Exploring Planning of Redistribution Trajectories for Profile Grading in Amorphous Materials

Louis Hanut[⋆,1,4], Yurui Du[2], Renaud Detry[1,2], Herman Bruyninckx[1,3,4]

***Abstract*— This paper addresses the problem of planning actions to reshape an amorphous material surface to match a desired distribution. This task is complicated by the uncertain and dynamic properties of such materials. Inspired from state-of-the-art planning for the manipulation of amorphous materials, we explore a next-best trajectory planning algorithm for sand redistribution. The method generates multiple candidate sweeps informed by domain-specific priors, including Optimal Transport, a Max-to-Min heuristic, and Adjacent Sweep strategies. Each candidate's outcome is predicted using a simplified sand-tool interaction model, and the trajectory yielding the lowest predicted error is executed. Simulation-based experiments demonstrate that combining multiple candidate generation strategies improves redistribution accuracy and robustness across diverse scenarios.**

I. INTRODUCTION

Automating construction tasks is crucial to meet growing demands while improving the working conditions of construction workers. However, the robotic manipulation of amorphous materials –such as sand, gravel, or plaster –remains challenging due to their uncertain, time-varying, and often heterogeneous properties.

Humans typically address these challenges through iterative and adaptive behavior, adjusting their actions based on observed outcomes. Inspired by this strategy, we study the planning of redistribution actions to shape an amorphous material surface.

Recent works have shown that including feedback, tool-material interaction models, and domain-specific heuristics can significantly enhance performance in amorphous material manipulation. Building on these insights, we propose a next-best-sweep planning framework, inspired by [6], that integrates material modeling and rule-based trajectory generation to produce candidate actions.

Through simulation experiments, we show that combining multiple candidate generation strategies leads to improved surface grading performance.

⋆ Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or HADEA. Neither the European Union nor the granting authority can be held responsible for them.

Funded by the European Union

⋆ Corresponding author, louis.hanut@kuleuven.be

[1] Department of Mechanical Engineering, KU Leuven, Belgium.

[2] Department of Electrical Engineering, KU Leuven, Belgium.

[3] Department of Mechanical Engineering, TU/e Eindhoven, the Netherlands.

[4] Flanders Make @ KU Leuven, Belgium.

II. RELATED WORK

Different methods have been employed in the literature for the manipulation of amorphous materials. First, predictive models were developed to estimate the outcome of actions and build plans accordingly [1], [2].

These models vary in complexity, from simple binary-grid representations [3] to detailed models based on solid/fluid mechanics or discrete element methods. However, for real-time robotic applications, simplified models are generally preferred due to their lower computational demands [1], [4]. Including feedback to update the material state allowed to compensate model simplifications and inaccuracies [5], [6].

Alternatively, model-free approaches were proposed, using domain-specific heuristics to guide planning (e.g., vision-based contour shaping heuristics) [7].

Such rules have also been integrated as priors within model-based frameworks, e.g., for tasks like plastering [5] or cleaning [3]. Notably, [6] used Optimal Transport as a prior to propose candidate actions, which were then evaluated using a model of the material.

Finally, as prediction models depend on several parameters that are difficult to obtain (material, tool, environment conditions ...), learning-based methods have been used to model the material behavior [4], [8], [9]. Learning has also been applied to derive action-selection policies for tasks such as bulldozing [10], [11] and everyday material manipulation [12].

In this work, we present a planning algorithm that combines material state feedback, a simple predictive model and domain knowledge as informative prior to guide action selection.

III. PLANNING MATERIAL REDISTRIBUTION

We address the problem of planning redistribution actions to shape an initial material distribution H_{init} into a desired one H_{des}. Drawing from the literature on amorphous material manipulation, we identify three key elements to enhance planning: (1) intermediate feedback, (2) tool–material interaction modeling, and (3) domain-informed planning priors.

To integrate these components, we propose a "next-best sweep" planning algorithm similar to the one presented in [6]. In between each executed action, the algorithm generates a set of candidate sweep trajectories and evaluates them based on the predicted material redistribution outcome. The candidate generation and selection are essential in the planning and include domain knowledge for efficient redistribution. The planning algorithm is outlined in Algorithm 1, with its main components detailed in the following sections.

A. Simplifying Assumptions and Action Space Restriction

To reduce the complexity of the planning problem, the action space is limited to straight-line sweeps, each defined by the start and end coordinates in the workspace:

$$a = (x_{\text{start}}, y_{\text{start}}, x_{\text{end}}, y_{\text{end}})$$

We use discrete sweeps rather than continuous interactions to simplify execution. Continuous tool motion poses greater challenges for robotic kinematics and results in the constant accumulation and transport of material in front of the tool. In contrast, discrete sweeps allow localized accumulation, enabling material to be deposited temporarily and handled in subsequent actions. Furthermore, the total material quantity is assumed equal to the material quantity of the desired distribution such that no additive or subtractive operations are necessary.

B. Sweep Candidates Generation

We propose three domain-informed methods to generate candidate sweep trajectories:

1) Optimal Transport (OT): Following [6], we leverage Optimal Transport to compute a material redistribution plan. The OT algorithm computes a transport matrix T indicating the optimal flow of material between grid cells. Due to the high computational cost on large grids, we apply OT to a downsampled version of the heightmap (see Fig. 1), yielding a coarse redistribution plan between aggregated zones.
OT assumes direct, non-interfering transport from source to destination cells (as for scoop-and-dump operations). However, our context involves sweeping actions that push material across the surface, which do not strictly adhere to the OT plan. Nevertheless, we assume that the OT plan still provides a strong prior for rearranging the material distribution. Deviations from the OT plan are leveraged by recomputing the transport plan in between each executed sweep, using updated heightmaps from sensor feedback.

2) Max-to-Min Heuristic: This approach defines trajectories from the maximum to minimum of the relative heightmap $H_{\text{rel}} = (H_{\text{mes}} - H_{\text{des}})$. The resulting trajectory connects the maximum material surplus to the maximum deficit points.

3) Adjacent Sweeps: During sweep execution, material accumulates in front of the tool and spills laterally, leaving residual traces beside the tool path. We generate candidate trajectories, a_{adj}, offset by $\pm\delta$ from the previous sweep, where δ is half the tool width. These parallel paths overlap with prior sweeps and allow to remove side-spilled material (see Fig. 1).

Candidate Set Construction: The final candidate set $\mathcal{A} = \{a_{\text{OT}} + \epsilon, a_{\text{max_min}} + \epsilon, a_{\text{adj}}\}$ comprises:

- a_{OT}: 5 actions generated from the maximum element in the OT plan.
- $a_{\text{max_min}}$: 5 actions generated from the Max-to-Min heuristic.
- a_{adj}: 2 paths adjacent to the previous sweep.

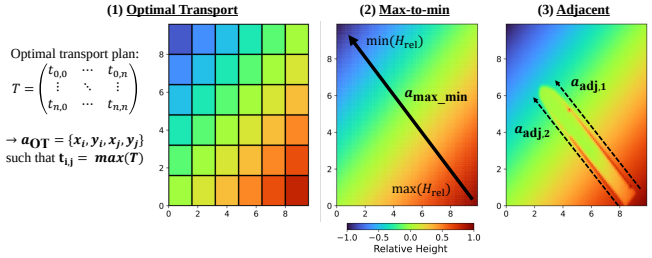

Fig. 1: Illustration of the candidate generation methods. (1) The Optimal transport action uses the maximum element of the transport plan T. (2) The max-to-min sweep is defined between $max(H_{\text{rel}})$ and $min(H_{\text{rel}})$. (3) Adjacent method produces two candidates along the side spills left by the previous sweep.

For a_{OT} and $a_{\text{max_min}}$, a stochastic perturbation ϵ is introduced for the exploration of nearby trajectories. The candidate generation process is performed in line 1 of Algorithm 1.

C. Sweep prediction - Sand-tool interaction modeling

The next step in Algorithm 1 (line 3) predicts the resulting heightmap for each candidate $a \in \mathcal{A}$, using a model similar to [1] (see Fig. 2). The interaction between the tool and the material is modeled in two consecutive steps, applied pixel-wise along the sweep path:

1) **Step 1 - Tool-induced Motion**: All material volume within the tool's swept area is displaced to neighboring pixels.
2) **Step 2 - Natural Material Flow**: The material returns to its equilibrium state, defined by its angle of repose [13]. Flow only occurs in directions not obstructed by the tool (indicated in purple in Fig. 2).

Although more complex models for various amorphous materials exist in the literature, this simplified model is adequate for our work: the repose angle is its only parameter and could be estimated online through sensor data. While the model introduces approximations, it can still provide a good prior for action selection. Model inaccuracies are mitigated by sensor-based updates between planning iterations.

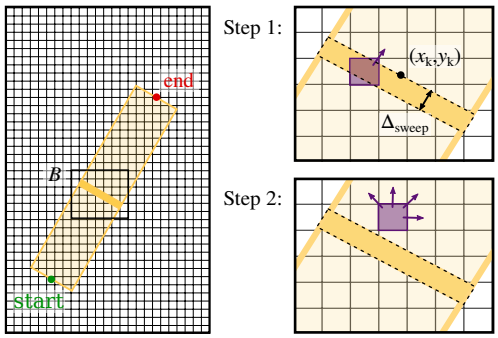

Fig. 2: Sand-tool interaction modeling [1]. For each pixel along the tool trajectory, two prediction steps are computed. Step 1: Tool-induced Motion. Step 2: Natural Material Flow.

D. Sweep Candidates Evaluation

Each predicted heightmap resulting from a candidate action $a \in \mathcal{A}$ is evaluated using a cost function that reflects task performance:

$$C_a = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(z_{i,\text{des}} - z_{i,\text{mes}})^2} \tag{1}$$

Here, $z_{i,\text{des}}$ and $z_{i,\text{mes}}$ denote the desired and measured heights at pixel i, respectively. This root mean square error (RMSE) quantifies the deviation from the target heightmap.

Candidate evaluation occurs in line 4 of Algorithm 1. The cost C_a of each action is compared to the current best cost C_{best}; if lower, the candidate becomes the new best action a_{best}.

E. Best Sweep Execution

The final step of Algorithm 1 (line 10) executes the selected best candidate a_{best}. During execution, control variables are adapted in real time to address unforeseen deviations and ensure compliance with task constraints, as described in [14]. After the execution, the obtained material surface is evaluated to decide if another iteration is necessary. In that case, Algorithm 1 is repeated.

Algorithm 1 Grading sweep planning algorithm based on a "next-best-sweep" approach [6]. From a set of sweep candidates, the best action is determined using the prediction model in Fig. 2.

1: $\mathcal{A} \leftarrow$ **Generate_sweep_candidates**$(H_{\text{des}}, H_{\text{mes}})$
2: **for** a in $\mathcal{A}$ **do**
3: $\quad H_{\text{obtained}} \leftarrow$ **Sweep_prediction**(a)
4: $\quad C_a \leftarrow$ **Evaluate_prediction**$(a,\ H_{\text{obtained}})$
5: $\quad$ **if** $C_a < C_{\text{best}}$ **then**
6: $\quad\quad a_{\text{best}} \leftarrow a$
7: $\quad\quad C_{\text{best}} \leftarrow C_a$
8: $\quad$ **end if**
9: **end for**
10: **Execute_sweep**(a_{best})

IV. PRELIMINARY RESULTS AND DISCUSSION

To evaluate the performance of Algorithm 1, we conducted a series of simulation experiments using the model illustrated in Fig. 2. All experiments were performed on 100x100 pixels heightmaps. Fig. 3 contains the task progression for two redistribution scenarios. In Scenario (a), the initial and desired heightmaps correspond to vertical and horizontal linear gradients, respectively. In Scenario (b) both initial and desired distributions are bimodal, featuring two distinct summits.

For each scenario, we compare our method –which uses a candidate set combining all three candidate generation methods –with each method individually. Performance is evaluated based on the RMSE between the final and desired heightmaps after 50 sweep executions, the total distance traveled by the tool, and the total planning time (see Table I). From these results, we observe the following:

- **Improved Redistribution Accuracy**: Our combined approach yields the lowest RMSE in both scenarios, demonstrating the benefit of combining diverse candidate generation strategies.
- **Planning Time Considerations**: As expected, planning time increases with the number of simulated candidates and is higher for our method. Nevertheless, it remains acceptable compared to the real sweep execution time. In practice, planning time can be significantly reduced through parallelization of the prediction step.
- **Role of the Adjacent Method**: On its own, the Adjacent method performs poorly due to its dependence on the initial sweep direction (chosen arbitrarily along the diagonal). However, it complements other strategies well by enabling local refinement of the previous sweep.
- **OT vs. Max-to-Min Trade-offs**: The Max-to-Min method ignores spatial context and may result in inefficient long-range transport. In Scenario (a), this leads to higher tool travel distance compared to Optimal Transport (OT), despite similar RMSE values. Conversely, in Scenario (b), Max-to-Min outperforms OT in redistribution accuracy. This highlights the benefit of keeping both methods to improve robustness against different scenarios.

Overall, these preliminary findings underscore the advantage of combining multiple candidate generation heuristics to enhance robustness and redistribution quality across varying scenarios.

TABLE I: Results of the simulated redistribution after 50 sweeps for two scenarios: (a) Linear distributions and (b) Bimodal distributions (see Fig. 3). Performance is evaluated based on the RMSE between final and desired heightmaps, the total distance traveled by the tool, and the total planning time.

(a) Linear distributions	RMSE	Traveled distance	Planning time
OT	0.2056	368.67	164.3
Max-to-Min	0.2060	446.57	199.37
Adjacent	0.3033	256.4	68.15
Our method	**0.1420**	448.94	383.02

(b) Bimodal distributions	RMSE	Traveled distance	Planning time
OT	0.3154	263.44	129.94
Max-to-Min	0.2455	285.97	141.8
Adjacent	0.3950	206.33	66.07
Our method	**0.2074**	285.73	279.6

V. FUTURE WORK

A first direction for future work is to integrate the proposed planner with the real-time adaptive execution presented in [14]. This would allow adjustment of model parameters (e.g., material's repose angle) based on online measurements. Additionally, during adjacent sweep execution, controlling the tool orientation could help direct material flow more effectively, particularly by pushing material outward from the sweep path.

Next, the experimental scenarios considered in this paper are

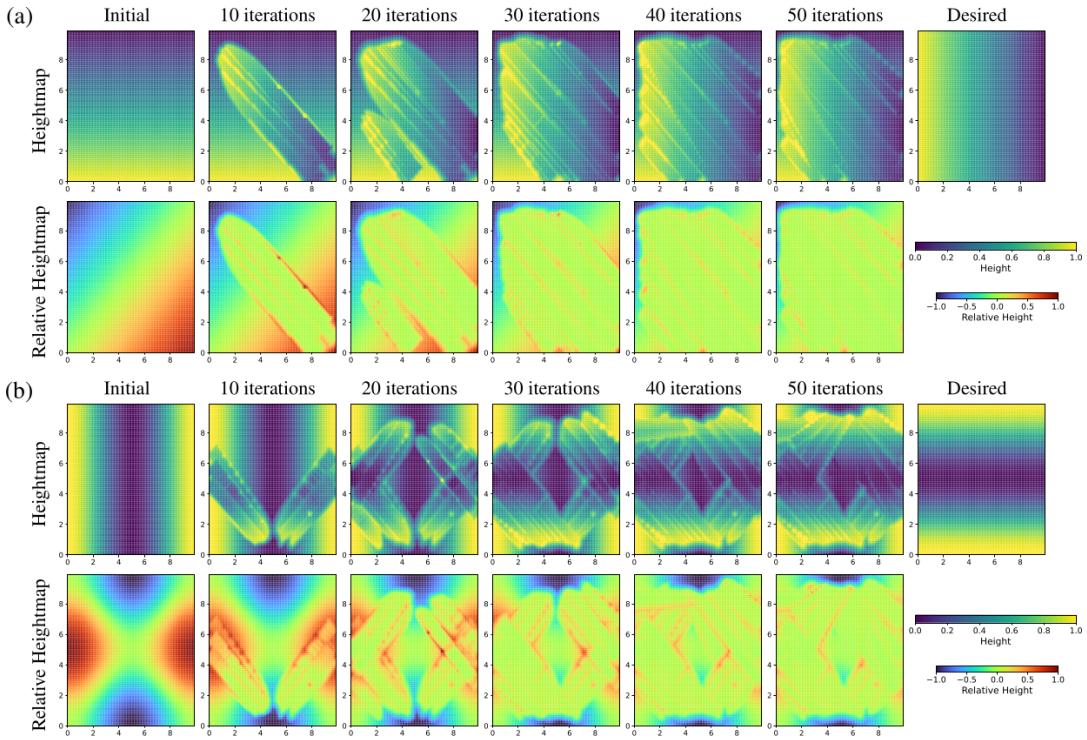

Fig. 3: Results of the simulation experiments, showing task progression when using our planning method (combining all three candidate generation heuristics) in the two redistribution scenarios: (a) Linear distributions and (b) Bimodal distributions.

relatively simple. To be applicable in real-world construction settings, the planner must be extended to handle more complex surface geometries and interactions with rigid obstacles. Finally, the current planning implementation selects the next action by minimizing the immediate cost. Planning over a longer horizon may enhance the overall redistribution performance [5].

REFERENCES

[1] W. Kim, C. Pavlov, and A. M. Johnson, "Developing a Simple Model for Sand-Tool Interaction and Autonomously Shaping Sand," arXiv, Aug. 2019. doi: 10.48550/arXiv.1908.02745.

[2] D. Leidner, G. Bartels, W. Bejjani, A. Albu-Schäffer, and M. Beetz, "Cognition-enabled robotic wiping: Representation, planning, execution, and interpretation," Robotics and Autonomous Systems, vol. 114, pp. 199–216, Apr. 2019, doi: 10.1016/j.robot.2018.11.018.

[3] S. Elliott and M. Cakmak, "Robotic Cleaning Through Dirt Rearrangement Planning with Learned Transition Models," in 2018 IEEE International Conference on Robotics and Automation (ICRA), May 2018, pp. 1623–1630. doi: 10.1109/ICRA.2018.8460915.

[4] K. Tsuruta, S. J. Griffioen, J. M. Ibáñez, and R. L. Johns, "Deep Sandscapes: Design Tool for Robotic Sand-Shaping with GAN-Based Heightmap Predictions," in 2022 Annual Modeling and Simulation Conference (ANNSIM), Jul. 2022, pp. 730–741. doi: 10.23919/ANNSIM55834.2022.9859472.

[5] M. A. Kuhlmann-Jørgensen, J. Pankert, L. L. Pietrasik, and M. Hutter, "PLASTR: Planning for Autonomous Sampling-Based Trowelling," IEEE Robotics and Automation Letters, vol. 8, no. 8, pp. 5069–5076, Aug. 2023, doi: 10.1109/LRA.2023.3291894.

[6] N. Alatur, O. Andersson, R. Siegwart, and L. Ott, 'Material-Agnostic Shaping of Granular Materials with Optimal Transport', in 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Oct. 2023, pp. 10652–10659. doi: 10.1109/IROS55552.2023.10342505.

[7] A. Cherubini, V. Ortenzi, A. Cosgun, R. Lee, and P. Corke, "Model-free vision-based shaping of deformable plastic materials," The International Journal of Robotics Research, vol. 39, no. 14, pp. 1739–1759, 2020, doi: 10.1177/0278364920907684.

[8] C. Schenck, J. Tompson, S. Levine, and D. Fox, "Learning Robotic Manipulation of Granular Media," PMLR, Oct. 2017, pp. 239–248. Accessed: Mar. 13, 2023. [Online]. Available: https://proceedings.mlr.press/v78/schenck17a.html

[9] Z. Xian et al., "FluidLab: A Differentiable Environment for Benchmarking Complex Fluid Manipulation," in International Conference on Learning Representations, 2023.

[10] C. Ross, Y. Miron, Y. Goldfracht, and D. D. Castro, "AGPNet - Autonomous Grading Policy Network," ISARC Proceedings, pp. 40–43, May 2022.

[11] Y. Miron, Y. Goldfracht, C. Ross, D. D. Castro, and I. Klein, "Autonomous Dozer Sand Grading Under Localization Uncertainties," IEEE Robotics and Automation Letters, vol. 8, no. 1, pp. 65–72, Jan. 2023, doi: 10.1109/LRA.2022.3222990.

[12] Y. Zhang, W. Yu, C. K. Liu, C. Kemp and G. Turk, "Learning to manipulate amorphous materials", ACM Trans. Graph., vol. 39, no. 6, 2020.

[13] S. Nowak, A. Samadani, and A. Kudrolli, 'Maximum angle of stability of a wet granular pile', Nature Phys, vol. 1, no. 1, pp. 50–52, Oct. 2005, doi: 10.1038/nphys106.

[14] L. Hanut, Y. Du, A. Vande Moere, R. Detry, and H. Bruyninckx, "Robotic Framework for Iterative and Adaptive Profile Grading of Sand," in IEEE International Conference on Robotics and Automation, 2025.