

# Data-Driven Bulldozer Blade Control for Autonomous Terrain Leveling

Harry Zhang<sup>†</sup> Ganesh Arivoli<sup>†</sup> Huzaifa Unjhawala Luning Bakke Radu Serban Dan Negrut

<sup>†</sup> Equal Contribution

Department of Mechanical Engineering, University of Wisconsin-Madison

1513 University Ave, Madison, Wisconsin

Email Address: hzhang699, arivoli, unjhawala, lfang9, serban, negrut@wisc.edu

**Keywords:** *High-Fidelity Simulation, Reduced Order Modeling, Autonomous Construction*

Autonomous leveling of granular materials is a ubiquitous yet challenging operation in automated construction due to the complex physics governing the soil-tool interaction. This paper outlines a simulation-driven framework for optimizing low-level bulldozer blade control (pitch and height) to enhance leveling performance. The approach uses high-fidelity, physics-based simulations to generate training data. This data informs a Neural Network (NN) based reduced-order model that accurately predicts both the terrain evolution and the leveling operation duration in response to blade actions. A gradient-based, multi-objective optimization algorithm then utilizes the reduced-order model to determine optimal, time-varying blade control profiles, managing the trade-off between leveling flatness and operation time. The proposed method augments the state-of-the-art by producing policies that can readily level arbitrary soil pile configurations while avoiding vehicle immobilization and achieving better leveling efficiency. The system exhibits robustness to variations in initial pile geometry, and offers explicit control over the trade-off between leveling quality and operational efficiency. By integrating high-fidelity physics into the controller design and providing an open-source simulation pipeline, this work provides a low-level control solution that complements existing global path planning algorithms for autonomous construction operations. The project's resources, including code and media demonstration, are available at: [link<sup>1</sup>](#).

# 1 Introduction

Autonomous construction operations, including site preparation, substructure construction, and superstructure assembly, have garnered significant attention in recent years [1, 2]. Granular material manipulation, a close second to fluid handling in industrial applications relevance, presents significant challenges due to the intricate physics that govern the interaction between the material and the manipulation blade. In the field of autonomous leveling, Osaka et al. [3] demonstrated that a heightmap model, utilizing Enhanced Sand Pile Models (ESPM), can effectively describe bulldozer-soil interactions. Combined with reinforcement learning techniques, their algorithm was shown to perform placement tasks. This heightmap-based approach was also employed to address stockpile mixing through safe and efficient path planning. Inspired by practical bulldozing operations, Hirayama et al. [4] integrated integer linear programming for granular material segmentation, subsequently used for bulldozer path planning. Further work by Hirayama et al. [5] detailed a topological path planning methodology to determine efficient multi-push sequences for bulldozers in dumping operations, focusing on the overall material handling strategy. Beyond path planning, Wagner et al. [6] proposed a scheme for optimized bulldozer blade control during leveling operations. While the aforementioned studies were primarily validated through simulation, other research has demonstrated efficacy in real-world experiments. Addressing sensor fusion and blade state estimation, Kim et al. [7] developed a multi-sensor algorithm demonstrating high tracking accuracy, although this estimation was not integrated into a feedback control loop. Lee et al. emphasized the use of blade control by developing a Cartesian-space system aimed at enhancing bulldozing performance through improved blade positioning [8]. Highlighting the importance of blade control in autonomous bulldozing, Mononen et al. developed adaptive blade control strategies that enable surface profile tracking and support semi-autonomous operation [9, 10]. While these studies advanced blade control capabilities, complementary research has focused on higher-level planning strategies. Adopting a simulation-to-reality (sim2real) approach, Miron et al. [11, 12, 13] employed a heightmap model for terrain representation and rule-based route labeling to generate a training dataset. Subsequently, a behavioral cloning technique was used to train a Neural Network (NN) for predicting agent routes for ground leveling. In real-world testing, a top-view RGBD camera pro-

vided heightmap observations, and the trained NN predicted bulldozer routes for leveling ground piles. Heightmap-based terrain models are the predominant choice in state-of-the-art simulations of bulldozer-soil interactions. This approach is embraced herein, as it displays two key characteristics: it efficiently captures essential terrain height evolution during operations, and has demonstrated sim2real transferability when coupled with depth sensors in real-world experiments [3, 6, 11, 12, 13]. Another common characteristic of the state of the art is the utilization of simplified or empirical models for the soil-bulldozer terramechanics interaction. Existing research has employed empirical force-based blade-soil interaction models (e.g., ESPM [14, 15]), often times captured in simple rules that emulate the interplay between soil heightmap profiles and blade movement [6, 11, 16]. Inspired by Wagner et al. [6], we employ a NN as a reduced-order model, diverging from rule-based or empirical equation-based approaches. By using heightmap frames as NN inputs and outputs, complex terramechanics interactions can be embedded in the NN by training the model with high-quality data. Furthermore, the differentiability of NN models facilitates gradient-based optimization tailored to meet operational requirements [17, 18]. This strategy of using differentiable NN surrogates, trained on high-fidelity simulation data for complex physical systems, is also finding applications in other domains involving particle dynamics. For instance, Liu et al. employed a Graph Neural Network (GNN) trained on DEM-SPH simulation data to model particle-fluid flows, leveraging the GNN's differentiability for an inverse design framework to optimize flow control structures [19]. While some studies [11, 6] note the high computational cost as a disadvantage of using physics-based terramechanics simulation, we embrace this approach for two reasons. Firstly, recent hardware (GPU acceleration [20]) and modeling (continuum hypo-elasto-plastic terramechanics [21]) advancements have improved both the accuracy of the synthetic data and the speed of producing it. As a rule of thumb, the real time factor for a typical bulldozing simulation in this paper is around 35; i.e., to simulate 1 second of the vehicle bulldozing deformable terrain, one has to wait approximately 35 seconds. Secondly, the higher fidelity simulation comes into play exclusively in the synthetic data collection (data is subsequently used in NN model training). Recently, machine learning techniques have been increasingly employed to address the computational challenges in demanding fields such

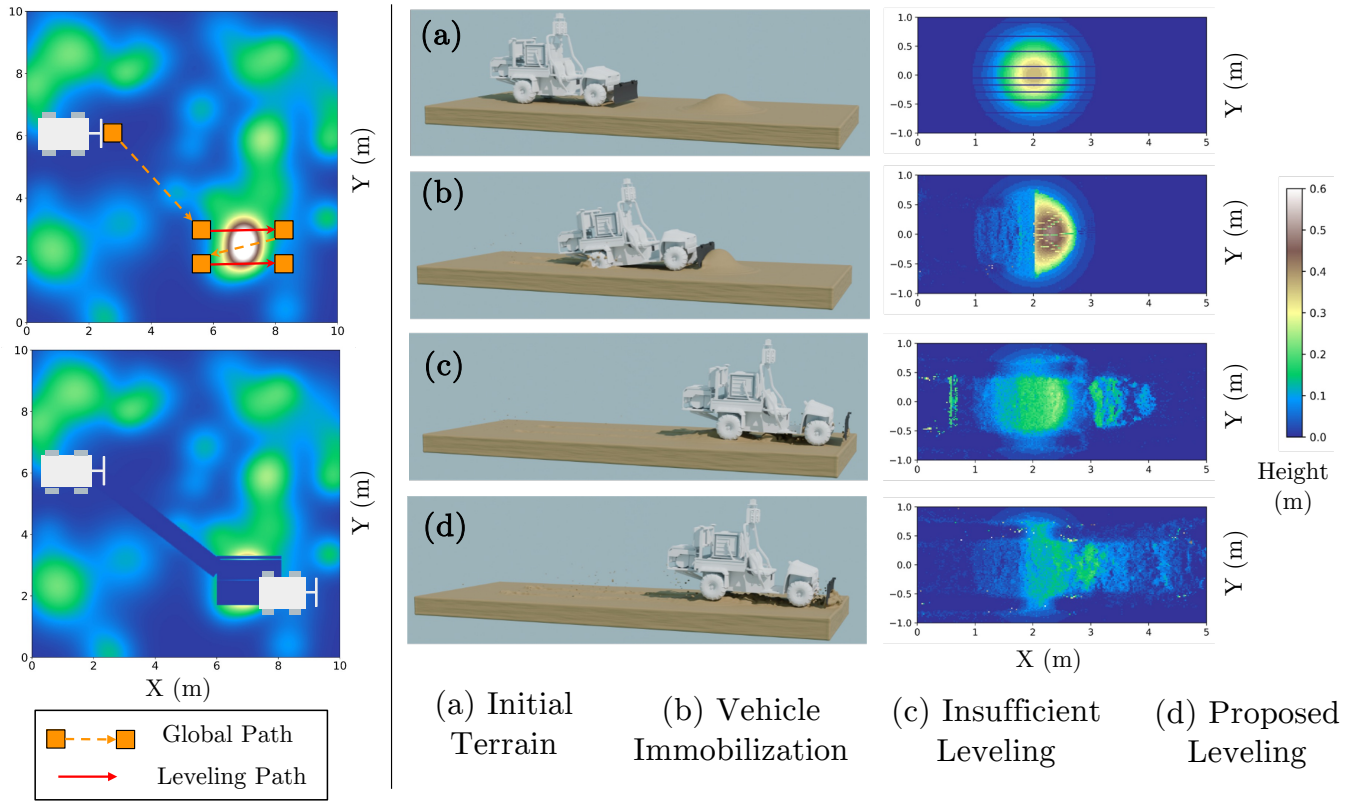


Figure 1: Unlike other studies that aim at producing a leveling path or plan [3, 11, 12] (left side of the figure), this effort optimized the blade control (height, pitch) for a given path for optimal leveling performance. (a) Initial configuration: The vehicle is positioned before a soil pile targeted for leveling. (b) Fixed blade parameters: Suboptimal blade height and pitch leading to high resistance and vehicle immobilization. (c) Baseline control (fixed parameters): The vehicle traverses the terrain but achieves suboptimal leveling due to a static blade configuration. (d) Proposed learning-based control: Optimized blade pitch and height based on the initial terrain state, achieving efficient soil displacement and enhanced leveling performance.

as computational fluid dynamics [22, 23]. Thus, the increased simulation time primarily impacts the training phase without affecting the operational aspects of the autonomy solution. To generate high-fidelity synthetic training data, this effort uses the Project Chrono simulation engine [24]. The two modules employed are Chrono::Vehicle [25] and Chrono::FSI [26]. Chrono's vehicle and terrain co-simulation capabilities enable high-fidelity simulation of vehicle powertrain dynamics, tire slip, and sinkage phenomena [25]. The terramechanics and vehicle co-simulation environment have been previously validated against real vehicle data, demonstrating the simulation engine's accuracy [21, 27].

### Problem Statement and Contributions

Recent advances in autonomous bulldozing have focused on path planning methodologies [11, 3, 16]. These approaches often rely on simplified representations of the complex dynamics at the bulldozer-terrain interface. Although active blade control during operation has received attention [9, 10, 28], the systematic optimization of adaptive, multi-degree-of-freedom

blade actions—crucial for achieving high-quality terrain flatness—remains a key challenge, especially in a form that complements global path planning strategies. To address this, we leverage a high-fidelity physics engine that accurately models the complex terramechanics involved. However, its computational cost makes it difficult for use in real-time control loops or iterative optimization algorithms. To overcome this limitation, we develop a reduced-order neural network (NN) model [6], trained on data generated from high-fidelity simulations to capture terrain deformation dynamics efficiently. This differentiable NN surrogate enables a model-based optimization framework that maximizes ground leveling performance via precise, low-level control of blade actions—specifically pitch and elevation. As illustrated in Fig. 1, our focus is on the blade-soil interaction segment, i.e., determining how an autonomous bulldozer should dynamically actuate its blade for optimal leveling. Our contributions are threefold:

1. Integration of physics-based terramechanics simulation into the autonomous bulldozing design framework.




Figure 2: Overview of the data collection process involving a simulated two-pass terrain leveling operation. (Top-left) An initial heightmap is recorded at  $t = 0.0$  s. First operation: The bulldozer moves forward (0.0s-1.0s), then operates with blade control parameters set 1 (1.0s-3.5s) followed by set 2 (3.5s- 6.0s). When it reached 6s or rightmost terrain limit, the bulldozer stops leveling and moves back. This point in time is referred to as the first operation duration  $t_d^1$ . Bottom-left: An intermediate heightmap is recorded after the bulldozer returns to the starting area at  $t_{back}$ . Second Pass: Starting from  $t_{back}$ , the bulldozer again moves forward (for 1.0s), then operates with blade parameters set 3 (active for 2.5s), followed by set 4 (active for another 2.5s). When the vehicle reaches  $t_{back} + 6$  s or the rightmost terrain limit, the simulation stops, at end time  $t_{end}$ ; the second operation duration is  $t_d^2 = t_{end} - t_{back}$ . Bottom-right: A final heightmap is recorded at  $t_{end}$ . Each simulation yields three heightmaps: initial  $t = 0$ , intermediate  $t_{back}$ , and final  $t_{end}$ .

2. Enhancement of autonomous bulldozing efficiency through a data-driven optimal blade control strategy.
3. Sharing in an open-source repository the high-fidelity simulation setup and training pipeline to facilitate reproducibility and further research, see [link<sup>1</sup>](#).

## 2 Method

### 2.1 Reduced Order Modeling for Leveling Task

#### 2.1.1 Terramechanics Simulation

The “physics-based” attribute of the terramechanics model in 1) above refers to the simulator’s ability to capture soil deformation using fundamental continuum mechanics principles - conservation of mass, momentum, and stress evolution as shown in Eq. 1. The material model considered is hypo-elasto-plastic, which is a continuum model that captures the elastic and plastic deformation of the soil. The equations governing the

deformation of the soil are:

$$\frac{d\rho}{dt} = -\rho \nabla \cdot \mathbf{u} \quad (1a)$$

$$\frac{d\mathbf{u}}{dt} = \frac{\nabla \cdot \boldsymbol{\sigma}}{\rho} + \mathbf{f}_b \quad (1b)$$

$$\frac{d\boldsymbol{\sigma}}{dt} = \dot{\boldsymbol{\phi}} \cdot \boldsymbol{\sigma} - \boldsymbol{\sigma} \cdot \dot{\boldsymbol{\phi}} + \hat{\boldsymbol{\sigma}}, \quad (1c)$$

where  $\rho$  is the soil density,  $\mathbf{u}$  is the soil velocity field,  $\mathbf{f}_b$  is the body force,  $\hat{\boldsymbol{\sigma}}$  is the objective Jaumann stress rate tensor, and  $\dot{\boldsymbol{\phi}} = \frac{1}{2}(\nabla \mathbf{u} - \nabla \mathbf{u}^T)$  is the rotation rate tensor. We solve the system of partial differential equations above numerically using the Smoothed Particle Hydrodynamics method, which is a Lagrangian method. The simulation proceeds at time steps of approximately  $10^{-4}$  seconds, and the number of particles associated with the soil dynamics is in the millions. The main parameters associated with the soil model – friction angle, cohesion coefficient, and bulk density – are available based on the type of material that the bulldozer interacts with, e.g., sand, clay, loam, silts, snow. More parameters enter the model, but they are of secondary importance and can be safely approximated. For details, see [21].

The vehicle used was a four-wheel Gator with a two degrees-of-freedom blade – elevation and pitch. One actuator prescribes the blade’s pitch angle (in radians),  $b_p \in [0, \frac{\pi}{6}]$ , and another one governs the blade’s vertical movement (in meters),  $b_z \in [-0.075, 0.05]$ . At zero pitch, the blade is perpendicular to the ground; at zero vertical displacement, the blade is positioned 0.35m above ground level. The control ranges consider factors such as blade width, mounting position relative to the vehicle, and vehicle engine power. The soil piles in their initial state are generated using a 2D Gaussian distribution, with the height  $h$  at a point  $(x, y)$  given by:

$$h(x, y) = h_{max} \cdot \exp \left[ -\frac{(x - x_c)^2 + (y - y_c)^2}{2\sigma^2} \right], \quad (2)$$

where the pile center is at  $(x_c, y_c) = (2, 0)$  with a constant standard deviation of  $\sigma = 0.4\text{m}$ . The most important parameter is the maximum soil pile height  $h_{max}$ . This Gaussian distribution approach for soil pile modeling has been widely adopted in the literature [3, 11, 6, 12, 13].

### 2.1.2 Data Collection Process

Figure 2 illustrates the training data collection process. We split each operation into two intervals, each 2.5 seconds long. Each of the two intervals has blade control input  $(b_p, b_z)$ . The choice of two control intervals 2.5 seconds long considered both empirical evidence and data collection efficiency. Indeed, during real bulldozing operation, an operator would adjust the commands at about this frequency [10]. Intuitively, the more control intervals, the more effective and smoother the leveling process. However, the data collection process to support the training of the NN model would become challenging in terms of storage and number of simulations. To generate a diverse training set, 2000 random blade control pairs ( $b_p \in [0, \frac{\pi}{6}]$  and  $b_z \in [-0.075, 0.05]$ ) and three different initial pile heights,  $h_{max}$ , (0.3m, 0.4m, and 0.5m) were used, resulting in 6000 simulation tests for training data collection. Each simulation test yields two distinct training samples: the first captures transition from initial to intermediate state, and second from intermediate to final state, see Fig. 2). Each sample’s input consists of the pre-operation heightmap and two sets of random blade control parameters, while the output includes the resulting heightmap and the corresponding operation duration  $t_d$ . The data collection process required approximately 90 hours to complete, leveraging the ACCESS GPU cluster [29] to run multiple simulations

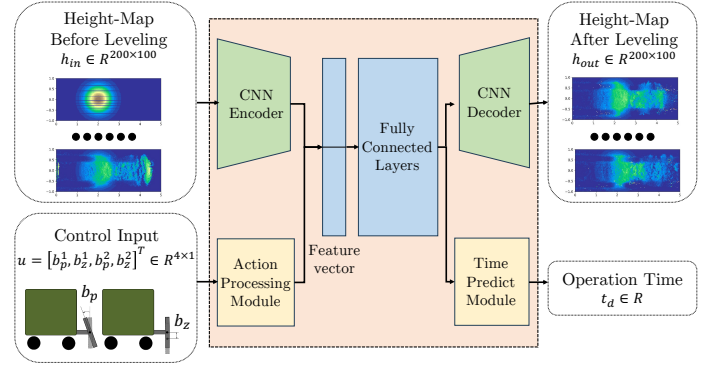


Figure 3: Architecture of the NN used for the gradient-based optimization.

in parallel with GPU-accelerated Chrono terramechanics modeling.

### 2.1.3 NN-based Differentiable Model




Figure 4: Model prediction performance comparison, showing the trade-off between model size and training loss. Although the steady-state training losses for the large and medium models are comparable, the medium model is preferred due to its superior inference speed and relative simplicity.

Figure 3 illustrates the selected NN architecture. The underlying principle is that NNs can learn to approximate complex physical interactions [30]. The initial terrain heightmap is fed into a Convolutional Neural Network (CNN) based encoder, while two sets of blade control inputs are processed by an action-processing module based on Fully Connected Layers (FCL). The choice to use CNN layers and an encoder-decoder structure is based on their proven ability to extract features from input images and predict output images [31, 32]. The image and action input streams are then concatenated to form a feature vector, which is processed through FCL-based layers to create a latent representation. This latent representation branches into two heads: one feeds into a CNN decoder to predict



the post-operation heightmap, while the other predicts the operation duration. The NN learning problem involves finding parameters  $\theta$  such that:

$$\{h_{out}, t_d\} = f_\theta(h_{in}, u), \quad (3)$$

where  $h_{in}, h_{out} \in \mathbb{R}^{200 \times 100}$  are heightmaps before and after the leveling operations, respectively,  $u \in \mathbb{R}^{4 \times 1}$  is a set of two pairs of blade control actions in different time intervals during operation (blue and purple section in Fig. 2), and  $t_d \in \mathbb{R}$  is operation duration. To determine an efficient NN structure (e.g., number of layers or neurons per layer in the encoder, decoder, and other modules), a comparative study was conducted with three different NN sizes: a large model (39 million parameters), a medium model (2 million parameters), and a small model (78 000 parameters). As depicted in Fig. 4, the medium model achieved a trade-off between training loss and model complexity. The NN model proposed diverges from Wagner et al. [6] in two aspects. Firstly, it leverages training data of enhanced physical fidelity from detailed Chrono simulations, which enables the NN to learn richer physics, encompassing the prediction of both resultant terrain heightmaps and operation times (Eq. 3). Secondly, the model adopts an operational-level prediction granularity, forecasting the outcome of an entire forward leveling pass rather than continuous, step-by-step terrain evolution. This operational focus, in turn, significantly reduces training data requirements—a critical factor given lengthy high-fidelity data generation—allowing for a streamlined NN architecture. This simpler architecture eschews time-dependent modules (e.g., LSTMs, RNNs), thereby accelerating design iterations and prototyping.

## 2.2 The Blade Optimal Control Problem

With the NN-based reduced-order model  $f_\theta$  of Eq. 3 capturing the leveling operation, a multi-objective gradient-based optimization problem is formulated to determine the optimal blade control actions  $u^*$ . Given  $h_{in}$ , the objective is to achieve a flat post-leveling terrain  $h_{out}$  that closely matches a desired post-leveling flat terrain profile  $h_{des} \in \mathbb{R}^{200 \times 100}$ , while minimizing the operation time  $t_d$ :

$$\begin{aligned} \min_{u^*} \mathcal{L}(h_{in}, u) &= \|h_{out} - h_{des}\|_2^2 + \lambda t_d \\ \text{s.t. } \{h_{out}, t_d\} &= f_\theta(h_{in}, u). \end{aligned} \quad (4)$$

The differentiability of the NN model  $f_\theta$  allows for the computation of the cost function's gradient with respect to the control actions  $u$ , enabling gradient-based

optimization. By adjusting the time penalty parameter  $\lambda$ , the optimization can prioritize either the flatness of the leveled ground, or the total operation time; i.e., it addresses the trade-off between operational duration and leveling quality.

The optimization is solved using the Adam algorithm, which has proven effective for non-convex problems [33, 34]. The blade control update at iteration  $t + 1$  is:

$$u_{t+1} = \text{Adam}(u_t, \nabla_u \mathcal{L}(u_t), \alpha, \beta_1, \beta_2, \varepsilon), \quad (5)$$

where  $\mathcal{L}(u)$  is the loss function in Eq. 4,  $\alpha$  is the learning rate,  $\beta_1 = 0.9$  and  $\beta_2 = 0.999$  are the exponential decay rates for the first and second moment estimates, and  $\varepsilon = 10^{-8}$  is a numerical stability term. Convergence typically requires 2000 to 3000 iterations – approximately 3.6 – 5.5 seconds, with a learning rate of  $5 \times 10^{-4}$ .

## 2.3 Measuring Metrics for Earth-moving Task

To evaluate leveling performance, we use the percentage change in Mean Squared Error ( $\Delta\text{MSE}$ ) between the initial and final terrain heightmaps, after one bulldozer pass. This metric quantifies how much the terrain reshaped toward the desired flat distribution  $h_{des} \in \mathbb{R}^{200 \times 100}$  relative to its original state. The Mean Squared Error (MSE) of a heightmap  $h$  relative to  $h_{des}$  is computed as  $\text{MSE}(h, h_{des}) = \frac{1}{N} \sum_{i,j} (h_{ij} - h_{des,ij})^2$ . If  $\text{MSE}_{\text{initial}} = \text{MSE}(h_{in}, h_{des})$  and  $\text{MSE}_{\text{final}} = \text{MSE}(h_{out}, h_{des})$ , the change toward the desired heightmap is

$$\Delta\text{MSE} (\%) = \frac{\text{MSE}_{\text{initial}} - \text{MSE}_{\text{final}}}{\text{MSE}_{\text{initial}}} \times 100\%. \quad (6)$$

A higher  $\Delta\text{MSE}$  value indicates better leveling performance, with negative values indicating degraded terrain conditions. The perfect MSE score is 100% if  $\text{MSE}_{\text{final}}$  equals zero, which means final heightmap perfectly matched the desired flat terrain distribution. This metric is aligned with the flatness objective in our optimization framework, Eq. 4.

## 3 Simulation Experiments

The autonomous leveling scheme employed is summarized in Fig. 2. The initial pile has random heights, and the optimization problem in Eq. 4 is solved to plan control action profiles. Subsequently, after the first leveling pass, the bulldozer receives an updated heightmap

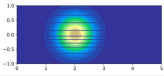
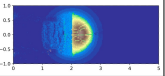
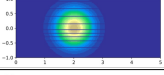
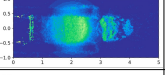
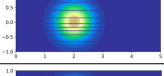
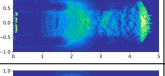
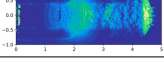
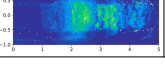
Terrain Heightmaps		$\Delta$ MSE	Scenarios
		-61.35%	Vehicle Immobilization
		23.82%	Insufficient Leveling
		26.20%	1st Optimized Level
		48.13%	2nd Optimized Level

Figure 5: Algorithm performance demonstration with an initial soil pile height of 0.37m. First row: A 'stuck' scenario where fixed blade parameters (control profile: red line, Fig. 6) result in leveling failure. Second row: Fixed blade setting without optimal control (control profile: green line, Fig. 6), resulting in insufficient material displacement despite completing one pass. Third row: First pass using the proposed algorithm ( $\Delta$ MSE=26.2%; control profile: blue line, Fig. 6). Fourth row: Second pass using the proposed algorithm, achieving a  $\Delta$ MSE of 48.13%.

reflecting the newly formed soil distribution and applies optimal blade control again to perform a second leveling operation. Currently, a specific stopping criterion for the operation has not been defined; instead, the bulldozer performs two forward leveling passes in each trial (forward-backward-forward as action sequence). This approach is adopted because leveling quality standards vary dynamically with construction requirements. Therefore, it is more pertinent to demonstrate the pipeline's capability to conduct recursive leveling operations based on updated terrain profile information. Low-level vehicle control commands were implemented for bulldozer movements such as moving forward, stopping at terrain boundaries, and moving backward to the initial position. Since this work centers on optimizing low-level blade control, the control policies for vehicle movement (such as throttle, braking, and steering), which were not the primary focus, remained consistent throughout all experiments. We present three categories of simulation results: a comparison of leveling quality with and without the proposed algorithm; an assessment of the algorithm's robustness across random initial pile configurations; and an analysis of the trade-off between leveling quality and operation duration.

### 3.1 Autonomous Leveling

Figure 5 compares the performance of the proposed algorithm against leveling with constant blade settings without active low-level adjustments. The latter approach leads to vehicle immobilization (vehicle gets stuck and spins its wheels in place), or scenarios where




Figure 6: Blade control action profiles logged during experiments referenced in Fig. 5 and Fig. 8.

only small amounts of soil are displaced. The proposed blade control strategy, which integrates data-driven modeling with gradient-based optimization, resulted in successful task completion and improved leveling metrics as quantified by the  $\Delta$ MSE score.

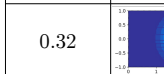
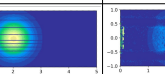
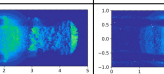
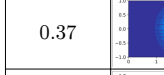

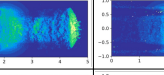
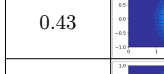

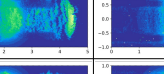
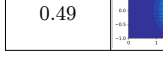
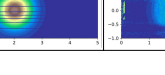
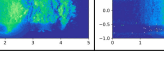
Height (m)	Initial	After 1st Level	After 2nd Level	$\Delta$ MSE
0.32				42.76%
0.37				48.13%
0.43				43.72%
0.49				35.75%

Figure 7: Algorithm performance across different initial pile heights.

To assess the algorithm's robustness, three additional tests were conducted with different random initial pile heights,  $h_{max}$ , supplementing the 0.37m pile height results of Fig. 5. As shown in Fig. 7, the proposed algorithm demonstrates effective performance across various initial pile configurations. The lowest  $\Delta$ MSE score was observed for the case with the highest initial pile. It is hypothesized that a larger volume of granular material presents a greater challenge for the leveling process, and an additional leveling pass might be beneficial in such scenarios.

### 3.2 The Trade-off Between Level Quality and Operation Time

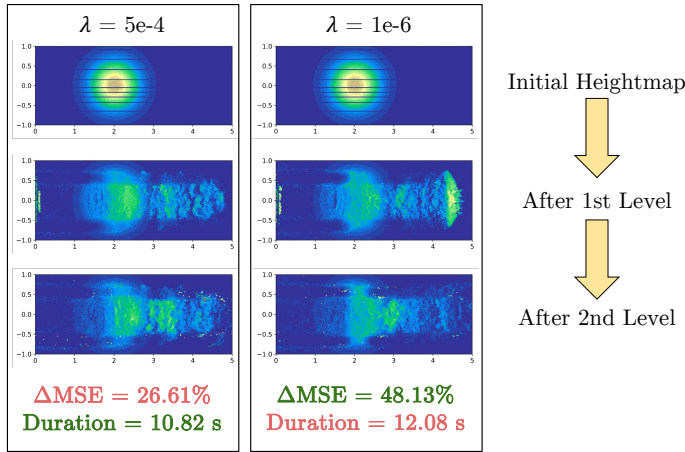


Figure 8: Illustration of the trade-off between leveling quality ( $\Delta\text{MSE}$ ) and operation time by varying the penalty parameter  $\lambda$ .

As per the optimization formulation in Eq. 4, a smaller value of the penalty parameter  $\lambda$  yields a solution that prioritizes terrain flatness over leveling speed. Figure 8 illustrates this trade-off. The default value for  $\lambda$  is  $1e-6$ , which produced the results shown in Fig. 5 and Fig. 7. When a larger  $\lambda = 5e-4$  is used (Fig. 8), the flatness score ( $\Delta\text{MSE}$ ) decreases to 26.61%, while the operation time is reduced by approximately 10%. In Fig. 6, the action profile for this faster first pass (larger  $\lambda$ ) is plotted as an orange line, contrasted with the blue line representing the slower leveling process (default  $\lambda$  value). The plots capture a trade-off between the two objectives: for slower, more thorough leveling (smaller  $\lambda$ ), the blade's vertical control tends to lower the blade to gather more material, using pitch control to “hold and spill” it effectively. Conversely, faster leveling (larger  $\lambda$ ) employs higher blade positions to minimize contact and resistance between the soil and the bulldozer.

## 4 Discussion

This work integrates physics-based terramechanics with vehicle simulation to enhance the design process for autonomous leveling policies. To mitigate the computational burden associated with direct simulation, a reduced-order, NN model was formulated, embedding knowledge and training experience derived from high-fidelity simulations. Leveraging the differentiability of the NN, a gradient-based optimization problem was formulated to control low-level blade actions for specific terrain profiles. The multi-objective optimization

framework provides flexibility to adjust blade control, balancing operational time and leveling quality. The proposed algorithm, by controlling low-level bulldozer blade actions to optimize the leveling process, complements existing research focused on global route planning for bulldozers [11, 3, 4]. While the existing global route planning solutions guide the bulldozer's global movement, our work concentrates on the blade's direct contact and interaction with the terrain. The entire pipeline – including the high fidelity simulator, NN-based reduced-order model, and the testing environment – has been open-sourced due to its relevance in other types of off-road robotics tasks (quadrupeds and bipeds, tracked vehicles, etc.) Indeed, integrating high-fidelity terramechanics and agent simulation into simulation-based autonomy design holds broad applicability.

Looking forward, developing a real-world testing facility to demonstrate the proposed algorithm's effectiveness combined with existing global routes planning algorithms in physical environments is a key priority. This will involve constructing a scaled bulldozer testbed, implementing a real-time ROS-based control framework, and bridging the simulation-to-reality (sim2real) gap, particularly for utilizing realistic depth sensor data [12, 11]. Secondly, the current state-of-the-art autonomous grading policies primarily address vehicle route planning. Combining the proposed low-level control with these advanced global planning algorithms could yield robust and efficient performance for real-world applications. Lastly, while this work utilized the NN-based reduced-order model within a gradient-based control framework, its computational efficiency also opens possibilities for experimentation with other control paradigms, such as Reinforcement Learning or Imitation Learning.

## Acknowledgments

The authors gratefully acknowledge the financial support provided by the National Science Foundation under grant number CMMI2153855. The authors also extend their appreciation to Komatsu America Corporation for their generous funding support and valuable technical insights that contributed to this research.

## Conflicts of Interest

The authors declare no conflict of interest.



## Data Availability Statement

The software reproducibility study and media demonstration is available in the following link.

## References

- [1] Y. Liu, A. AH, N. A. Haron, B. NA, H. Wang, Robotics in the construction sector: Trends, advances, and challenges, *Journal of Intelligent & Robotic Systems* **2024**, 110, 2 72.
- [2] N. Melenbrink, J. Werfel, A. Menges, On-site autonomous construction robots: Towards unsupervised building, *Automation in construction* **2020**, 119 103312.
- [3] Y. Osaka, N. Odajima, Y. Uchimura, Route optimization for autonomous bulldozer by distributed deep reinforcement learning, In *2021 IEEE International Conference on Mechatronics (ICM)*. **2021** 1–6.
- [4] M. Hirayama, M. Whitty, J. Katupitiya, J. Guivant, An optimized approach for automatic material distribution operations of bulldozers, *International Journal of Advanced Robotic Systems* **2018**, 15, 2 1729881418764716.
- [5] M. Hirayama, J. Guivant, J. Katupitiya, M. Whitty, Path planning for autonomous bulldozers, *Mechatronics* **2019**, 58 20.
- [6] W. J. Wagner, K. Driggs-Campbell, 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** 6531–6538.
- [7] S.-H. Kim, Y.-S. Lee, D.-I. Sun, S.-K. Lee, B.-H. Yu, S.-H. Jang, W. Kim, C.-S. Han, Development of bulldozer sensor system for estimating the position of blade cutting edge, *Automation in construction* **2019**, 106 102890.
- [8] Y.-S. Lee, S.-H. Kim, J. Seo, J. Han, C.-S. Han, Blade control in cartesian space for leveling work by bulldozer, *Automation in Construction* **2020**, 118 103264.
- [9] T. Mononen, J. Mattila, A. Kolu, Blade control for surface profile tracking by heavy-duty bulldozers, In *Fluid Power Systems Technology*, volume 85239. American Society of Mechanical Engineers, **2021** V001T01A017.
- [10] T. Mononen, A. Kolu, J. Mattila, Semi-autonomous bulldozer blade control using real-time terrain mapping, In *2022 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM)*. **2022** 775–782.
- [11] Y. Miron, C. Ross, Y. Goldfracht, C. Tessler, D. Di Castro, Towards autonomous grading in the real world, In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. **2022** 11940–11946.
- [12] Y. Miron, Y. Goldfracht, C. Ross, D. D. Castro, I. Klein, Autonomous dozer sand grading under localization uncertainties, *IEEE Robotics and Automation Letters* **2023**, 8, 1 65.
- [13] Y. Miron, D. Navon, Y. Goldfracht, D. D. Castro, I. Klein, Decentralized and asymmetric multi-agent learning in construction sites, *IEEE Open Journal of Vehicular Technology* **2024**, 5 1587.
- [14] M. Pla-Castells, I. Garcia, R. Martinez, Visual representation of enhanced sand pile models, In *Industrial Simulation Conference. Universidad Polit cnica de Valencia, Valencia, Spain*. **2003** 141–146.
- [15] M. Pla-Castells, I. Garc a-Fernandez, R. J. Martinez-Dura, et al., Physically-based interactive sand simulation., In *Eurographics (Short Papers)*. **2008** 21–24.
- [16] M. Nakatani, Z. Sun, Y. Uchimura, Autonomous grading work using deep reinforcement learning based control, In *IECON 2018 - 44th Annual Conference of the IEEE Industrial Electronics Society*. **2018** 5068–5073.
- [17] H. J. Suh, M. Simchowicz, K. Zhang, R. Tedrake, Do differentiable simulators give better policy gradients?, In K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, S. Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*. PMLR, **2022** 20668–20696, URL <https://proceedings.mlr.press/v162/suh22b.html>.
- [18] Y. Hu, L. Anderson, T.-M. Li, Q. Sun, N. Carr, J. Ragan-Kelley, F. Durand, DiffTaichi: Differentiable programming for physical simulation, *ICLR* **2020**.

- [19] C. Liu, T. Liu, K. Yang, K. Hong, X. Chen, Differentiable simulation and optimization of particle-fluid flows using graph neural networks, *Powder Technology* **2025**, 121143.
- [20] H. Unjhwala, L. Bakke, H. Zhang, G. Arivoli, M. Taylor, L. Yang, R. Serban, D. Negrut, A physics-based continuum model for versatile, scalable, and fast terramechanics simulation.
- [21] W. Hu, M. Rakhsha, L. Yang, K. Kamrin, D. Negrut, Modeling granular material dynamics and its two-way coupling with moving solid bodies using a continuum representation and the SPH method, *Computer Methods in Applied Mechanics and Engineering* **2021**, 385 114022.
- [22] D. Kochkov, J. A. Smith, A. Alieva, Q. Wang, M. P. Brenner, S. Hoyer, Machine learning-accelerated computational fluid dynamics, *Proceedings of the National Academy of Sciences* **2021**, 118, 21 e2101784118.
- [23] D. Coscia, N. Demo, G. Rozza, Generative adversarial reduced order modelling, *Scientific Reports* **2024**, 14, 1 3826.
- [24] A. Tasora, R. Serban, H. Mazhar, A. Pazouki, D. Melanz, J. Fleischmann, M. Taylor, H. Sugiyama, D. Negrut, Chrono: An open source multi-physics dynamics engine, In T. Kozubek, R. Blaheta, J. Šístek, M. Rozložník, M. Čermák, editors, *High Performance Computing in Science and Engineering*. Springer International Publishing, Cham, ISBN 978-3-319-40361-8, **2016** 19–491.
- [25] R. Serban, M. Taylor, D. Negrut, A. Tasora, Chrono::Vehicle template-based ground vehicle modeling and simulation, *International Journal of Vehicle Performance* **2019**, 5, 1 18.
- [26] R. Serban, L. Bakke, H. Unjhwala, D. Negrut, Design of a generic Chrono framework for fluid-solid interaction **2025**.
- [27] W. Hu, P. Li, A. Rogg, A. Schepelmann, S. Chandler, K. Kamrin, D. Negrut, A study demonstrating that using gravitational offset to prepare extraterrestrial mobility missions is misleading, *Journal of Field Robotics*.
- [28] R. Ozaki, T. Motomura, Y. Nakayama, S. Mori, Blade control for bulk dozer push by bulldozer using reinforcement learning, In *IECON 2024 - 50th Annual Conference of the IEEE Industrial Electronics Society*. **2024** 1–6.
- [29] National Science Foundation, Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS, <https://access-ci.org>, **2024**, Accessed: [28th November 2024].
- [30] A. Sanchez-Gonzalez, J. Godwin, T. Pfaff, R. Ying, J. Leskovec, P. Battaglia, Learning to simulate complex physics with graph networks, In *International conference on machine learning*. PMLR, **2020** 8459–8468.
- [31] X. Guo, W. Li, F. Iorio, Convolutional Neural Networks for Steady Flow Approximation, In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16. Association for Computing Machinery, New York, NY, USA, ISBN 978-1-4503-4232-2, **2016** 481–490, URL <https://doi.org/10.1145/2939672.2939738>.
- [32] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, In *International Conference on Medical image computing and computer-assisted intervention*. Springer, **2015** 234–241.
- [33] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, A. Lerer, Automatic differentiation in PyTorch, In *NIPS 2017 Workshop Autodiff*. **2017**.
- [34] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, *CoRR* **2017**, abs/1412.6980.