

Specifically, we define the description of 3D Gaussian:

$$G_{3D}(\mathbf{x}) = e^{-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}_{3D}^{-1}(\mathbf{x} - \boldsymbol{\mu})}, \quad (1)$$

$$\boldsymbol{\Sigma}_{3D}^{-1} = \begin{pmatrix} c_{xx} & c_{xy} & c_{xz} \\ c_{xy} & c_{yy} & c_{yz} \\ c_{xz} & c_{yz} & c_{zz} \end{pmatrix}, \quad (2)$$

where $\mathbf{x} = (x, y, z)$. Then, given a height z , the sliced 2D Gaussian is obtained as (detailed derivation in the Supplementary Material):

$$G_{2D}(\mathbf{x}, z) = e^{-\frac{1}{2}\lambda(z - \mu_z)^2} e^{[\mathbf{x} - \boldsymbol{\mu}(z)]^T \boldsymbol{\Sigma}_{2D}^{-1}[\mathbf{x} - \boldsymbol{\mu}(z)]}, \quad (3)$$

where $\mathbf{x} = (x, y)$,

$$\begin{cases} \alpha c_{xx} + \beta c_{xy} = c_{xz} \\ \alpha c_{xy} + \beta c_{yy} = c_{yz} \end{cases}, \quad \boldsymbol{\Sigma}_{2D}^{-1} = \begin{pmatrix} c_{xx} & c_{xy} \\ c_{xy} & c_{yy} \end{pmatrix}, \quad (4)$$

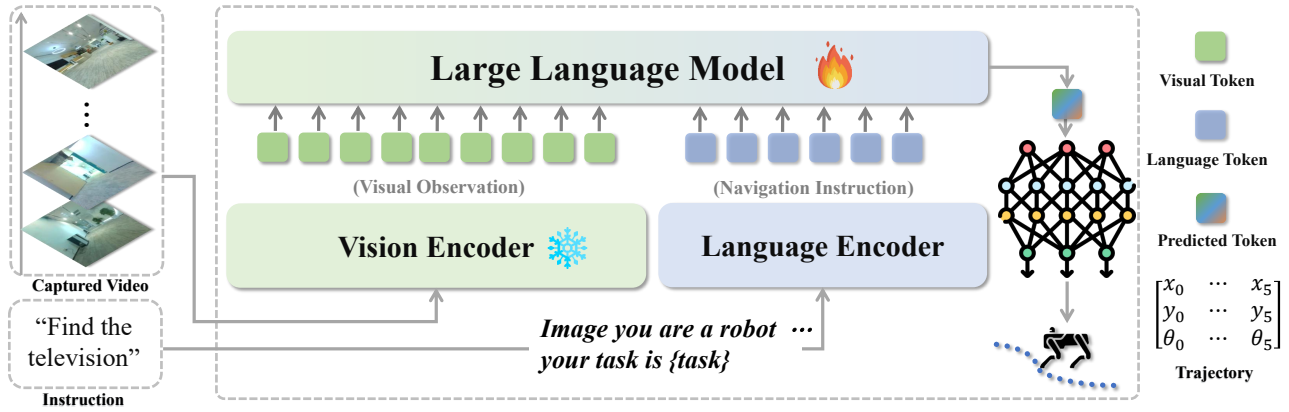


Fig. 4: **VLA-NavGSim architecture.** Our VLA model takes multiple frames of images and language instructions as input. The images are processed by a Vision Encoder to extract visual features. The resulting visual tokens, along with the language tokens, are fed into the LLM to generate a predicted token. This token is then passed through a lightweight MLP to generate a trajectory.

$$\lambda = c_{zz} - \alpha^2 c_{xx} - \beta^2 c_{yy} - 2\alpha\beta c_{xy}, \quad (5)$$

$$\mu(z) = (\mu_x, \mu_y)^T - (z - \mu_z) (\alpha, \beta)^T. \quad (6)$$

In Equation 3, we introduced a decay term, denoted as $\lambda(z - \mu_z)^2$, which determines the size of the 2D Gaussian slice. When $z = \mu_z$, the maximum 2D slice of the Gaussian is achieved, and as z moves further away from μ_z , the slice gradually shrinks. Simultaneously, $\mu(z)$ also varies with z , influencing the position of the 2D Gaussian.

Although 3D Gaussian Splatting serves as a probabilistic representation and cannot precisely capture the volume, due to the small size of our Gaussian volumes, we are able to limit the error in the collision volume to within a 5 cm range.”

IV. NAVGSIM FOR VLA

A. VLA architecture

We extend an existing navigation Vision-Language-Action (VLA) model called Uni-NaVid [48], with trajectory output. Uni-NaVid is a unified video-language model tailored for long-horizon embodied navigation, which fuses visual observations and natural language instructions to predict agent actions. In this work, we adapt and extend Uni-NaVid to enable continuous trajectory prediction, which is essential for realistic navigation in high-fidelity environments.

Visual Encoding and Token Compression. The model processes a sequence of egocentric RGB frames and a navigation instruction to enable effective multi-modal inference. Each frame is first encoded by an EVA-CLIP[49] vision encoder, which divides it into 16×16 non-overlapping patches, producing 256 visual tokens. To ensure computational efficiency during long-horizon navigation, we adopt the online token merging technique from Uni-NaVid, which temporally organizes visual tokens into three categories: current tokens from the most recent frame with

the highest spatial resolution, short-term tokens from the past 64 frames with grid pooling, and long-term tokens from earlier frames with aggressively pooling and token merging based on a cosine similarity threshold. These compressed visual tokens are then projected into the language embedding space via a Cross-Modality Projector (2-layer MLP) to match the latent space of the Vicuna-7B LLM[50]. The navigation instruction is tokenized and embedded using the same LLM. Finally, the visual and language tokens are concatenated into a unified input sequence: $\{\text{Long_term_tokens}\} \{\text{Short_term_tokens}\} \{\text{Current_tokens}\} \langle \text{NAV} \rangle \{\text{Instruction}\}$, allowing the LLM to jointly attend to both visual observations and instructions.

Modified Output Representation. Unlike the original Uni-NaVid design, which outputs discrete action tokens, inspired by [51], we reformulate the prediction target to suit continuous trajectory planning. Specifically, the LLM performs only a single autoregressive step, and we extract the hidden state from its final layer as the condition, which is subsequently passed through an MLP-based action head to predict the agent’s future trajectory in the form: $(x, y, \theta) \in \mathbb{R}^3$, where (x, y) and θ denote the agent’s position and yaw angle in its current coordinate frame, respectively. Compared to the discrete action space used in Uni-NaVid, our continuous trajectory output enables smoother and more realistic motion, better aligning with the requirements of large-scale real-world or simulated environments, where discretization may lead to unnatural movements or overshooting.

B. Data Collection

We collect a diverse set of 4,500 navigation trajectories within the NavGSim simulator. Each trajectory starts from a randomly sampled navigable position and heads toward a semantically meaningful goal, ensuring broad coverage of spatial configurations and behaviors.

Targets Annotation. Within the reconstructed environment, we manually annotate 9 semantic landmarks (*e. g.*,




Fig. 5: **Visualisation of reconstructed Gaussian Splatting and rendering examples.** (Left) We illustrate the process of subdividing the scene into spatial chunks using H3DGS for efficient reconstruction. (Right) We compare rendering results of NavGSim with real-world image at seven distinct viewpoints (A-F).

elevator, television, meeting room, etc.) as navigation goals. For each goal, we collect 200 expert trajectories generated by a shortest-path planner. Each trajectory starts from a randomly sampled position and rotation and ends at the selected specified landmark, ensuring a wide variety of spatial configurations and path topologies. To address the compounding error problem in policy inference, we employ the DAGger algorithm [52] to iteratively augment the dataset. In each round of DAGger, the agent navigates using its own policy, while an expert policy provides supervision at each step. In this work, we conduct two rounds of DAGger, collecting 150 additional trajectories per landmark in each round, resulting in a total of 2,700 trajectories. This significantly enhances the robustness and diversity of the training dataset.

Imitation Data Collection. Initially, the dataset is constructed using trajectories generated by a shortest-path planner operating under idealized simulation conditions. Although this supervised data helps the model learn fundamental navigation priors, we observe a significant drop in performance during inference. Specifically, the model tends to accumulate small prediction errors over time, which eventually lead to navigation failure—particularly in long-horizon tasks or cluttered environments.

Dagger Data Collection. To address this issue, we incorporate Dataset Aggregation (DAGger)[52], a technique that iteratively refines the model by collecting data under its own policy. During rollout, the agent executes its predicted actions, and trajectory collection is terminated under one of the following conditions: (1) the agent collides with the environment for n consecutive steps (we set $n=8$); (2) the number of inference steps exceeds a predefined budget; (3) the agent successfully reaches the goal. For each collected trajectory, the expert policy is queried to provide corrective supervision when the agent’s behavior diverges. This enables the model to learn from its own failure cases, improving robustness in out-of-distribution states and enhancing recovery capabilities.

Together, these strategies yield a training corpus that combines expert planning, DAGger-augmented behavior, and multi-modal supervision. This hybrid approach equips the VLA model with both accurate motion prediction and high-level scene comprehension, enabling reliable long-horizon navigation in complex environments.

V. EXPERIMENTS

We evaluate the effectiveness of NavGSim in both simulation and real-world environments through comprehensive experiments. Our goal is to validate whether the trained VLA model can generalize to novel start positions, maintain high success rates during long-horizon navigation, and transfer its policy effectively from simulation to the real world.

A. Experimental Setup

Gaussian Splatting Scene Construction. To evaluate NavGSim, we constructed a realistic large-scale indoor simulation environment of an office factory, covering approximately 700 m². We used a handheld LiDAR-RGB camera device to perform a walk-through scan of the entire scene, capturing multi-view images and LiDAR point clouds along natural human trajectories. These data were then processed using our H3DGS pipeline to reconstruct a photorealistic 3D environment.

Training Details of VLA. We train the VLA model using a combination of supervised learning and DAGger-augmented imitation data, as described in Sec. IV-B. To further improve generalization, we also incorporate video question-answering (VQA) data from publicly available datasets[53], [54], [55] into the training process. During training, the vision encoder (EVA-CLIP [49]) is frozen to preserve its pre-trained visual features, while the large language model (Vicuna-7B [50]) is fully fine-tuned along with the cross-modal projection layers and the trajectory regression head. Training is conducted on a cluster of 16 NVIDIA H800 GPUs for 20 hours, totaling 320 GPU-hours.




Fig. 6: **Visualization of real-world results.** Visualization First-person perspective display of real-world testing.

Real-World Deployment and Evaluation. We deploy our NavGSim-trained VLA agent on a Unitree Go2 quadruped robot in office environments, using a front-facing RGB camera to stream egocentric vision via WebSocket to an NVIDIA A100 server. The purely simulated-trained model fuses real-time visual data with language instructions to generate navigation waypoints, achieving 8 Hz closed-loop control through the robot’s onboard planner (Unitree extension dock). This modular architecture decouples semantic planning from low-level execution while maintaining responsiveness, with physical deployment demonstrations captured in Fig. 6.

Metrics. During evaluation, we randomly sample initial positions and orientations and measure the success rate of navigation to the 9 annotated landmarks and two unannotated novel objects. A trajectory is considered successful if the agent stops within 1.5 meters of the target. We use two metrics: Success Rate (SR) and Distance to Goal (DTG).

B. Performance Analysis




Fig. 7: **Real machine deployment architecture** The agent interacts with the server via the Internet.

Performance on simulator and real-world environments. We evaluate the navigation performance of the NavGSim-trained VLA model across both simulated and real-world environments. Table I summarizes the navigation success rates and stop distances across seen and unseen targets in the simulation environment. The model

demonstrates strong performance on seen landmarks such as *Television* (100.0%) and *Laboratory* (100.0%), achieving precise and efficient trajectories as indicated by low stop distances. Furthermore, it generalizes effectively to unseen targets, with success rates of 78.3% on *Table* and 83.3% on *Sofa*, suggesting robust spatial understanding and goal generalization beyond the training distribution. In real-world deployment, despite the presence of sensor noise and domain discrepancies, the trained model maintains high success rates on key targets, demonstrating the model’s ability to transfer navigation skills learned in simulation to physical environments. See the Supplemental Material for the results video.

Comparison of rendering quality. Moreover, we compare the training performance in a simulator built on original 3DGS. As shown in Table II, the rendering quality of the original 3DGS simulator is lower than that built by H3DGS. Although the model performs reasonably well in simulation, its low-quality rendering results in a significant performance gap when transferred to the real world, resulting a sharp decrease (81% ↓) in real-world performance. These results confirm the practical feasibility of sim-to-real transfer enabled by the NavGSim framework.

Influence of training strategy. We compared the progression of the success rate in different training stages for models trained with and without pre-training, as shown in Figure 8. The model utilizing pretraining shows significantly accelerated learning and higher final success rates, achieving over 90% success after two rounds of DAgger. In contrast, models trained from scratch exhibit slower improvement and plateau at substantially lower performance levels. This comparison highlights the critical role of pretraining in improving data efficiency, reducing compounding errors, and enabling more effective long-term navigation policy learning. The pre-trained models leverage previously learned features, allowing them to make faster adjustments and achieve a higher success rate much earlier in the training process.

Target Category	Target	H3DGS / 3DGS		
		Sim SR(%) \uparrow	Sim DTG(m) \downarrow	Real SR(%) \uparrow
Annotated Targets	Pantry	96.7 / 83.3	0.61 / 1.75	66.7 / 0.00
	Elevator	98.3 / 83.3	0.58 / 2.73	71.4 / 16.7
	Restroom	96.7 / 86.1	1.77 / 2.10	73.3 / 23.1
	Television	100 / 91.6	0.39 / 1.28	73.3 / 18.1
	Laboratory	100 / 83.3	0.30 / 1.64	75.0 / 18.1
	Meeting Room I	100 / 80.6	0.88 / 2.62	80.0 / 23.1
	Meeting Room II	100 / 77.8	0.50 / 3.82	63.0 / 15.4
	Reception Room	98.3 / 72.2	1.21 / 4.29	63.3 / 9.10
	Water Dispenser	95.0 / 69.4	1.92 / 5.42	57.1 / 0.00
	Avg	98.3 / 80.8	0.91 / 2.85	69.2 / 13.7
Novel Targets	Table	78.3 / 52.8	3.91 / 7.62	68.8 / 16.7
	Sofa	83.3 / 55.6	2.56 / 6.97	70.4 / 0.00
	Avg	80.8 / 54.2	3.23 / 7.29	69.6 / 8.35

TABLE I: Comparison of different rendering techniques. We report navigation performance across seen and unseen targets in the simulator and real-world environments.

Method	PSNR \uparrow	LPIPS \downarrow	SSIM \uparrow
H3DGS	25.396	0.274	0.866
3DGS [12]	23.221	0.426	0.789

TABLE II: Comparison on different methods

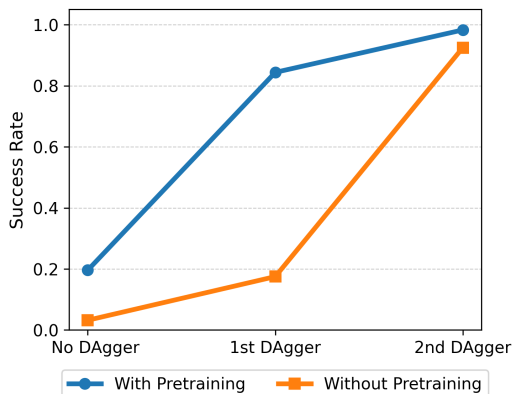


Fig. 8: Performance on NavGSim with different training strategy. Success rate comparison with and without pretraining across different training stages.

VI. CONCLUSIONS

In this work, we present NavGSim, a simulator based on Hierarchical Gaussian Splatting for high-fidelity, large-scale navigation. Compared to traditional mesh-based simulators, NavGSim achieves photorealistic rendering and precise collision modeling via the proposed GS-Slice technique. We integrate NavGSim into a real2sim2real navigation pipeline, constructing large-scale environments and adopting an adapted Uni-NaVid architecture with continuous pose regression. Experiments in both simulation and real-world settings demonstrate improved scene understanding for VLA models and effective sim-to-real transfer. By bridging high-quality scene reconstruction and embodied learning, NavGSim enables scalable and generalizable robot navigation.

REFERENCES

- [1] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, “A survey of robot learning from demonstration,” *Robotics and autonomous systems*, vol. 57, no. 5, pp. 469–483, 2009.
- [2] A. A. Rusu, M. Večerík, T. Rothörl, N. Heess, R. Pascanu, and R. Hadsell, “Sim-to-real robot learning from pixels with progressive nets,” in *Conference on robot learning*. PMLR, 2017, pp. 262–270.
- [3] G. B. Margolis, G. Yang, K. Paigwar, T. Chen, and P. Agrawal, “Rapid locomotion via reinforcement learning,” *The International Journal of Robotics Research*, vol. 43, no. 4, pp. 572–587, 2024.
- [4] Y. Yang, T. Zhang, E. Coumans, J. Tan, and B. Boots, “Fast and efficient locomotion via learned gait transitions,” in *Conference on Robot Learning*. PMLR, 2022, pp. 773–783.
- [5] O. Kroemer, S. Niekum, and G. Konidaris, “A review of robot learning for manipulation: Challenges, representations, and algorithms,” *Journal of machine learning research*, vol. 22, no. 30, pp. 1–82, 2021.
- [6] R. Liu, F. Nageotte, P. Zanne, M. de Mathelin, and B. Dresp-Langley, “Deep reinforcement learning for the control of robotic manipulation: a focussed mini-review,” *Robotics*, vol. 10, no. 1, p. 22, 2021.
- [7] L. Shao, T. Migimatsu, Q. Zhang, K. Yang, and J. Bohg, “Concept2robot: Learning manipulation concepts from instructions and human demonstrations,” *The International Journal of Robotics Research*, vol. 40, no. 12-14, pp. 1419–1434, 2021.
- [8] X. Xiao, B. Liu, G. Warnell, and P. Stone, “Motion planning and control for mobile robot navigation using machine learning: a survey,” *Autonomous Robots*, vol. 46, no. 5, pp. 569–597, 2022.
- [9] K. Zhu and T. Zhang, “Deep reinforcement learning based mobile robot navigation: A review,” *Tsinghua Science and Technology*, vol. 26, no. 5, pp. 674–691, 2021.
- [10] M. Savva, A. Kadian, O. Maksymets, Y. Zhao, E. Wijmans, B. Jain, J. Straub, J. Liu, V. Koltun, J. Malik *et al.*, “Habitat: A platform for embodied ai research,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 9339–9347.
- [11] E. Kolve, R. Mottaghi, W. Han, E. VanderBilt, L. Weihs, A. Herrasti, M. Deitke, K. Ehsani, D. Gordon, Y. Zhu *et al.*, “Ai2-thor: An interactive 3d environment for visual ai,” *arXiv preprint arXiv:1712.05474*, 2017.
- [12] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, “3d gaussian splatting for real-time radiance field rendering,” *ACM Trans. Graph.*, vol. 42, no. 4, pp. 139–1, 2023.
- [13] H. He, Y. Ma, W. Wu, and B. Zhou, “From seeing to experiencing: Scaling navigation foundation models with reinforcement learning,” *arXiv preprint arXiv:2507.22028*, 2025.
- [14] Y. Deng, S. Yuan, G. C. R. Bethala, A. Tzes, Y.-S. Liu, and Y. Fang, “Hierarchical scoring with 3d gaussian splatting for instance image-goal navigation,” *arXiv preprint arXiv:2506.07338*, 2025.
- [15] D. Ong, Y. Tao, V. Murali, I. Spasojevic, V. Kumar, and P. Chaudhari, “Gaussian splatting as a unified representation for autonomy in unstructured environments,” *arXiv preprint arXiv:2505.11794*, 2025.

- [16] K. Honda, T. Ishita, Y. Yoshimura, and R. Yonetani, "Gsplattnm: Point-of-view synthesis for visual navigation models using gaussian splatting," *arXiv preprint arXiv:2503.05152*, 2025.
- [17] D. Ong, Y. Tao, V. Murali, I. Spasojevic, V. Kumar, and P. Chaudhari, "Atlas navigator: Active task-driven language-embedded gaussian splatting," *arXiv preprint arXiv:2502.20386*, 2025.
- [18] V. Ress, J. Meyer, W. Zhang, D. Skuddis, U. Soergel, and N. Haala, "3d gaussian splatting aided localization for large and complex indoor-environments," *arXiv preprint arXiv:2502.13803*, 2025.
- [19] B. Kerbl, A. Meuleman, G. Kopanas, M. Wimmer, A. Lanvin, and G. Drettakis, "A hierarchical 3d gaussian representation for real-time rendering of very large datasets," *ACM Transactions on Graphics (TOG)*, vol. 43, no. 4, pp. 1–15, 2024.
- [20] J. Zhang, K. Wang, S. Wang, M. Li, H. Liu, S. Wei, Z. Wang, Z. Zhang, and H. Wang, "Uni-navid: A video-based vision-language-action model for unifying embodied navigation tasks," *arXiv preprint arXiv:2412.06224*, 2024.
- [21] N. Koenig and A. Howard, "Design and use paradigms for gazebo, an open-source multi-robot simulator," in *2004 IEEE/RSJ international conference on intelligent robots and systems (IROS)(IEEE Cat. No. 04CH37566)*, vol. 3. Ieee, 2004, pp. 2149–2154.
- [22] C. Beattie, J. Z. Leibo, D. Teplyashin, T. Ward, M. Wainwright, H. Küttler, A. Lefrancq, S. Green, V. Valdés, A. Sadik *et al.*, "Deepmind lab," *arXiv preprint arXiv:1612.03801*, 2016.
- [23] C. Yan, D. Misra, A. Bennet, A. Walsman, Y. Bisk, and Y. Artzi, "Chalet: Cornell house agent learning environment," *arXiv preprint arXiv:1801.07357*, 2018.
- [24] X. Puig, K. Ra, M. Boben, J. Li, T. Wang, S. Fidler, and A. Torralba, "Virtualhome: Simulating household activities via programs," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 8494–8502.
- [25] X. Gao, R. Gong, T. Shu, X. Xie, S. Wang, and S.-C. Zhu, "Vrkitchen: an interactive 3d virtual environment for task-oriented learning," *arXiv preprint arXiv:1903.05757*, 2019.
- [26] F. Xiang, Y. Qin, K. Mo, Y. Xia, H. Zhu, F. Liu, M. Liu, H. Jiang, Y. Yuan, H. Wang *et al.*, "Sapien: A simulated part-based interactive environment," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 097–11 107.
- [27] C. Gan, J. Schwartz, S. Alter, D. Mrowca, M. Schrimpf, J. Traer, J. De Freitas, J. Kubilius, A. Bhandwadar, N. Haber *et al.*, "Threedworld: A platform for interactive multi-modal physical simulation," *arXiv preprint arXiv:2007.04954*, 2020.
- [28] P. Anderson, Q. Wu, D. Teney, J. Bruce, M. Johnson, N. Sünderhauf, I. Reid, S. Gould, and A. Van Den Hengel, "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.
- [29] J. Krantz, E. Wijmans, A. Majumdar, D. Batra, and S. Lee, "Beyond the nav-graph: Vision-and-language navigation in continuous environments," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVIII 16*. Springer, 2020, pp. 104–120.
- [30] D. S. Chaplot, D. P. Gandhi, A. Gupta, and R. R. Salakhutdinov, "Object goal navigation using goal-oriented semantic exploration," *Advances in Neural Information Processing Systems*, vol. 33, pp. 4247–4258, 2020.
- [31] A. Majumdar, G. Aggarwal, B. Devnani, J. Hoffman, and D. Batra, "Zson: Zero-shot object-goal navigation using multimodal goal embeddings," *Advances in Neural Information Processing Systems*, vol. 35, pp. 32 340–32 352, 2022.
- [32] A. Das, S. Datta, G. Gkioxari, S. Lee, D. Parikh, and D. Batra, "Embodied question answering," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 1–10.
- [33] A. Majumdar, A. Ajay, X. Zhang, P. Putta, S. Yenamandra, M. Henaff, S. Silwal, P. Mcvay, O. Maksymets, S. Arnaud *et al.*, "Openeqa: Embodied question answering in the era of foundation models," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2024, pp. 16 488–16 498.
- [34] H. Delingette, "General object reconstruction based on simplex meshes," *International journal of computer vision*, vol. 32, pp. 111–146, 1999.
- [35] D. Terzopoulos and M. Vasilescu, "Sampling and reconstruction with adaptive meshes," in *CVPR*, vol. 91, 1991, pp. 70–75.
- [36] S. K. Ramakrishnan, A. Gokaslan, E. Wijmans, O. Maksymets, A. Clegg, J. Turner, E. Undersander, W. Galuba, A. Westbury, A. X. Chang *et al.*, "Habitat-matterport 3d dataset (hm3d): 1000 large-scale 3d environments for embodied ai," *arXiv preprint arXiv:2109.08238*, 2021.
- [37] A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Niessner, M. Savva, S. Song, A. Zeng, and Y. Zhang, "Matterport3d: Learning from rgb-d data in indoor environments," *arXiv preprint arXiv:1709.06158*, 2017.
- [38] F. Xia, A. R. Zamir, Z. He, A. Sax, J. Malik, and S. Savarese, "Gibson env: Real-world perception for embodied agents," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 9068–9079.
- [39] E. H. T. Collaboration *et al.*, "Repair of groin hernia with synthetic mesh: meta-analysis of randomized controlled trials," *Annals of surgery*, vol. 235, no. 3, pp. 322–332, 2002.
- [40] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "Nerf: Representing scenes as neural radiance fields for view synthesis," *Communications of the ACM*, vol. 65, no. 1, pp. 99–106, 2021.
- [41] D. Hahn, P. Banzet, J. M. Bern, and S. Coros, "Real2sim: Visco-elastic parameter estimation from dynamic motion," *ACM Transactions on Graphics (TOG)*, vol. 38, no. 6, pp. 1–13, 2019.
- [42] D. Robinson, F. Haldi, P. Leroux, D. Perez, A. Rasheed, and U. Wilke, "Citysim: Comprehensive micro-simulation of resource flows for sustainable urban planning," in *Proceedings of the Eleventh International IBPSA Conference*, 2009, pp. 1083–1090.
- [43] X. Ren, Y. Lu, H. Liang, Z. Wu, H. Ling, M. Chen, S. Fidler, F. Williams, and J. Huang, "Scube: Instant large-scale scene reconstruction using voxplats," *arXiv preprint arXiv:2410.20030*, 2024.
- [44] Y. Liu, C. Luo, L. Fan, N. Wang, J. Peng, and Z. Zhang, "Citygaussian: Real-time high-quality large-scale scene rendering with gaussians," in *European Conference on Computer Vision*. Springer, 2024, pp. 265–282.
- [45] Y. Liu, C. Luo, Z. Mao, J. Peng, and Z. Zhang, "Citygaussianv2: Efficient and geometrically accurate reconstruction for large-scale scenes," *arXiv preprint arXiv:2411.00771*, 2024.
- [46] J. Guo, S. Yang, Y. Huang, Y. Long, X. Xia, X. Su, B. Zhao, Z. Xie, and L. Nie, "Utilgen: Utility-centric generative data augmentation with dual-level task adaptation," *arXiv preprint arXiv:2510.24262*, 2025.
- [47] Y. Duan, F. Wei, Q. Dai, Y. He, W. Chen, and B. Chen, "4d-rotor gaussian splatting: towards efficient novel view synthesis for dynamic scenes," in *ACM SIGGRAPH 2024 Conference Papers*, 2024, pp. 1–11.
- [48] J. Zhang, K. Wang, R. Xu, G. Zhou, Y. Hong, X. Fang, Q. Wu, Z. Zhang, and H. Wang, "Navid: Video-based vlm plans the next step for vision-and-language navigation," *arXiv preprint arXiv:2402.15852*, 2024.
- [49] Q. Sun, J. Wang, Q. Yu, Y. Cui, F. Zhang, X. Zhang, and X. Wang, "Eva-clip-18b: Scaling clip to 18 billion parameters," *arXiv preprint arXiv:2402.04252*, 2024.
- [50] W.-L. Chiang, Z. Li, Z. Lin, Y. Sheng, Z. Wu, H. Zhang, L. Zheng, S. Zhuang, Y. Zhuang, J. E. Gonzalez *et al.*, "Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality," See <https://vicuna.lmsys.org> (accessed 14 April 2023), vol. 2, no. 3, p. 6, 2023.
- [51] D. Shah, A. Sridhar, A. Bhorkar, N. Hirose, and S. Levine, "Gnm: A general navigation model to drive any robot," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 7226–7233.
- [52] S. Ross, G. Gordon, and D. Bagnell, "A reduction of imitation learning and structured prediction to no-regret online learning," in *Proceedings of the fourteenth international conference on artificial intelligence and statistics*. JMLR Workshop and Conference Proceedings, 2011, pp. 627–635.
- [53] D. Azuma, T. Miyashita, S. Kurita, and M. Kawanabe, "Scanqa: 3d question answering for spatial scene understanding," in *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 19 129–19 139.
- [54] T.-S. Chen, A. Siarohin, W. Menapace, E. Deyneka, H.-w. Chao, B. E. Jeon, Y. Fang, H.-Y. Lee, J. Ren, M.-H. Yang *et al.*, "Panda-70m: Captioning 70m videos with multiple cross-modality teachers," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 13 320–13 331.
- [55] Y. Li, C. Wang, and J. Jia, "Llama-vid: An image is worth 2 tokens in large language models," in *European Conference on Computer Vision*. Springer, 2024, pp. 323–340.