

Localized Graph-Based Neural Dynamics Models for Terrain Manipulation

Chaoqi Liu¹, Yunzhu Li², and Kris Hauser¹
chaoqi-liu.com/scoopbot

Abstract—Predictive models can be particularly helpful for robots to effectively manipulate terrains in construction sites and extraterrestrial surfaces. However, terrain state representations become extremely high-dimensional especially to capture fine-resolution details and when depth is unknown or unbounded. This paper introduces **L-GBND**, a learning-based approach for terrain dynamics modeling and manipulation, leveraging the Graph-based Neural Dynamics (GBND) framework to represent terrain deformation as motion of a graph of particles. Based on the principle that the moving portion of a terrain is usually localized, our approach builds a large terrain graph (potentially millions of particles) but only identifies a very small active subgraph (hundreds of particles) for predicting the outcomes of robot-terrain interaction. To minimize the size of the active subgraph we introduce a learning-based approach that identifies a small region of interest (RoI) based on the robot’s control inputs and the current scene. We also introduce a novel domain boundary feature encoding that allows GBND to perform accurate dynamics prediction in the RoI interior while avoiding particle penetration through RoI boundaries. Our proposed method is both orders of magnitude faster than naive GBND and it achieves better overall prediction accuracy. We further evaluated our framework on excavation and shaping tasks on terrain with different granularity.

I. INTRODUCTION

Terrain manipulation is essential in construction industry and outer space exploration [1], [2]. Autonomous robots that can create desired structure in the construction site by interpreting sensory data and making decision on where and how to manipulate the terrain, can greatly improve efficiency and be deployed to place where lacks human labor, e.g., other planets. Prediction is a key challenge in terrain manipulation, since robots need to predict outcomes of excavated volume and the movement of terrain shape after scooping, dumping, or pushing material. People have traditionally built approximate models using analytical expressions and geometric calculations to predict excavation outcomes. However, these models are often limited in adapting to other action types and typically require predefined parameters [3]. Simulation-based methods use terramechanics physics simulators which can handle a wider variety of situations, but these are extremely computationally expensive and still require solving difficult parameter identification problems [4], [5], [6]. Learning-based approaches have also been investigated since early work in autonomous excavation [7], but recent attention has been focused on deep learning techniques, such as image-based prediction of granular media manipulation [8]. A

¹University of Illinois at Urbana-Champaign ²Columbia University. This work was partially funded by NSF Grant #FRR-2409661. Corresponding author: chaoqi12@illinois.edu.




Fig. 1: **Evaluation platforms.** Our terrain manipulation system includes a UR5e robotic arm equipped with an overhead RGB-D camera. We are also extending this system to a mobile scooping platform for in-the-wild studies.

line of recent work has studied the use of Graph-Based Neural Dynamics (GBND), a technique that uses Graph Neural Networks (GNNs) to model system dynamics of a scene represented as set of 3D particles, with each particle representing a small volume of matter [9], [10], [11]. GBND has been applied to various nontrivial deformable object manipulation tasks [12], [13], [14], and can generalize well across different viewpoints, objects and robots poses, shapes and sizes. However, all prior such works have addressed table-top scenarios and objects with bounded volume that can be observed through static cameras.

In this work, we propose a framework, *Localized Graph-Based Neural Dynamics (L-GBND)*, for modeling deformable terrains at fine resolutions over large spatial scales, such as a construction site in which a mobile excavator uses 3D mapping to generate a terrain representation. In particular, our method extends the GBND to handle spatially large-scale and even unbounded terrains which may be represented by graphs of millions of particles that cannot fit in GPU memory. To predict the deformation of the terrain under a candidate interaction in a computationally efficient manner suitable for planning, our method restricts the attention of dynamics prediction to a dynamically-selected local neighborhood. We introduce a learning-based 3D region-of-interest (RoI) proposer that identifies a small subset of particles from the large graph that are likely to move. Assuming particles that are outside the RoI stay static during robot-terrain interaction improves speed by orders of magnitude in addition to higher dynamics prediction accuracy. **L-GBND** simultaneously learns the GBND dynamics and the RoI from

simulation data, and both are fine-tuned on real data. Moreover, we introduce a novel boundary feature representation to address the problem that naïvely applying GBND to the identified subgraph leads to erroneous predictions. In our experiments, we verify both the efficiency and effectiveness of **L-GBND** for manipulating terrains of different materials.

Our key innovations include: (1) learning graph dynamics and RoI to minimize **L-GBND** memory usage and computation time while retaining high accuracy; and (2) introducing novel node features that enable learning **L-GBND** over arbitrary volumetric domains, ensuring that accuracy is preserved at the boundary of RoI, and that translation invariance is maintained.

II. RELATED WORK AND DISCUSSION

Terrain Modeling and Manipulation. Effective modeling of terrain movement induced by external forces, and its use in manipulation and locomotion, has been studied for years. Resistive Force Theory (RFT) and 3D-RFT, for example, is generic but reduced-order force models widely used in modeling locomotion in granular media [15]. However, RFT lacks the ability to model movements where inertial effects become significant, in addition to assumption of over-simplified geometry [7]. Other work instead use heightmap as the representation, and regress a latent linear dynamics model [16] for bulldozing. Linear dynamics models lack capability for more complicated interactions [8] and discretized heightmaps lack the flexibility of change of resolution, while the optimal resolution varies across tasks and underlying material [17]. In addition, particle-based representation provides high-fidelity geometry for tool equipped, focusing more on geometry during interaction, which provides extra precision. Alternatively, model-free methods are utilized widely [18], but these approaches often grapple with high sample complexity and exhibit a task-specific nature.

Dynamics Model Learning. Physics-based models have shown effectiveness in a wide range of robotic manipulation tasks [19], [20], [21], but accurate model building requires full state estimation. To address this issue, recent advances in machine learning are applied to learn the system dynamics models directly from sensory data [22], [23], [24]. The learned dynamics model generally demonstrates superior effectiveness relative to its model-free counterparts [25]. Graph-based scene representation processed by GNNs have showcased effectiveness in dynamics modeling due to the spatial relational inductive bias injected [26], [27], [9], [10], [11]. Prior work explored a wide range of applications of graph-based dynamics model, including but not limited to rigid bodies [9], [10], elastic-plastic materials [13], [12], fluids [9], [11], and granular materials [17], [11], [28]. However, the vast majority of prior work model systems that are either confined to a 3D box or atop a 2D plane, and our work demonstrates an approach to extend graph-based approaches to terrains of unbounded breadth and depth using regions of interest and RoI boundary features.

Model-Free Policy Learning. While model-free methods such as reinforcement learning (RL) and behavior cloning

(BC) have achieved notable success in robotic manipulation [29], [30], [31], [32], they typically suffer from limited generalization, particularly in environments involving complex, high-dimensional, and long-horizon interactions like terrain deformation. These approaches require extensive task-specific data collection and tuning, often making them impractical for settings involving diverse or previously unseen terrain geometries. In the context of deformable terrain manipulation, the inability of model-free policies to extrapolate beyond the training distribution presents a fundamental limitation [33]. While such methods can learn reactive strategies within the training distribution, they lack an explicit mechanism for generalizing to new goals or terrain states. In contrast, dynamics model-based approaches retain predictive structure: given a learned forward model, planning can be performed zero-shot for unseen terrain states by simulating the effects of actions. This capability enables flexible and task-agnostic generalization, which is essential for terrain-aware manipulation. Because of these limitations, we do not consider model-free policies as baselines in this work. Our focus is exclusively on dynamics model learning, as it enables both general-purpose planning and scalability across diverse terrain interaction scenarios.

III. LOCALIZED GRAPH-BASED NEURAL DYNAMICS

We present a framework for efficient and accurate terrain dynamics prediction under robot interaction, based on *localized* GBND (**L-GBND**) models. Our goal is to learn a computationally-efficient model to predict the shape of a terrain after robot-terrain interaction, such that **L-GBND** can be used in terrain manipulation planning. Specifically, letting x_t denote the state of the terrain, we model $\hat{x}_{t+1} = f(x_t, u_t)$ where u_t is the robot control and t is the time index. We learn fine-grained prediction where the time step is 0.1 sec and u_t indicates a small robot end-effector displacement. The dynamics model is trained on data recorded during three types of actions (scooping, pushing, and dumping), and we use a unified dynamics model for all actions. Pretraining is performed from simulation data and then the model is fine-tuned on real data.

We choose **L-GBND** as our terrain model, which is a particle-based scene representation that provides flexibility in inter-particle interactions, an inductive bias toward locality in interactions, mass conservation, and preserves fine details of deforming solids and fluids [9], [11], [10]. To manage the high-dimensional state space of the terrain, we assume that only a small region of interest (RoI) $R_t(u_t) \subset \mathbb{R}^3$ of the volume is affected by the interaction. Letting $p_t^{(i)}$ denote the i^{th} particle in the terrain at time t , its future state $\hat{p}_{t+1}^{(i)} = p_t^{(i)}$ for any point $p_t^{(i)} \notin R_t(u_t)$. Our method simultaneously learns f and R_t . The RoI proposer is formulated as a super-level set of a learned function $g(p, u_t)$ such that $g(p, u_t) \geq 0$ if and only if $p \in R_t(u_t)$. Using the RoI, we only predict the motion of a subgraph of the particle graph. The proposed framework is shown in Fig. 2.

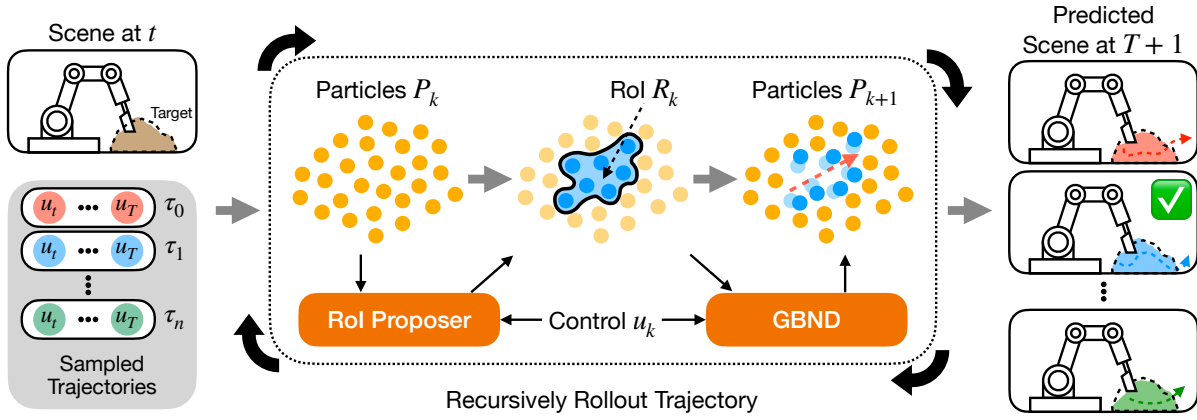


Fig. 2: **System overview.** The proposed **dynamics** model is shown in the central block: scene particles and the control input are fed into the RoI proposer, which selects particles likely to move based on the interaction caused by the control input. This is followed by a rollout using GBND. When integrated with **planning**, multiple trajectories are sampled and rolled out in parallel, and the best trajectory, determined by a pre-defined metric, is executed.




Fig. 3: **Visualization of a typical scooping simulation.** Our method (left) only predicts the dynamics of the highlighted particles in the predicted RoI. Predictions are computed for an order of magnitude fewer particles as the full graph (right) while retaining similar accuracy.

A. Graph-Based Scene Representation

We adopt a particle-based representation of the environment, in which both terrain and robot end-effectors are modeled as sets of particles. This representation is particularly suitable for learning physical dynamics in deformable and granular media, as each particle implicitly represents a small volume of matter. Terrain particles $P_t = \{p_t^{(i)}\}_{i=1}^{N_p}$ are instantiated from the observed terrain surface map, volumetrically down to a sufficiently large depth. Tool particles

$T_t = \{p_t^{(i)}\}_{i=N_p+1}^{N_p+N_t}$ are sampled on the robot’s end-effector. For each particle, we connect it to the k closest particles that meet a distance threshold using a nearest-neighbor algorithm, resulting in a sparse graph.

Follow prior work [10], [9], [11], we can define nodes in the graph as

$$v_t^{(i)} \triangleq \left[\mathbf{v}_{t-H:t-1}^{(i)}, u_t^{(i)}, c^{(i)} \right]^\top, \quad (1)$$

where $\mathbf{v}_{t-H:t-1}^{(i)}$ is particle $p^{(i)}$ ’s velocity over the last H time steps, $u_t^{(i)}$ is the impulse on the tool particle directly caused by robot action (displacement of the end-effector, hence, naturally zeros for terrain particles), and $c^{(i)}$ is the particle class (terrain or tool) embedding. The edge features for an edge $e_t^{(k)}$ connecting particles i and j are

$$e_t^{(k)} \triangleq \left[p_{t-H:t}^{(i)} - p_{t-H:t}^{(j)}, c^{(i)}, c^{(j)} \right]^\top. \quad (2)$$

B. Graph-Based Neural Dynamics

The dynamics of the particle system are modeled using a Graph Neural Network (GNN), which represents the environment state as a directed graph $x_t \equiv G_t(V_t, E_t)$ at each time step, where $V_t \triangleq \{v_t^{(i)}\}_{i=1}^{|V_t|}$ is the set of nodes, and $E_t \triangleq \{e_t^{(k)}\}_{k=1}^{|E_t|}$ is the set of edges, with each edge $e_t^{(k)}$ denoting a potential direct interaction from $v_t^{(i)}$ to $v_t^{(j)}$. Indirect interactions between particles are modeled in this framework using a message passing algorithm.

The dynamics are defined as follows. Each node and edge are first encoded into a latent representation,

$$z_0[v_t^{(i)}] = f_V^{\text{enc}}(v_t^{(i)}), \quad z_0[e_t^{(k)}] = f_E^{\text{enc}}(e_t^{(k)}). \quad (3)$$

Then, message propagation is run for L iterations to propagate, aggregate, and update latent representations across the graph. Mathematically, for $l \in \{0, 1, \dots, L-1\}$,

$$z_{l+1}[e_t^{(k)}] = f_E^{\text{prop}}(z_l[e_t^{(k)}], z_l^i, z_l^j) \quad (4)$$

for an edge $e_t^{(k)} \triangleq (i, j)$, and messages along edges to nodes are propagated and aggregated via

$$z_{l+1}[v_t^{(i)}] = f_V^{\text{prop}}\left(z_l[v_t^{(i)}], \sum_{e \in E^+(v_t^{(i)})} z_l[e]\right). \quad (5)$$

Algorithm 1 L-GBND

```

procedure PROPOSE( $\{p_{t-H:t}^{(i)}\}_{i=1}^N, u_t$ )
   $M_{t-H:t} \leftarrow \text{PROJECT2D}(\{p_{t-H:t}^{(i)}\}_{i=1}^N)$ 
  Predict  $\tilde{g}$  based on  $M_{t-H:t}, u_t$   $\triangleright R_t$  in Sec. III-C
  for  $i \in [N]$  do
    if  $p_t^{(i)}.z \geq \tilde{g}(p_t^{(i)}.x, p_t^{(i)}.y)$  then
       $indices \leftarrow indices \cup \{i\}$ 
  return  $indices$ 

procedure STEP( $x_{t-H:t}, u_t$ )
  Instantiate scene particles  $\{^W p_{t-H:t}^{(i)}\}_{i=1}^N$ 
  Align coordinate frame  $L$  with scoop-tip
  Localize particles  $^L p_{t-H:t}^{(i)} \leftarrow L^T W^W p_{t-H:t}^{(i)}$ 
   $indices \leftarrow \text{PROPOSE}(\{^L p_{t-H:t}^{(i)}\}_{i=1}^N, u_t)$ 
  Rollout dynamics  $\forall i$ ,
  
$$^L \hat{p}_{t+1}^{(i)} = \begin{cases} f_\theta(^L p_{t-H:t}^{(i)}, u_t^{(i)}) & \text{if } i \in indices \\ ^L p_t^{(i)} & \text{otherwise} \end{cases}$$

  Restore particles  $^W p_{t+1}^{(i)} \leftarrow W^T L^L \hat{p}_{t+1}^{(i)}$ 
  return  $\{^W \hat{p}_{t+1}^{(i)}\}_{i=1}^N$ 

```

Algorithm 2 MPPI Planning with L-GBND

Require: score S , scene x_t , goal description G , motion parameter space Ω , batch size b , time budget B

```

 $u_{0:T}^* \leftarrow nil$ 
while time used <  $B$  do
  Sample a batch of trajectory parameters  $\omega_1, \dots, \omega_b$ 
  Convert batch of  $\omega_i$  to control sequences  $u_{i,0:T_i}, \forall i$ 
  Predict batch  $\hat{x}_{i,t+T_i} \leftarrow \text{Rollout}(x_t, u_{i,0:T_i}, \forall i)$ 
  Evaluate scores  $S(\hat{x}_{i,t+T_i}, G), \forall i$ 
  Update best control sequence  $u_{0:T}^*$ 
return  $u_{0:T}^*$ 

```

Finally, a node decoder f_V^{dec} predicts displacement of the particle, $\Delta p_t^{(i)} = f_V^{\text{dec}}(z_L[v_t^{(i)}])$. The ultimate dynamics prediction is given by $\hat{p}_{t+1}^{(i)} = p_t^{(i)} + \Delta \hat{p}_t^{(i)}$. GBND learns encoder, decoder, and propagation networks $f_V^{\text{enc}}, f_E^{\text{enc}}, f_V^{\text{dec}}, f_V^{\text{prop}}, f_E^{\text{prop}}$ via backpropagation through message passing.

C. Adaptive Fine-Grained Region of Interest Identification

GBND’s computational and memory costs scale with the number of particles and edges; therefore, constraining the graph size is critical for efficiency. We learn an implicit function $g(p, u_t)$ that defines the region of interest (RoI) $R_t(u_t) = \{p \in \mathbb{R}^3 \mid g(p, u_t) \geq 0\}$.

We first extract the planar tool position and yaw from u_t and define a 2D reference frame T_{tool} . We then predict a tool-centered 2D height function $\tilde{g}(x, y)$ such that $g(p, u_t) = p_z - \tilde{g}(T_{\text{tool}}^{-1}(p_x, p_y))$. In other words, points whose z exceeds $\tilde{g}(x, y)$ (in the tool-centered frame) are included in the RoI.

Our approach uses a terrain heightmap and instantiates \tilde{g} as a Convolutional Neural Network (CNN) architecture. A heightmap with grid resolution 2.5 cm is centered about the tool and cropped to a 0.4 m \times 0.4 m region. At each time step, particles in the bounding box are squeezed into a 32 \times 32 heightmap, with each pixel concatenated with a 16D latent representation of the current control input. Our CNN is a U-Net-like [34] architecture with three levels of pooling, nonlinearity, and upsampling steps with skip connections. The output is a 16 \times 16 depth map representation of \tilde{g} .

	GBND	GBND w/ z features	L-GBND
In Distribution	1.85	1.07	1.09
Out of Distribution	1.85	2.85	1.09

TABLE I: Comparison of node feature encodings for boundaries. GBND fails to generalize, and adding absolute height (z) overfits to training terrain elevation. **L-GBND** with normal features performs consistently across distributions.

Training uses simulation data. Moving particles are identified using a velocity threshold, and connected component filtering is applied to remove noise. The target heightmap is defined by combining the upper heightmap of non-moving particles (inflated to produce a bowl shape) with the bottom-up heightmap of moving particles. An L_1 loss is applied to non-empty cells, while a hinge loss at 0.5 m is used for empty cells. The pseudocode of **L-GBND** with scale reduction is presented in Algo. 1, and a typical simulation rollout with adaptive RoI scale reduction is presented in Fig. 3.

D. Boundary-Aware Node Features

GBND is best suited for tasks with similar training / test distributions, but our setting introduces the complication of unboundedness and varying regions of interest. As an example, past works using GBND for fluid and pile manipulation have used a fixed domain of particles, such that the network could *implicitly* learn the boundary conditions of a floor and walls [13], [12], [35], [17]. In our case, “boundary conditions” are not strictly-enforced boundaries, but are rather just more material that is unaffected by the current interactions due to energy dissipation.

In effect, particles adjacent to the RoI boundary should “feel” resistance as though the boundary were fixed. With naive node feature encoding (Equ. 1) as used in prior work, these particles will simply pass through the boundary as if there was nothing there. Alternatively, we experimented with adding “wall particles”, a boundary set of nodes, directly connected to nodes in the RoI, whose latent features are frozen. Not only does this increase memory usage due to massive wall particles, but in our experiments, inner particles were able to penetrate gaps in the sampled wall particles which introduces unnatural artifacts.

To encode appropriate boundary behavior, our approach incorporates an estimated geometric normal of particles within RoI R_t [36]. In the interior, the normal information is not likely to help much, but normals are likely to be strongly aligned near the boundary. This allows our network to encode the effect that nodes near the RoI boundary should “feel” resistance in the direction of the normal. This encoding also has the favorable effect of translational invariance, which is that for the overall dynamics model f ,

$$\forall \Delta p \in \mathbb{R}^3, f(p_t^{(i)} + \Delta p, u_t) = f(p_t^{(i)}, u_t) + \Delta p. \quad (6)$$

This ensures that the model’s predictions are independent of the absolute positions of particles.

Experiments in Table I present validation losses of learned GBND models for different sets of node features. As baselines we compare standard GBND encoding as well as augmenting with particle height (z features). In the *in distribution* case,

the training and validation sets used terrains whose particles were sampled in a fixed z range of [0 cm, 5 cm]. In the *out of distribution* case, the validation particle sets were randomly shifted in the z axis by randomly sampled offsets from -100 to 100 m. Encoding the z value significantly improves loss in the in-distribution case, demonstrating that the network is implicitly learning a “floor effect”. This implicit relationship is detrimental in the out-of-distribution case. In contrast, our **L-GBND** with proposed normal encoding leads to better generalization than the baseline in both regimes.

E. Sim-to-Real Transfer

To transfer the pretrained **L-GBND** to the real world, we allow the robot to autonomously collect its own interaction data through self-play. Real-world trajectories are monitored using an overhead RGB-D camera, from which scene point clouds are reconstructed. Robot proprioception is used to remove robot particles and isolate in-tool particles by re-sampling within the tool’s convex hull. Since real-world data lacks explicit particle correspondences across time, we follow standard practices in the **GBND** literature and fine-tune the model using Chamfer Distance [37] and Earth Mover’s Distance [38] between predicted and observed point clouds. This enables effective learning without requiring explicit tracking of individual particles across frames.

F. Terrain Manipulation Planning with Localized **GBND**

We use our **L-GBND** in a planner for terrain manipulation tasks. Our planner selects among parameterized scooping, pushing, and dumping trajectories to minimize the difference between the predicted terrain shape and a target heightmap. For scooping and pushing, we use the penetrate-drag-scoop (PDS) trajectory used in excavation [1], [39], [7], which maintains a 0 roll angle. This yields a 9D motion parameter space. Dumping uses a 6D parameterization (dumping location, yaw angle, pre- and post-dump pitch angles). Scooping and pushing are only allowed if the scoop is empty. The planner is based on Model Predictive Path Integral (MPPI) control [40]. A batch of trajectories is sampled in each iteration. The corresponding control sequences are rolled out with the dynamics model, and the outcomes are scored based on proximity to the goal shape. The best trajectory is selected and executed, and the terrain state is updated. Pseudocode is shown in Algo. 2.

IV. EXPERIMENTS AND RESULTS

The computational performance and accuracy of the proposed method were thoroughly evaluated and contrasted against simpler baselines. Furthermore, the capability of our method was assessed through its application to extended-range excavation tasks. All experimental evaluations were conducted using an Intel i9-13900K CPU and a single NVIDIA GeForce RTX 4080 GPU.

A. Implementation Details

The architectures of **GBND** are MLPs, trained over 200,000 backpropagation steps, unless stated otherwise. For data




Fig. 4: Rollouts (left to right) from a learned 2D CNN heightmap-based dynamics model compared to **L-GBND** and ground truth. In the CNN, fine-grained features are lost due to the compressed representation and smoothing bias of CNNs. In contrast, **L-GBND** preserves volumetric structure and local interactions for 37% lower prediction error on a test dataset.

collection during pre-training, the SAPIEN simulator [41] was utilized. Attempts to employ Nvidia Flex [42] were made, but it was abandoned due to the increasing prevalence of unrealistic behaviors as the particle count escalated. In the terrain generation phase, the diamond-square algorithm [43] was employed to produce landscapes comprised of approximately 3,000 miniature cubes. Each cube was then populated with 5 to 10 particles, and further reduction of terrain particle sets was achieved through farthest-point sampling [44].

Approximately 1,000 robot-terrain interaction trajectories, each uniquely parameterized, were simulated. Throughout these simulations, the downsampled terrain particles were meticulously tracked to facilitate the learning of the dynamics model. Training on this simulation data involved regressing the motion trajectories of the particles using the Mean-Square Error (MSE) criterion, aided by the perfect association of scene particles. Concurrently, the CNN proposer was trained as detailed in Section III-C. The model’s forward simulation steps were set at 0.1 seconds.

B. Particle-based v.s. 2D Heightmap Representations

We compare **L-GBND** against a baseline that represents terrain using a 2D heightmap and learns dynamics directly in heightmap space. Both models are trained on the same dataset using the same control inputs. For consistency and implementation efficiency, the 2D baseline reuses the architecture of our RoI proposing CNN: a U-Net-style convolutional network that takes as input a stack of past terrain heightmaps and current control input, and autoregressively predicts the future heightmap. To ensure a fair comparison, we project the output of the particle-based model onto a heightmap before computing evaluation metrics. We evaluate prediction accuracy using the L_1 loss over the heightmap grid. Quantitatively, **L-GBND** achieves approximately 37% lower prediction error compared to the 2D heightmap model. Fig. 4 presents qualitative results comparing heightmap-based dynamics prediction and ground-truth terrain heightmaps.

This comparison highlights the critical role of representation in deformable terrain dynamics modeling. The 3D particle-based representation used in **GBND** captures volu-

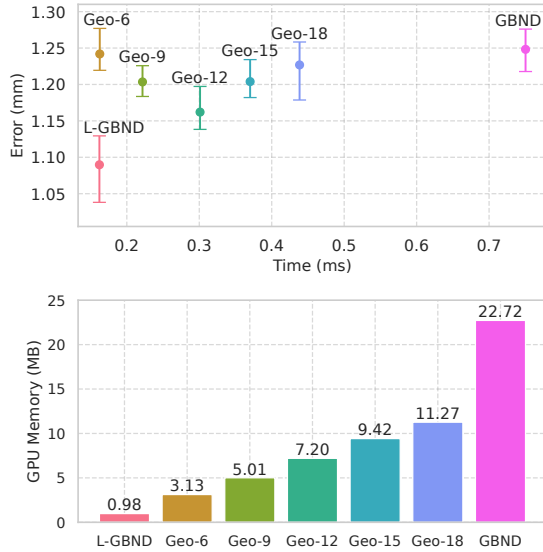


Fig. 5: Comparing our **L-GBND** against GBND and geometric region proposers with different size (Geo-X). Our method demonstrates significant advantages in speed and GPU memory. $\approx 3,000$ particles and batch size 256 are used in these experiments, per sample measurements are reported (i.e., divided by 256). Colors and labels are shared by both figures.

metric structure and encodes fine-grained geometry, providing strong inductive biases toward spatial locality, directional interaction patterns, and mass conservation. These properties enable the model to represent contact-rich interactions such as scooping, penetration, and material buildup along tool edges—phenomena that are difficult to express in 2D. In contrast, 2D heightmap representations collapse the terrain into a single surface, discarding sub-surface interactions and occluding contact geometry. Moreover, convolutional architectures tend to produce overly smoothed predictions, which dampens sharp discontinuities and local terrain features. These limitations manifest even more clearly when directly predicting future terrain shape.

C. **L-GBND**'s Efficiency and Effectiveness

The computational performance and accuracy of our method are compared against two baselines. The first baseline utilizes all particles within a large-scale terrain, while the second employs a naïve geometric Region of Interest (RoI), specifically, a cylinder centered at the scoop tip with a predetermined radius. To maintain computational feasibility with the full-terrain baseline, terrains are restricted to 3,000 particles — significantly fewer than what might be encountered by an excavator operating on an large worksite. Figure 5 presents the results on a terrain measuring $50\text{ cm} \times 50\text{ cm} \times 5\text{ cm}$, with a particle density of 0.5 particles per cm^3 . Here, ‘Geo-X’ represents the geometric RoI with a cylinder diameter of X cm. Each RoI-based approach is developed using the methodology described in Section III-C alongside our novel node boundary features.

In terms of per-sample processing time and GPU memory utilization, our method matches the speed of the smallest geometric RoI and requires even less GPU memory, while delivering superior accuracy. It is important to note that enhanced RoI prediction significantly improves the accuracy

Batch	L-GBND	Geo-6	Geo-12	Geo-18	GBND	SAPIEN
1	4.624	3.543	4.459	5.157	5.443	3.309 (CPU)
2	2.454	1.913	2.280	2.616	2.775	2.104
4	1.400	1.036	1.250	1.463	1.642	1.627
8	0.860	0.586	0.755	0.931	1.097	1.563
16	0.522	0.366	0.501	0.650	0.872	1.713
32	0.337	0.242	0.364	0.520	0.885	1.991
64	0.232	0.181	0.332	0.498	0.809	2.211
128	0.179	0.169	0.318	0.456	0.774	1.755
256	0.162	0.163	0.301	0.438	0.750	OoM
512	0.181	0.181	0.293	0.431	OoM	OoM
1024	0.198	0.198	0.294	OoM	OoM	OoM

TABLE II: Average prediction time (ms) per sample as the batch size varies, showing that 256 is an optimal batch size for our framework. For SAPIEN, 3000 cubes are simulated.

Batch	1	2	4	8	16	32	64	128	256	512	1024
# rollouts	216	407	714	1163	1916	2967	4310	5587	6173	5525	5051
$L_1(10^{-3})$	15.32	14.53	12.68	11.37	10.91	7.22	5.72	3.93	3.42	4.11	4.31

TABLE III: For given time budget (1 sec), more samples being simulated with **L-GBND** in MPPI planning leads to better planning performance.

of dynamics predictions. This is because the GBND is prone to inaccuracies, such as predicting minimal movements of static particles. An RoI introduces a strong inductive bias that is absent in GBND, leading to jitter among distant particles. For smaller values of X, Geo-X exhibits large errors due to the exclusion of moving particles. Conversely, as X increases, the inclusion of numerous static particles leads to performance degradation due to GBND jitter.

Furthermore, since the planning speed and quality are heavily influenced by the batch size, memory efficiency becomes a critical objective for our dynamics model. Table II illustrates that our method reaches optimal performance at a batch size of $b = 256$, achieving an average processing time of less than 0.2 ms per sample. Under these conditions, the planner is capable of rolling out hundreds of trajectories across numerous time steps in seconds. We also compare against the recently updated GPU-based simulator in SAPIEN 3.0 [41] in scenes with 3,000 cube-shaped particles, without a simulated robot. Our method is an order of magnitude faster and can handle larger batch sizes within available GPU memory. As a result of superior efficiency, **L-GBND** leads to better planning performance in terrain shaping tasks as shown in Table III, where we analyze the scaling of planning performance with respect to number of candidate trajectories being simulated.

D. Real World Terrain Manipulation

Real-world experiments were conducted using the platform shown in Fig. 1, left, on two different materials as depicted in Fig. 6: (1) pebbles with an approximate grain size of 0.8 - 1.0 cm, and (2) play sand with a grain size significantly less than 1 mm. To ensure experimental repeatability, the terrain was contained within a box measuring $0.9\text{ m} \times 0.6\text{ m} \times 0.2\text{ m}$, though the methodologies developed are equally applicable to natural terrains.

To evaluate real-world transfer, we fine-tune the pretrained model on 100 real-world interaction trajectories collected via autonomous self-play and overhead RGB-D sensing. Following the procedure described in Section III-E, we com-




Fig. 6: Materials used in real world experiments, U.S. quarter coin on the bottom right for scale.




Fig. 7: Fine-tuning (FT) the pre-trained (PT) L-GBND model on real-world datasets significantly improves the model stability and precision in practical scenarios, as evidenced by a reduced median error and a more compact interquartile range.

pute supervision signals using Chamfer and Earth Mover’s Distance between predicted and observed point clouds. This fine-tuning enables the model to adapt effectively to real-world sensor data and contact dynamics.

The experimental setup focused on two specific terrain shaping tasks: (1) creating a hole, and (2) forming a rectangular moat. The desired terrain configurations were represented as 2D heightmaps, with the L_1 loss between the current scene heightmap and the target heightmap serving as the performance metric. These target shapes have sharp discontinuities and cannot be physically realized due to the material’s settling behavior (angle of repose is approximately $25\text{-}40^\circ$). For each manipulation task, our planner evaluated 1500 trajectories with an average horizon of 30 steps (equivalent to 3 seconds), requiring 20-30 seconds of planning time. Qualitative results are shown in Fig. 8, demonstrating successful terrain shaping.

For a qualitative comparison, we implement a heuristic baseline that employs dynamic programming to determine the scooping trajectory that maximizes the intersection of the swept volume and the excess material in the current terrain heightmap, similar to the volume maximization approach developed by Yang et al. [45]. We also constrain the solver to use axis-aligned trajectories to boost its computational efficiency. Fig. 9 illustrates that the heuristic makes rapid initial progress but is unable to shape the terrain’s finer details by anticipating settling behavior.

V. CONCLUSION AND LIMITATION

We presented **L-GBND** for handling unbounded terrains in graph-based neural dynamics models that simultaneously learns a RoI and particle dynamics. Novel node features are key to success of this model. Experiments show that this method achieves orders of magnitude faster prediction and GPU memory usage compared with naïve GBND models or GPU-based physics simulators. It can be trained with both simulated and real data, and when used as a model for




Fig. 8: Sequences of heightmaps recorded during real-world manipulation.

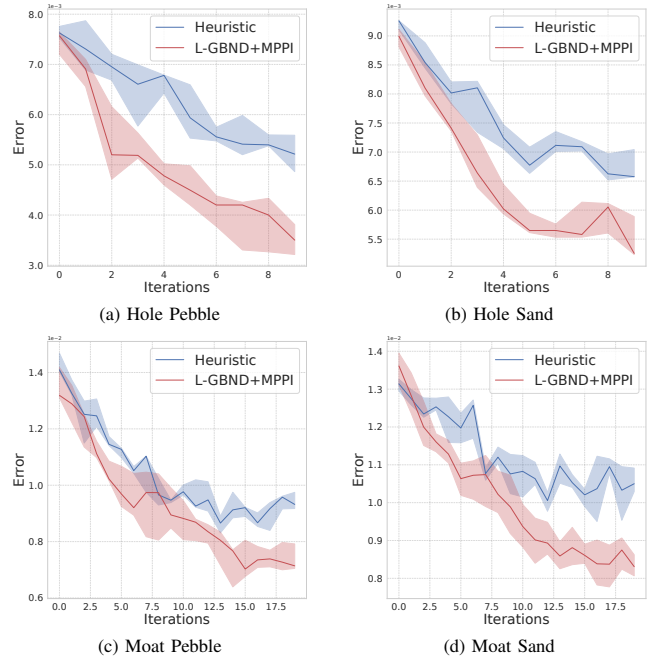


Fig. 9: Errors between the observed terrain and the target terrain after each PDS trajectory. Interquartile ranges are shaded. 10 trials per scenario.

planning our method can achieve desired target shapes on real-world excavation platforms.

There remains some limitations of our work that we would like to address in future work. First, our RoI proposer does not accurately model long-range terrain effects (e.g., landslides). Second, our current platform uses an overhead camera to update the terrain, and we would like to extend this to egocentric 3D mapping, in which our model may need to model occluded parts of the terrain. Third, our models assume terrain homogeneity, and future work should address heterogeneous materials including terrains with embedded obstacles. Lastly, the current dynamics model requires full-parameter fine-tuning when moving to unseen materials, one potential solution is to adapt physics-parameter conditioning [46] for online adaptation.

REFERENCES

- [1] Y. Zhu, P. Thangeda, M. Ornik, and K. Hauser, “Few-shot adaptation for manipulating granular materials under domain shift,” *Robotics: Science and Systems*, 2023.

- [2] R. Lee, B. Younes, A. Pletta, J. Harrington, R. Q. Wong, and W. R. Whittaker, "Cratergrader: Autonomous robotic terrain manipulation for lunar site preparation and earthmoving," *Robotics: Science and Systems*, 2024.
- [3] A. R. Reece, "Paper 2: The fundamental equation of earth-moving mechanics," *Proceedings of the Institution of Mechanical Engineers, Conference Proceedings*, vol. 179, no. 6, pp. 16–22, 1964.
- [4] H. Herrmann, "Simulation of granular media," *Physica A: Statistical Mechanics and its Applications*, vol. 191, no. 1, pp. 263–276, 1992.
- [5] N. Bell, Y. Yu, and P. J. Mucha, "Particle-based simulation of granular materials," in *Proceedings of the 2005 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, ser. SCA '05. New York, NY, USA: Association for Computing Machinery, 2005, p. 77–86.
- [6] A. Sommer, U. Schwanecke, and E. Schoemer, "Interactive high-resolution simulation of granular material," *Journal of WSCG*, vol. 30, no. 1–2, p. 9–15, 2022.
- [7] S. Sing, "Synthesis of tactical plans for robotic excavation," Ph.D. dissertation, Citeseer, 1995.
- [8] H. J. T. Suh and R. Tedrake, "The surprising effectiveness of linear models for visual foresight in object pile manipulation," *Workshop on Algorithmic Foundations of Robotics*, 2020.
- [9] Y. Li, J. Wu, R. Tedrake, J. B. Tenenbaum, and A. Torralba, "Learning particle dynamics for manipulating rigid bodies, deformable objects, and fluids," *arXiv preprint arXiv:1810.01566*, 2018.
- [10] P. W. Battaglia, R. Pascanu, M. Lai, D. Rezende, and K. Kavukcuoglu, "Interaction networks for learning about objects, relations and physics," *Conference on Neural Information Processing Systems*, 2016.
- [11] A. Sanchez-Gonzalez, J. Godwin, T. Pfaff, R. Ying, J. Leskovec, and P. W. Battaglia, "Learning to simulate complex physics with graph networks," *International Conference on Machine Learning*, 2020.
- [12] H. Shi, H. Xu, S. Clarke, Y. Li, and J. Wu, "Robocook: Long-horizon elasto-plastic object manipulation with diverse tools," *arXiv preprint arXiv:2306.14447*, 2023.
- [13] H. Shi, H. Xu, Z. Huang, Y. Li, and J. Wu, "Robocraft: Learning to see, simulate, and shape elasto-plastic objects with graph networks," *Robotics: Science and Systems*, 2022.
- [14] X. Lin, Y. Wang, J. Olkin, and D. Held, "Softgym: Benchmarking deep reinforcement learning for deformable object manipulation," *Conference on Robot Learning*, 2021.
- [15] Q. Yu, C. Pavlov, W. Kim, and A. M. Johnson, "Modeling wheeled locomotion in granular media using 3d-rft and sand deformation," *Journal of Terramechanics*, vol. 115, p. 100987, 2024.
- [16] W. J. Wagner, K. Driggs-Campbell, and A. Soylemezoglu, "Model learning and predictive control for autonomous obstacle reduction via bulldozing," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 6531–6538.
- [17] Y. Wang, Y. Li, K. Driggs-Campbell, L. Fei-Fei, and J. Wu, "Dynamic-resolution model learning for object pile manipulation," *Robotics: Science and Systems*, 2023.
- [18] J. Xu, Y. Jia, D. Yang, P. Meng, X. Zhu, Z. Guo, S. Song, and M. Ciocarlie, "Tactile-based object retrieval from granular media," *arXiv*, 2024.
- [19] T. Pang, H. T. Suh, L. Yang, and R. Tedrake, "Global planning for contact-rich manipulation via local smoothing of quasi-dynamic contact models," *IEEE Transactions on robotics*, 2023.
- [20] K.-T. Yu, M. Bauza, N. Fazeli, and A. Rodriguez, "More than a million ways to be pushed. a high-fidelity experimental dataset of planar pushing," in *2016 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2016, pp. 30–37.
- [21] F. R. Hogan and A. Rodriguez, "Feedback control of the pusher-slider system: A story of hybrid and underactuated contact dynamics," in *Algorithmic Foundations of Robotics XII: Proceedings of the Twelfth Workshop on the Algorithmic Foundations of Robotics*. Springer, 2020, pp. 800–815.
- [22] A. Nagabandi, K. Konolige, S. Levine, and V. Kumar, "Deep dynamics models for learning dexterous manipulation," in *Conference on Robot Learning*. PMLR, 2020, pp. 1101–1112.
- [23] H. Hu, F. Qian, and D. Seita, "Learning granular media avalanche behavior for indirectly manipulating obstacles on a granular slope," *Conference on Robot Learning*, 2024.
- [24] L. Wang, K. Zhao, C. Liu, and X. Chen, "Learning real-world action-video dynamics with heterogeneous masked autoregression," *arXiv*, 2025.
- [25] Z. Liu, G. Zhou, J. He, T. Marcucci, L. Fei-Fei, J. Wu, and Y. Li, "Model-based control with sparse neural dynamics," *Conference on Neural Information Processing Systems*, 2023.
- [26] P. W. Battaglia, J. B. Hamrick, V. Bapst, A. Sanchez-Gonzalez, V. Zambaldi, M. Malinowski, A. Tacchetti, D. Raposo, A. Santoro, R. Faulkner, et al., "Relational inductive biases, deep learning, and graph networks," *arXiv preprint arXiv:1806.01261*, 2018.
- [27] T. Pfaff, M. Fortunato, A. Sanchez-Gonzalez, and P. W. Battaglia, "Learning mesh-based simulation with graph networks," *arXiv preprint arXiv:2010.03409*, 2020.
- [28] N. Tuomainen, D. Blanco-Mulero, and V. Kyrki, "Manipulation of granular materials by learning particle interactions," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, p. 5663–5670, Apr. 2022.
- [29] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song, "Diffusion policy: Visuomotor policy learning via action diffusion," *Robotics: Science and Systems*, 2024.
- [30] Y. Du, M. Yang, B. Dai, H. Dai, O. Nachum, J. B. Tenenbaum, D. Schuurmans, and P. Abbeel, "Learning Universal Policies via Text-Guided Video Generation," *Conference on Neural Information Processing Systems*, Nov. 2023, arXiv:2302.00111 [cs].
- [31] M. Janner, Y. Du, J. B. Tenenbaum, and S. Levine, "Planning with diffusion for flexible behavior synthesis," in *International Conference on Machine Learning*, 2022.
- [32] J. Carvalho, A. T. Le, M. Baierl, D. Koert, and J. Peters, "Motion Planning Diffusion: Learning and Planning of Robot Motions with Diffusion Models," Aug. 2023, arXiv:2308.01557 [cs].
- [33] N. Rajaraman, L. Yang, J. Jiao, and K. Ramchandran, "Toward the fundamental limits of imitation learning," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 2914–2924.
- [34] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," *Medical Image Computing and Computer Assisted Intervention*, 2015.
- [35] B. Ai, S. Tian, H. Shi, Y. Wang, C. Tan, Y. Li, and J. Wu, "Robopack: Learning tactile-informed dynamics models for dense packing," *Robotics: Science and Systems*, 2024.
- [36] F. Tombari, S. Salti, and L. Di Stefano, "Unique signatures of histograms for local surface description," in *Computer Vision – ECCV 2010*, K. Daniilidis, P. Maragos, and N. Paragios, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 356–369.
- [37] H. Fan, H. Su, and L. J. Guibas, "A point set generation network for 3d object reconstruction from a single image," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 605–613.
- [38] Y. Rubner, C. Tomasi, and L. J. Guibas, "The earth mover's distance as a metric for image retrieval," *International journal of computer vision*, vol. 40, pp. 99–121, 2000.
- [39] Q. Lu and L. Zhang, "Excavation learning for rigid objects in clutter," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 7373–7380, 2021.
- [40] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Aggressive driving with model predictive path integral control," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 1433–1440.
- [41] F. Xiang, Y. Qin, K. Mo, Y. Xia, H. Zhu, F. Liu, M. Liu, H. Jiang, Y. Yuan, H. Wang, L. Yi, A. X. Chang, L. J. Guibas, and H. Su, "SAPIEN: A simulated part-based interactive environment," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [42] M. Macklin, M. Müller, N. Chentanez, and T.-Y. Kim, "Unified particle physics for real-time applications," *ACM Transactions on Graphics (TOG)*, vol. 33, no. 4, pp. 1–12, 2014.
- [43] A. Fournier, D. Fussell, and L. Carpenter, *Computer rendering of stochastic models*. New York, NY, USA: Association for Computing Machinery, 1998, p. 189–202.
- [44] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "Pointnet++: Deep hierarchical feature learning on point sets in a metric space," *Conference on Neural Information Processing Systems*, 2017.
- [45] Y. Yang, P. Long, X. Song, J. Pan, and L. Zhang, "Optimization-based framework for excavation trajectory generation," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1479–1486, 2021.
- [46] K. Zhang, B. Li, K. Hauser, and Y. Li, "Adaptigraph: Material-adaptive graph-based neural dynamics for robotic manipulation," *Robotics: Science and Systems*, 2024.