Granular loco-manipulation: Repositioning rocks through strategic sand avalanche

Haodi Hu[1], **Yue Wu**[1], **Tian Xie**[2], **Daniel Seita**[1,†], **Feifei Qian**[1,†]
[1]University of Southern California, United States
[2]Harvey Mudd College
[†]Co-advising
haodihu@usc.edu

***Abstract*—Legged robots have the potential to leverage obstacles such as rocks and boulders to climb steep sand slopes. However, efficiently repositioning these obstacles to desired locations remains challenging. Here we present DiffusiveGRAIN, a learning-based method that enables a multi-legged robot to strategically induce localized sand avalanches using their legs, indirectly manipulating obstacle positions. Using a laboratory granular trackway, we perform 375 loco-manipulation experiments with varied obstacle distances, robot orientation, and robot actions. DiffusiveGRAIN includes a diffusion-based "environment predictor" to capture multi-obstacle movements under granular flow interferences. In addition, we develop a "robot state predictor" to estimate changes in robot state from various leg action patterns. Deployment experiments (40 trials) demonstrate that by integrating the environment and robot state predictors, a multi-legged robot can autonomously plan its leg movements based on loco-manipulation goals, successfully shifting closely located rocks and boulders to desired locations in over 70% of trials. Our study showcases the potential for multi-robot teams to collaboratively manipulate obstacles to achieve improved mobility on challenging terrains.**

***Index Terms*—Granular Medium, Loco-manipulation, Diffusion Model**

I. Introduction

Natural environments contain deformable sand, steep inclines, and large rocks and boulders, which present significant challenges for terrestrial robot locomotion. Inspired from mountain goats that can push against rocks to climb up steep cliffs, and snakes that can utilize grass stems to reduce slippage on soft sand [18, 19], recent robotics research has developed "obstacle-aided" strategies [7, 8, 12, 14, 16, 20] for robots to utilize interactions and collisions with large rocks and boulders to improve mobility on complex terrains. While obstacle-aided locomotion offers promising opportunities for robots [2, 2, 7, 15] to negotiate challenging terrains, these strategies often rely on specific leg-obstacle contact positions [13, 16]. As a result, their effectiveness depends heavily on the availability and spatial distribution of rocks and boulders, which can vary unpredictably across natural terrains.

To address this challenge, we propose a novel approach, DiffusiveGRAIN, for robot loco-manipulation where a U-Net [17] and a diffusion model [21] (using the noise scheduler from [6]) takes as *input* a set of spatially-aligned image representations of the environment and robot action, and *predicts* the environmental changes on the sand slope. This prediction involves reasoning about the obstacle movement atop granular slope as well as the robot state change caused by the robot's leg excavation during the robot's actions. See Fig. 1 for an overview. During experiments, we use a receding horizon prediction to find the robot's optimal action toward the robot task.

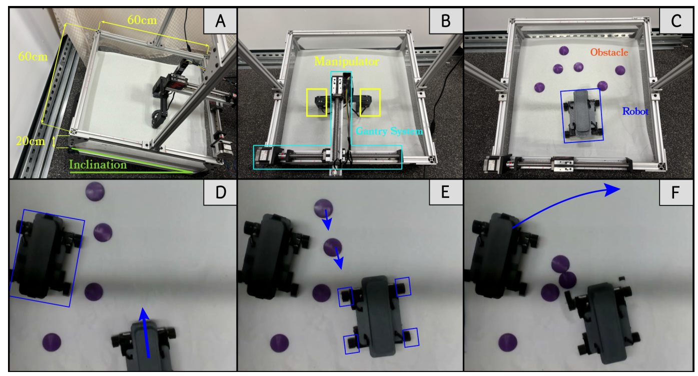

Fig. 1. Experiment environment. (A) shows the side view of the granular trackway with an inclination angle of $\Phi = 20$ degrees; (B) shows the granular trackway with two robotic legs mounted on an actuated gantry system; (C) shows the robot (*not* the manipulator in (B)) in the granular trackway, and 3D-printed obstacles (purple semi-spheres); (D, E, F) illustrates an example of a multi-robot collaborative obstacle manipulation scenario to demonstrate potential application. In (D), robot 1 (to the left, indicated with box) is trapped by surrounding obstacles with undesired locations that cannot aid its locomotion; (E) shows robot 2 (at the bottom) using its legs to trigger a localized sand avalanche, moving the obstacles around the trapped robot to free the trapped robot; (F) shows that the obstacle manipulation frees the (previously) trapped robot 1, allowing it to turn clockwise and move towards desired destination (in the direction of the curved arrow).

II. DiffusiveGRAIN: Learning-based Approach to Predict Obstacle Movement

A. Environment State Predictor using Diffusion

In this work, we built upon a prior method, GRAIN [9], which used a Vision Transformer (ViT) [3] to process image representations of granular dynamics and robot excavation actions. GRAIN was found to work well for single leg, single obstacle manipulation, yet did not generalize well for multi-obstacle scenarios when the obstacles are located adjacently, and did not consider robot state change when predicting and planning manipulation actions.

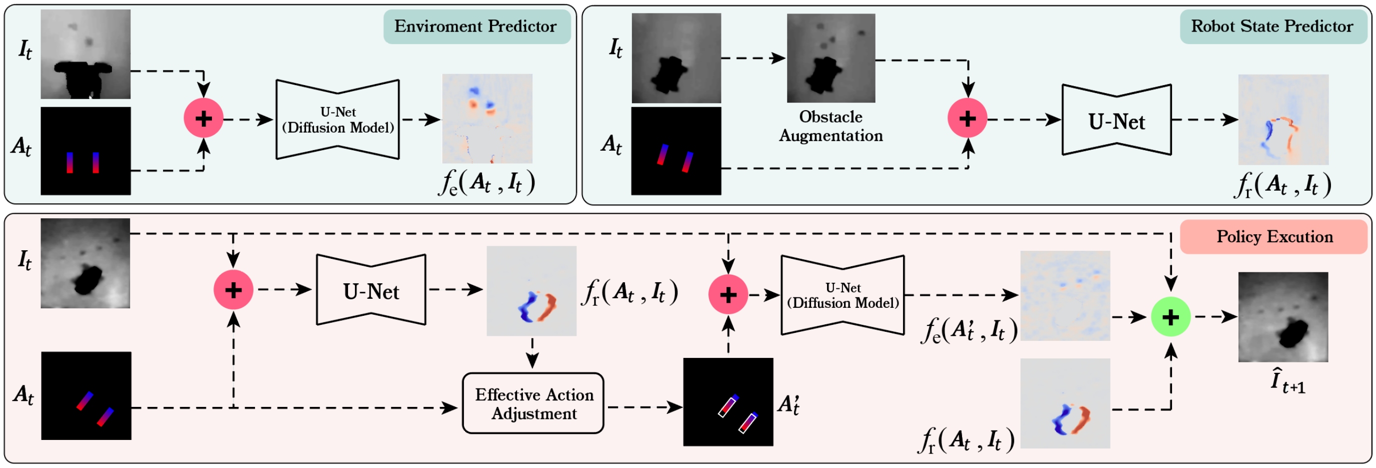

Fig. 2. System overview. The *environment predictor* f_e uses a diffusion model (with a U-Net backbone) to predict the depth change of the environment given the depth image and action. The *robot state predictor* f_r uses a U-Net to predict the robot state change given the robot state and action. During *policy execution*, given the predicted robot state change, we introduce an "Effective Action Adjustment (EAA)" to compensate for the affected leg-granular-media interaction by the robot state change. We then combine the updated robot action image with the depth image to the trained diffusion model and get the predicted depth image change. We combine this with the predicted robot state and the original depth image to get the predicted next depth image. The red addition symbols represent channel-wise image concatenation operation, and the green addition symbols represent an image combination method.

To address these deficits, we propose DiffusiveGRAIN, a learning-based method that enables multi-legged robots to reposition densely-distributed rocks and boulders on sand slopes to desired locations during its locomotion. The key feature of DiffusiveGRAIN is its novel granular media dynamics predictor that learns both obstacle and robot movement under leg excavation actions.

Our environment state predictor f_e consists of a U-Net backbone [17] diffusion model. The inputs are the depth image $\mathbf{I}_t$ and an RGB image $\mathbf{A}_t$ representing the robot action. For $\mathbf{A}_t$, we use a space-aligned gradient color region from blue to red to represent the robot leg interaction area with the granular slope. The color region's length is the effective action length (12.0 cm) and its width is the robot leg width (1.5 cm). We train f_e to predict the change in the depth image of the granular slope surface $f_e(\mathbf{I}_t, \mathbf{A}_t)$. The predicted image is converted to a grayscale image, which we can then add to the original input depth image to get the predicted depth image of the environment for the *next* state: $\mathbf{I}_t + f_e(\mathbf{I}_t, \mathbf{A}_t)$.

B. Robot State Predictor

We use a second U-Net, f_r, to predict the robot state change given its action and the environment state. As in f_e, the inputs are the depth image $\mathbf{I}_t$ and the RGB image action representation $\mathbf{A}_t$. During training, we use OpenCV code [1] to augment the input data to f_r by adding obstacles to the depth image of the collected dataset while keeping the same label. As a result, the U-Net learns from the input depth images with extra obstacles. We use this method based on our observation during the experiments that the obstacles do not noticeably affect robot's state change unless the robot leg directly contacts obstacles. The output is the predicted robot state change $f_r(\mathbf{I}_t, \mathbf{A}_t)$. We can similarly convert this to a grayscale and obtain the predicted depth image representing the robot, $\mathbf{I}_t + f_r(\mathbf{I}_t, \mathbf{A}_t)$, for the *next* state.

C. Effective Action Adjustment (EAA)

We observe that *robot* action-triggered sand avalanche behavior can be significantly different compared to the sand avalanche behavior triggered by the *manipulator* with the same excavation action. The reason is that robot leg excavation actions can lead to a different amount of advancement or slippage in each leg, resulting in significant changes to the robot state during excavation. We propose an Effective Action Adjustment (EAA) method to compensate for this prediction error. Based on the robot action, we know there are two leg-sand interaction events in a full rotation of a robot leg, and the robot state change is because the robot leg rotation provides a robot propulsion and rotation force to change its position and orientation. The robot has an initial state $\mathbf{x}_0$ and during the first leg-sand interaction the robot state changes to $\mathbf{x}_1$ and later changes to $\mathbf{x}_2$ during the second leg-sand interaction. The EAA assumes the leg excavation action triggered sand avalanche during the robot transition from $\mathbf{x}_0$ to $\mathbf{x}_2$ is equal to the leg excavation action triggered a sand avalanche at the fixed robot state $\mathbf{x}_1$. To get $\mathbf{x}_1$ in the DiffusiveGRAIN policy execution stage, we assume $\mathbf{x}_1 = \frac{\mathbf{x}_0+\mathbf{x}_2}{2}$.

III. Evaluation

A. Experiment Setup

Fig. 1A illustrates our experimental setup. The granular trackway is (60 cm L × 60 cm W × 20 cm D) and contains model granular medium (Grainger, 0.3 mm glass beads). The 0.3 mm particle size is similar to those observed in natural deserts, and behave qualitatively similar with natural sand, and thus has been used as an ideal model granular medium in many previous granular physics and robot locomotion studies [4, 10, 11]. The granular trackway can be tilted up to 35 degrees to emulate a wide variety of sand slopes in natural environments [5]. We mount an RGBD camera

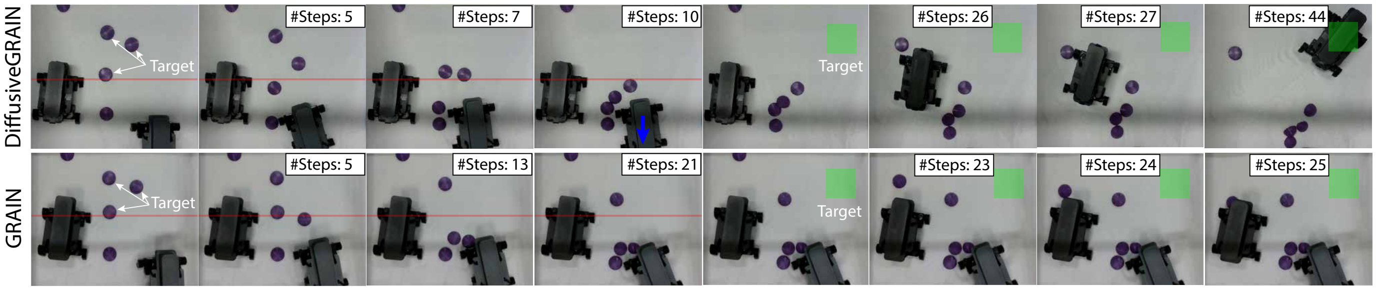

Fig. 3. One multi-robot loco-manipulation trial using DiffusiveGRAIN and GRAIN. We place one robot in the middle left, and surround it with obstacles. We place a second robot to the bottom right. The second robot must manipulate the 3 target obstacles to move below the red line. In the DiffusiveGRAIN trial, the bottom right robot achieves the manipulation task and moves backward to create room at Step 10. The first robot is free and then locomotes to the target shown by the green square. In the GRAIN trial, both the manipulation task and the locomotion task failed. Specifically, in Step 21 and Step 25 the robots stop the policy execution because GRAIN predicts there is no action that could further optimize the cost function.

(Intel RealSense 435-i) above the granular slope to record the granular flow and obstacle movement.

B. Baselines and Evaluation Protocol

Robot movement prediction Baseline: We used GRAIN GRAIN [9] as the locomotion baseline. GRAIN takes the depth image concatenated with the robot action representation image as input and outputs a 1×3 matrix, which corresponds to robot positions on the x-axis and y-axis and robot orientation.

Manipulation Planning Baseline: As a strong baseline, we adapted GRAIN [9] to this setup. We trained GRAIN with the new dataset and output the coordinates of each obstacle. The GRAIN framework was originally designed to train on depth images with only one obstacle, resulting in a 1×2 matrix corresponding to obstacle positions on the x-axis and y-axis. However, we have 1 to 5 obstacles. Consequently, we provided the ground truth number of obstacles to GRAIN and modified its output to be a 5×2 matrix. When there are $N < 5$ obstacles, GRAIN only considers the first N rows in the matrix.

Evaluation Protocol: We evaluate using Euclidean distance. To compute the "position" of each obstacle and robot, we estimate their Center of Mass (CoM). For manipulation, we sum the Euclidean distance among obstacle positions and their targets. For locomotion, we measure the distance between the robot and the locomotion target. In all trials, the robot stops acting if the policy predicts that no action can further reduce the cost function. In addition, an automatic failure applies if a robot robot flips over during policy execution.

C. Evaluation Results

Single Robot Manipulation: Using DiffusiveGRAIN with manipulation mode results in 8/10 success, while GRAIN results in 6/10 success. Among trials that fail, the mean errors of DiffusiveGRAIN are 12.8±2.4 cm, compared to 17.8±5.2 cm for GRAIN.

Single Robot Locomotion: Using DiffusiveGRAIN with locomotion mode results in 9/10 success, while GRAIN gets 8/10 success. Among trials that fail, DiffusiveGRAIN's errors are 8.6±0 cm while GRAIN's errors are 12.4±3.0 cm.

Single Robot Loco-manipulation: In the single robot loco-manipulation evaluation, the robot must move all 4 obstacles to a target region (below the red line). In parallel, it must *also* navigate to a target location (marked by a green square). Using DiffusiveGRAIN with loco-manipulation mode results in 7/10 success, compared to 2/10 for GRAIN. The average error for failed trials is 17.2±4.6 cm for DiffusiveGRAIN, and 28.4±9.0 cm for GRAIN.

Multi-robot collaboration: In this task, we surround a robot with obstacles, and place a second "free" robot nearby (see Fig. 3). The human specifies the 3 obstacles that the second robot needs to manipulate to free the first robot. The human also specifies a locomotion task for the trapped robot. During each trial, the free robot acts first, and executes its manipulation policy to adjust the location of the 3 obstacles. After achieving the manipulation task, that robot moves back to make room for the first—originally trapped—robot (see Fig. 3 at Step 10). As the first robot is no longer trapped, it must only locomote to its target. The robot successfully arrives at the target location at Step 44, ultimately suggesting the potential for DiffusiveGRAIN in a multi-agent robot system with locomotion and manipulation. Using DiffusiveGRAIN with locomotion and manipulation modes, the robot achieves 7/10 success, while GRAIN results in 4/10 success. The failure errors for DiffusiveGRAIN are 15.8cm±3.6 cm, versus 23.8±7.6 cm for GRAIN.

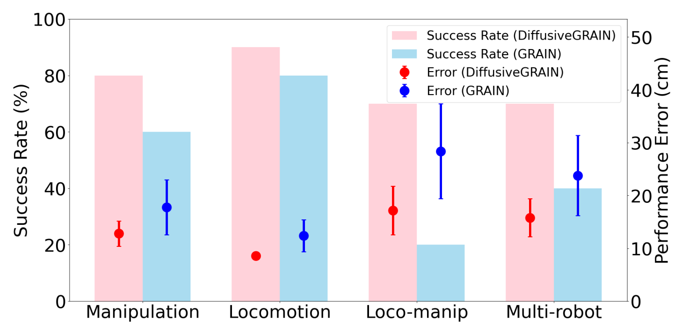

Fig. 4. Experiment results. Pink and light blue bars are the success rate of DiffusiveGRAIN and the baseline (GRAIN) respectively. The red and blue error bars are the mean and standard deviation of failed trials in each task respectively, measured by the sum of distances among the objects and target locations at the end of policy execution.

References

[1] G. Bradski. The OpenCV Library. *Dr. Dobb's Journal of Software Tools*, 2000.

[2] Kaustav Chakraborty, Haodi Hu, Matthew D Kvalheim, and Feifei Qian. Planning of obstacle-aided navigation for multi-legged robots using a sampling-based method over directed graphs. *IEEE Robotics and Automation Letters*, 7(4):8861–8868, 2022.

[3] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations (ICLR)*, 2021.

[4] Justin R Finn, Ming Li, and Sourabh V Apte. Particle based modelling and simulation of natural sand dynamics in the wave bottom boundary layer. *Journal of Fluid Mechanics*, 796:340–385, 2016.

[5] Nick Gravish and Daniel I Goldman. Effect of volume fraction on granular avalanche dynamics. *Physical Review E*, 90(3):032202, 2014.

[6] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models. In *Neural Information Processing Systems (NeurIPS)*, 2020.

[7] Haodi Hu and Feifei Qian. Obstacle-aided trajectory control of a quadrupedal robot through sequential gait composition. *IEEE Transactions on Robotics*, pages 1–15, 2024.

[8] Haodi Hu, Xingjue Liao, Wuhao Du, and Feifei Qian. Multi-robot connection towards collective obstacle field traversal, 2024. URL `https://arxiv.org/abs/2409.11709`.

[9] Haodi Hu, Feifei Qian, and Daniel Seita. Learning granular media avalanche behavior for indirectly manipulating obstacles on a granular slope. In *8th Annual Conference on Robot Learning*, 2024.

[10] Chen Li, Tingnan Zhang, and Daniel I Goldman. A terradynamics of legged locomotion on granular media. *science*, 339(6126):1408–1412, 2013.

[11] Ryan D Maladen, Yang Ding, Chen Li, and Daniel I Goldman. Undulatory swimming in sand: subsurface locomotion of the sandfish lizard. *science*, 325(5938): 314–318, 2009.

[12] Ratan Othayoth, George Thoms, and Chen Li. An energy landscape approach to locomotor transitions in complex 3d terrain. *Proceedings of the National Academy of Sciences*, 117(26):14987–14995, 2020.

[13] Feifei Qian and Daniel Goldman. Anticipatory control using substrate manipulation enables trajectory control of legged locomotion on heterogeneous granular media. In *Micro-and Nanotechnology Sensors, Systems, and Applications VII*, volume 9467, page 94671U. International Society for Optics and Photonics, 2015.

[14] Feifei Qian and Daniel E Koditschek. An obstacle disturbance selection framework: emergent robot steady states under repeated collisions. *The International Journal of Robotics Research*, 2020.

[15] Divya Ramesh, Anmol Kathail, Daniel E Koditschek, and Feifei Qian. Modulation of robot orientation via leg-obstacle contact positions. *IEEE Robotics and Automation Letters*, 5(2):2054–2061, 2020.

[16] Jennifer M Rieser, Perrin E Schiebel, Arman Pazouki, Feifei Qian, Zachary Goddard, Kurt Wiesenfeld, Andrew Zangwill, Dan Negrut, and Daniel I Goldman. Dynamics of scattering in undulatory active collisions. *Physical Review E*, 99(2):022606, 2019.

[17] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2015.

[18] Perrin E Schiebel, Jennifer M Rieser, Alex M Hubbard, Lillian Chen, and Daniel I Goldman. Collisional diffraction emerges from simple control of limbless locomotion. In *Conference on Biomimetic and Biohybrid Systems*, pages 611–618. Springer, 2017.

[19] Perrin E Schiebel, Jennifer M Rieser, Alex M Hubbard, Lillian Chen, D Zeb Rocklin, and Daniel I Goldman. Mechanical diffraction reveals the role of passive dynamics in a slithering snake. *Proceedings of the National Academy of Sciences*, 116(11):4798–4803, 2019.

[20] Tianyu Wang, Christopher Pierce, Velin Kojouharov, Baxi Chong, Kelimar Diaz, Hang Lu, and Daniel I Goldman. Mechanical intelligence simplifies control in terrestrial limbless locomotion. *arXiv preprint arXiv:2304.08652*, 2023.

[21] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *ACM Computing Surveys*, 56(4):1–39, 2023.