

Article World Modeling for Autonomous Wheel Loaders

Koji Aoshima ^{1,2}*^(D), Arvid Fälldin ²^(D), Eddie Wadbro ^{3,4}^(D) and Martin Servin ^{2,5}*^(D)

- ¹ Komatsu Ltd., 2-3-6, Akasaka, Minato-ku, Tokyo 107-8414, Japan
- ² Department of Physics, Umeå University, SE-901 87 Umeå, Sweden; arvid.falldin@umu.se
- ³ Department of Mathematics and Computer Science, Karlstad University, SE-651 88 Karlstad, Sweden; eddie.wadbro@kau.se
- ⁴ Department of Computing Science, Umeå University, SE-901 87 Umeå, Sweden
- ⁵ Algoryx Simulation AB, Kuratorvägen 2B, SE-907 36 Umeå, Sweden
- * Correspondence: koji.aoshima@umu.se (K.A.); martin.servin@umu.se (M.S.)

Abstract: This paper presents a method for learning world models for wheel loaders performing automatic loading actions on a pile of soil. Data-driven models were learned to output the resulting pile state, loaded mass, time, and work for a single loading cycle given inputs that include a heightmap of the initial pile shape and action parameters for an automatic bucket-filling controller. Long-horizon planning of sequential loading in a dynamically changing environment is thus enabled as repeated model inference. The models, consisting of deep neural networks, were trained on data from a 3D multibody dynamics simulation of over 10,000 random loading actions in gravel piles of different shapes. The accuracy and inference time for predicting the loading performance and the resulting pile state were, on average, 95% in 1.2 ms and 97% in 4.5 ms, respectively. Long-horizon predictions were found feasible over 40 sequential loading actions.

Keywords: wheel loader; earthmoving; automation; bucket-filling; world modeling; deep learning; multibody simulation

check for **updates**

Citation: Aoshima, K.; Fälldin, A.; Wadbro, E.; Servin, M. World Modeling for Autonomous Wheel Loaders. *Automation* 2024, *5*, 259–281. https://doi.org/10.3390/ automation5030016

Academic Editor: Felipe Martins

Received: 27 May 2024 Revised: 2 July 2024 Accepted: 4 July 2024 Published: 6 July 2024



Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https:// creativecommons.org/licenses/by/ 4.0/).

1. Introduction

The advances in artificial intelligence suggest that computerized control of construction and mining equipment has the potential to surpass that of human operators. Fully or semi-autonomous wheel loaders, not relying on experienced operators, can be an important solution to the increasing labor shortage. Wheel loaders typically operate on construction sites and quarries, repeatedly filling the bucket with soil and dumping it onto load receivers. They should move mass at a maximum rate with minimal operating cost without compromising safety. Recent research on autonomous loading control has focused on increasing performance and robustness using deep learning to adapt to soil properties [1–7]. However, these studies have been limited to a single bucket filling and have not considered the task of sequential loading from a pile. A challenge with sequential loading is that the pile state changes with every loading action [8–10]. The altered state affects the possible outcomes of the subsequent loading process and, ultimately, the total performance. A greedy strategy of always choosing the loading action that maximizes the performance for a single loading might be sub-optimal over a longer horizon. End-to-end optimization thus requires the ability to predict the cumulative effect of loading actions over a sequence of tasks. This involves having a model of the world and how it changes under the selected actions.

This paper introduces wheel loader world models for predicting the outcome of a loading action given the shape of the pile. The outcome includes the loaded mass, loading time, and work, as well as the resulting pile shape. The net outcome of a sequence of loading actions can then be predicted by repeated inferences of the model on the pile state, thus predicting its sequential evolution as well. The model aims at supporting optimal planning for autonomous wheel loaders, with the best sequence of loading actions computed from the initial pile surface only. We imagine that the plan would be updated with some regularity (e.g., daily, hourly, or after each loading cycle) from a new observation of the pile surface. Depending on the planning horizon and dimensionality of the action space, the optimizer requires numerous evaluations of the world model in a short time. Although a simulator based on multibody dynamics and the discrete element method (DEM) can predict the outcome of a loading [11], it is far too computationally intensive and slow for this optimization problem. Instead, we explore using a simulator to produce ground truth data and trained deep neural networks for "instantaneous" prediction of the outcome of particular loading actions on piles of different shapes. The learned model is informed specifically of the physics of the particular wheel loader and the type of soil used in the simulator.

The main contribution of this paper is a methodology for learning wheel loader world models and an investigation of how sensitive the model accuracy, inference speed, and required amount of training data are to the model complexity. The models predict the loading outcome, which includes the resulting pile state, loaded mass, time, and work, given the pile's initial state and a selected loading action. The net performance of a sequence of loading actions can thus be predicted by repeated model inference. The world models are shown to be differentiable with respect to the action parameters. This enables the use of gradient-based optimization algorithms for the planning of action sequences. The paper ends with an analysis of the models' computational properties, leaving the solution of the associated optimization problem to future work.

2. Related Work

Previous scientific work on data-driven models for predicting bucket-soil interactions have studied this from the perspective of automatic control [12–15] and bucket motion planning [16,17]. Model predictive control of earth-moving operations relies on a model to predict the dig force or soil displacement given a control signal acting over some time horizon and the current (and historic) system state [12,13]. A variational autoencoder (VAE) [15] and convolutional autoencoder [14] were used for learning a reduced representation of the terrain surface and combined with recurrent long short-term memory and mixture density neural networks to learn the time-evolution of the surface during an earth-moving task. Unfortunately, the errors grow during rollout, unless intermediate observations are added, making the resulting state unreliable and the method unstable for long horizons. Other researchers have developed models that predict the bucket fill factor and accumulated work from a planned bucket trajectory and initial soil pile shape [16,17]. This approach is infeasible in practical use-cases, where it is not possible to track a prescribed bucket trajectory [18]. In the present paper, we draw inspiration from the idea of using a VAE for representing the local soil surface but we focus on accurately predicting the end-state after breakout and on the accumulated mass, time, and work of the loading cycle. Instead of using a parameterized bucket trajectory as input, we use the control parameters of a force-based automatic bucket-filling controller.

3. Methodology

The method of developing a wheel loader world model includes creating a simulator [11,19] for the particular machine and its environment. The simulator is based on contacting 3D multibody dynamics with a real-time deformable terrain model that has been shown to produce digging forces and soil displacements with an accuracy close to that of resolved discrete elements and coupled multibody dynamics [20]. The simulated wheel loader is equipped with a type of admittance controller [21] for automatic bucket-filling, parameterized by four action parameters that determine how the boom and bucket actuation respond to the current dig force. An annotated dataset is created from simulated loading cycles carried out on soil piles of different shapes using different combinations of control parameters. Each simulation results in one data point, consisting of the pile's local heightmap before and after loading, loaded mass, time and work, and the set control parameters that were used. Models are trained to predict how the loading outcome depends on the input values. The models are finally tested on validation data withheld during training. The validated models can then be used to predict the result of new loading operations or be embedded in optimization routines for finding sequences of loading actions that are optimal with respect to productivity or energy efficiency. The method is illustrated in Figure 1.



Figure 1. Overview of the wheel loader world model, its development process, and intended usage. First, a dataset of the outcome of parametrized wheel loading actions *a* on a pile with local shape *h* is collected using a simulator. Two models are trained. One model predicts the expected loading performance \mathcal{P} in terms of loaded mass, time, and work, given *a* and *h*. Another model predicts the new shape *h'* that the pile transitions into after the completed loading. The outcome of a sequence of loading actions on a pile with global state *H* can be predicted by repeated model inference given the dig location *x* and heading *t* for each loading.

The learned models were investigated for accuracy, inference speed, and the required amount of training data. In general, model accuracy can be improved by increasing the input and output dimensions and the number of internal model parameters. However, larger models normally require more training data and may have lower inference speeds. To quantify the difference, we developed both a high-dimensional and a low-dimensional model using different representations of the pile state. The high-dimensional model takes a well-resolved heightmap of the pile state as input. This may capture many details of the pile surface at the cost of additional model parameters associated with convolutional layers. The low-dimensional model takes a heavily reduced representation of the pile surface, involving only four scalar parameters for its slope and curvature.

4. Wheel Loader World Models

4.1. Preface and Nomenclature

This paper defines the loading cycle as starting with a machine heading in direction t to dig location x in a pile, represented by the current pile state H. The cycle ends when the machine has returned to its initial location after digging in, filling the bucket,

breaking out, and leaving the pile in a new state H'. The wheel loader is equipped with an automatic bucket filling controller that is parametrized by some set of *action parameters* a. The controller is engaged when the machine reaches the dig location x. Each loading cycle can be assigned a *performance* \mathcal{P} , which in this study includes the loaded mass, loading time, and mechanical work. The performance is a consequence of the dynamics of the machine and the soil under the selected action a and the initial state of the pile. The expected performance of a loading cycle, indexed by $n \in \mathbb{N}$, is therefore given by some unknown *performance predictor* function Ψ . In other words,

$$\widehat{\boldsymbol{P}}_n = \boldsymbol{\Psi}(H_n, \boldsymbol{x}_n, \boldsymbol{t}_n, \boldsymbol{a}_n), \tag{1}$$

where we use the *hat* to distinguish predictions from actual values. The net effect of a sequence of *N* loading cycles is the accumulated outcome of a sequence of loading actions, a_1, \ldots, a_N , at pose $(x_1, t_1), \ldots, (x_N, t_N)$ applied on a sequence of pile states H_1, \ldots, H_N . Each loading cycle transforms the pile from its previous state to the next, $H_n \rightarrow H'_n \equiv H_{n+1}$, according to some unknown *pile state predictor* function Φ :

$$H'_n = \mathbf{\Phi}(H_n, \mathbf{x}_n, \mathbf{t}_n, \mathbf{a}_n). \tag{2}$$

The net outcome of N sequential loading actions is thus given by the sum

$$\widehat{\boldsymbol{\mathcal{P}}}_{1:N} \equiv \sum_{n=1}^{N} \widehat{\boldsymbol{\mathcal{P}}}_{n}(\widehat{H}'_{n-1}, \boldsymbol{x}_{n}, \boldsymbol{t}_{n}, \boldsymbol{a}_{n})$$
(3)

with initial pile state $H'_0 = H_1$. The predictions are associated with some error, that accumulates over repeated loadings from the evolving pile. The *accumulated error* in the pile state and loading performance over a horizon of *N* cycles is

$$\boldsymbol{\mathcal{E}}_{1:N}^{\hat{H}'} \equiv \left\| \hat{H}_N' - H_N' \right\|,\tag{4}$$

$$\boldsymbol{\mathcal{E}}_{1:N}^{\widehat{\boldsymbol{\mathcal{P}}}} \equiv \sum_{n=1}^{N} \left\| \widehat{\boldsymbol{\mathcal{P}}}_n - \boldsymbol{\mathcal{P}}_n \right\|.$$
(5)

When solving the problem of finding the sequence of dig locations and loading actions that maximize the net performance $\hat{\mathcal{P}}_{1:N}$, it is beneficial to use gradient-based optimization methods. These require the sought function Ψ to be differentiable with respect to *a*.

4.2. Global and Local Pile State

We represent the global pile state and the surrounding terrain by a height surface function z = H(x, y) in Cartesian coordinates. When making sequential predictions, the aim is to accurately predict the global state of the pile, which may take any shape consistent with the physics of the soil. To simplify matters, the predictor models take only the *local pile state* as input, assuming that the outcome of a loading depends only on the local state and not on the global shape of the pile. A dig location, x = [x, y] and heading t defines a local frame with basis vectors $\{e_x, e_y, e_z\}$, where e_x is aligned with the dig direction, and e_z is the vertical direction aligned with the gravitational field. We represent the local pile state with a discrete heightmap of $I \times J$ grid cells with the side length Δl . The heightmap, in the local frame, is

$$h_{ij} = H(\boldsymbol{x} - \delta \boldsymbol{e}_x + i\Delta l \boldsymbol{e}_x + j\Delta l \boldsymbol{e}_y) - H(\boldsymbol{x} - \delta \boldsymbol{e}_x)$$
(6)

with integers $i \in [0, I]$ and $j \in [-J/2, J/2]$ and constant displacement δ that ensures a certain fraction of the ground in front of the pile is present in the heightmap. The length and width of the heightmap, $I\Delta l$ and $J\Delta l$, must be large enough to cover the interaction region; see Section 6.3.

The simplified problem of predicting the loading outcome, Equations (1)–(3), using the local heightmap is

$$\widehat{\boldsymbol{\mathcal{P}}}_n = \psi(\boldsymbol{h}_n, \boldsymbol{a}_n), \tag{7}$$

$$\hat{h}_n' = \phi(h_n, a_n), \tag{8}$$

$$\widehat{\boldsymbol{\mathcal{P}}}_{1:N} \equiv \sum_{n=1}^{N} \widehat{\boldsymbol{\mathcal{P}}}_{n}(\widehat{\boldsymbol{h}}_{n-1}', \boldsymbol{a}_{n}).$$
(9)

with *local* predictor functions ψ and ϕ for the performance and pile state. The computational process is described in Algorithm 1 and illustrated in Figure 1.

Algorithm 1: Long-horizon prediction using world models						
Initialization: $\hat{H}'_0 = H_1$						
for $n = 1, \ldots, N$ do						
$x_n, t_n \leftarrow$ select dig pose along edge of \widehat{H}'_{n-1}						
$h_n \leftarrow ext{get local heightmap of } \widehat{H}'_{n-1} ext{ at pose } x_n, t_n$						
$a_n \leftarrow$ select appropriate dig action for h_n						
if high_dim_model then						
$\widehat{\boldsymbol{\mathcal{P}}}_n := \psi^{\text{high}}(\boldsymbol{h}_n, \boldsymbol{a}_n)$						
$\hat{h}'_n := \phi(h_n, a_n)$						
$\widehat{H}'_n \leftarrow ext{apply new local heightmap } \hat{h}'_n ext{ to } \widehat{H}_n ext{ at pose } x_n, t_n$						
else if low_dim_model then						
$\tilde{h}_n \leftarrow \text{compute local pile characteristics of } h_n$						
$\widehat{\boldsymbol{\mathcal{P}}}_n := \psi^{\text{low}}(\widetilde{\boldsymbol{h}}_n, \boldsymbol{a}_n)$						
$\widehat{H}'_n \leftarrow$ apply cellular automata to \widehat{H}_n after removing mass ($\widehat{\mathcal{P}}_n$) at x_n, t_n						
end						
$\widehat{\boldsymbol{\mathcal{P}}}_{1:N} = \sum_{n=1}^{N} \widehat{\boldsymbol{\mathcal{P}}}_n$						

4.3. Local Pile Characteristics

As an alternative to representing the local pile state in terms of a heightmap, we introduce a low-dimensional characterization, $\tilde{h} \equiv [\theta, \alpha, \kappa_x, \kappa_y]$, in terms of four scalar quantities: slope angle θ , incidence angle α , longitudinal curvature κ_x , and lateral curvature κ_y . These aim to capture the essence of the local pile shape from the perspective of bucket-filling [22,23]. First, the mean unit normal n is computed from the local heightmap. The slope angle relative to the horizontal plane is then computed $\theta \equiv \cos^{-1}[n \cdot e_z]$. The incidence angle is the angle between the attack direction $e_x = t/||t||$ and the pile normal projected on the horizontal plane $n_{\perp} = n - (n \cdot e_z)e_z$, that is, $\alpha \equiv \cos^{-1}([n_{\perp} \cdot e_x]/||n_{\perp}||)$. See Figure 2 for an illustration.

Taking inspiration from [24], we calculate the local mean curvature in the e_x and e_y directions by fitting a quadratic surface $h(x) \approx b + c^T x + \frac{1}{2}x^T Qx$ with surface parameters $Q = \text{diag}[-\kappa_x, -\kappa_y], c \in \mathbb{R}^2$, and $b \in \mathbb{R}$. This sign convention makes the curvatures positive for convex pile shape, which is the recommended shape for high-performance loading [23].



Figure 2. The mean normal *n* of the local heightmap *h* defines a slope angle θ relative to the horizontal plane. The attack angle α is the angle between the dig direction *t* and the normal projected onto the horizontal plane.

4.4. Performance Predictor Model

We developed two different models to predict the loading performance, referred to as the high-dimensional model and the low-dimensional model, respectively. The *high-dimensional model*, $\psi^{\text{high}}(h, a)$, takes the local heightmap h as input, while the *low-dimensional model*, $\psi^{\text{low}}(\tilde{h}, a)$, uses the local pile characteristics \tilde{h} . Both models take the same action parameters a as input, and they both output a performance vector $\hat{\mathcal{P}}$. The name distinction comes from dim $(h) \gg \dim(\tilde{h})$. The difference is reflected in the different neural network architectures of the models (Figure 3). The high-dimensional model uses three convolutional layers and a fully connected linear layer to encode h into a latent vector of length 32 before it is concatenated with a. h is interpolated to a size of 32×32 before encoding. They are followed by batch normalization and max pooling with window size 2×2 . The activation function is subject to hyperparameter tuning. The concatenation steps in both ψ^{high} and ψ^{low} are followed by multilayer perceptrons (MLPs) of identical architectures that are also subject to hyperparameter tuning.



Figure 3. Illustration of the model architectures. In the high-dimensional model (**a**), a convolutional neural network encodes the heightmap before being fed with the action vector to an MLP. In the low-dimensional model (**b**), the pile characteristics and action vector are inputs directly to the MLP.

4.5. Pile State Predictor Model

The pile state predictor model combines a VAE architecture with an MLP to make predictions via three steps: First, an encoder network compresses the initial pile state $h \in \mathbb{R}^{52 \times 52}$ into a lower-dimensional, regularized latent representation $z \in \mathbb{R}^{64}$. Next, given z along with the loading action a and a scale factor $\Delta h \equiv \max(h) - \min(h)$, an MLP predicts the regularized latent representation of the new pile state z'. Last, a decoder network constructs the resulting pile state $\hat{h}' \in \mathbb{R}^{52 \times 52}$ from z'. Figure 4 illustrates the inference process. For the VAE encoder/decoder blocks, we use the same CNN architecture as in [25].

The MLP block consists of two hidden layers with 1024 nodes and uses Leaky ReLU activation. Note that *h* is interpolated to size 64×64 before encoding and back to 52×52 after decoding to fit the off-the-shelf architecture.



Figure 4. The model architecture of the pile state predictor model. A VAE is paired with an MLP that learns state transition in the latent space.

4.5.1. Post-Processing

VAEs are known to produce somewhat blurry images [26]. In our case, this led to a loss of detail in the decoded pile states. In our preliminary tests, we found that we could improve the prediction accuracy on average by reusing some information from the initial pile state. By interpolating between h and \hat{h}' along the edges, the details of the outer area are retained. This makes the predicted local pile states blend seamlessly into the global pile state after the substitution. Figure 5 illustrates the post-processing.



Figure 5. The post-processing interpolates between the interior of the predicted \hat{h}' and the edges of the initial heightmap h to preserve the details along the boundary. The pile in the heightmaps, color-coded by height, was dug from the left.

4.6. Low-Dimensional Pile State Prediction Using Cellular Automata

The low-dimensional predictor model did not have direct access to the local heightmap. Instead, we constructed a pile state predictor method that acts directly on the global pile state. It works in two stages. First, the predicted load mass (part of the performance vector $\hat{\mathcal{P}}_n$) is removed from the global pile state at the dig location. This is performed by computing the corresponding volume and subtracting this from the global pile along a strip as wide as the bucket, starting at x_n and stretching along t_n . In the second stage, soil mass is redistributed to eliminate any slopes steeper than the set angle of repose. The mass-preserving cellular automata algorithm in [27] was employed, with the "velocity of flowing matter" set to $z_+ = 0.2\Delta l$.

4.7. Delimitations

The predictor models developed in this paper are limited to a single type of soil, gravel, which was assumed to be homogeneous. However, the process of learning world models is applicable to any other soil type supported by the simulator. This paper does not consider soil spillage on the ground after breaking out with an overfilled bucket. Spillage would affect subsequent performance, either by causing a loss in control precision or traction or by requiring the ground to be cleared. In this paper, "loading" refers to the bucket-filling phase, including approaching the pile straight ahead and reversing in the opposite direction after the breakout, following the definition in [16]. To account for the full loading performance, one should incorporate the V-cycle maneuver, emptying the bucket into the load receiver and clearing spillage from the ground.

5. Simulator and Loading Controller

5.1. Simulator

We collected data using a simulator developed [11] and validated [19] in a previous work. Images and videos from the simulator are found in Figures 1, 6 and Supplementary

Video S1. The simulator combines a deformable terrain model with a wheel loader model and runs approximately in real time. The terrain model, introduced in [20], combines the representation of soil as a continuous solid, distinct particles, and rigid multibodies. When a digging tool comes in contact with the terrain, the active zone of moving soil is predicted and resolved in terms of particles using DEM. When particles come to rest on the terrain surface, they merge back into the solid soil. The computational processes preserve the total mass. The wheel loader is modeled as a rigid multibody system, matching roughly the dimensions and physical properties of a Komatsu WA320-7. The model has five actuated joints, with hinge motors powering the driveline and steering and linear motors for the boom and bucket cylinders. The motors are controlled by specifying a momentaneous joint target speed and force limits. The actuators will run at the set target speed only if the system dynamics and required constraint force do not exceed their respective force limits. The limits are set to match the strength of the real machine [28]. The transmission driveline model includes the main, front, and rear differential couplings to the wheels. The force interaction between the vehicle and the terrain occurs through the tire-terrain and bucket-terrain contacts. The resulting forces on the bucket depend on the shape and amount of active soil, its dynamic state, and mechanical properties, including bulk density, internal friction angle, cohesion, and dilatancy angle. As in [19], these are set to the values 1727 kg/m^3 , 32° , 0 kPa, and 8° , which are intended to reflect the properties of dry gravel. The simulations were run with the physics engine AGX Dynamics [29] using a 0.01 s time-step and terrain grid cell-size of $0.1 \text{ m} \times 0.1 \text{ m}$.



Figure 6. Image sequence of a simulated loading cycle. A video version is found in Supplementary Video S1.

5.2. Loading Controller

The loading scenario starts with the machine driving at 8 km/h towards a pile 5 m from the target dig location. The bucket is lowered and held level to the ground during the approach. The target drive velocity is kept constant during the bucket-filling phase, although the actual velocity will vary due to digging resistance. Throughout the cycle, the machine maintains the same heading. When the bucket reaches the set dig location, the automatic bucket-filling controller is engaged. This controller was inspired by the admittance controller from [21], which regulates the bucket actuator velocity using the measured boom cylinder force. Our admittance controller uses the same method but applies it to both the boom and bucket actuators. The controller determines the target velocities of the boom and bucket cylinders, v_{bm} and v_{bk} , as follows:

$$v_{\rm bm} = \operatorname{clip}(k_{\rm bm}[f_{\rm bm}/f_{\rm b0} - \delta_{\rm bm}], 0, 1)v_{\rm bm}^{\rm max}, \tag{10}$$

$$v_{bk} = \operatorname{clip}(k_{bk}[f_{bm}/f_{b0} - \delta_{bk}], 0, 1)v_{bk}^{\max},$$
(11)

where f_{bm} is the momentaneous boom cylinder force and clip(*value, min, max*) limits *value* to the maximum and minimum values. The ramp function has four parameters (velocity gains k_{bm} and k_{bk} and force thresholds δ_{bm} and δ_{bk}) that parameterize the behavior of the controller. The actuator max speeds, v_{bm}^{max} and v_{bk}^{max} , and the normalizing boom force f_{b0} are set using specifications from the manufacturer. The control parameters are collected in the action vector

$$\boldsymbol{u} = [\delta_{\rm bm}, k_{\rm bm}, \delta_{\rm bk}, k_{\rm bk}]. \tag{12}$$

Parameters δ_{bm} and δ_{bk} regulate what force magnitude is required to trigger the lift and tilt reactions, while parameters k_{bm} and k_{bk} control how rapid the reaction is. Different values of the control parameter thus render different scooping motions, as illustrated by the examples in Figure 7. The challenge is to select the parameters most appropriate for the local pile state. Note that the controller operates only on the lift and tilt actuators. The vehicle keeps thrusting forward to reach the set target drive velocity. If there is insufficient traction, the wheels may slip.

The bucket-filling controller is stopped if the bucket achieves the tilt end position, reaches a penetration depth of 3.2 m, breaks out of the initial surface, or if the loading duration exceeds 15 s. The bucket is then tilted with maximum speed to its end position (if not there already). The brake is applied for at least 0.5 s to let the agitated soil settle. After that, the vehicle is driven in reverse with a target speed of 8 km/h, lift $v_{bm} = 0.6v_{bm}^{max}$ and tilt $v_{bk} = 0.6v_{bk}^{max}$ to reach a boom angle of -20° relative to the horizontal axis, and the bucket tilt ends. The scenario ends when the vehicle has reversed 5 m from the dig location. To simplify time and energy comparison between simulations, they all use the same start and end distance from the dig location and all end with identical boom and bucket angles. The target speeds and force measurements are smoothed using a 0.1 s moving average to avoid a jerky motion.



Figure 7. Examples of simulated trajectories using two different sets of action parameters $a_1 = [0.1, 1.0, 0.1, 2.5]$ and $a_2 = [0.5, 1.0, 0.4, 4.0]$ on five piles with a slope ranging between 30° and 10°.

6. Data Collection and Model Training

This section describes how data were collected from the simulator and used for learning and evaluating the predictor models.

6.1. Loading Outcome

The models were developed to predict the outcome of a loading cycle in terms of the performance \mathcal{P} and the resulting local pile state h' given the initial pile h and action a. The performance was measured using the three essential scalar metrics, namely load mass (ton), loading time (s), and work (kJ). Hence, $\mathcal{P} \in \mathbb{R}^3$. The load mass is measured as the amount of soil the bucket carries at the end of each loading cycle. The loading time is the time elapsed between each loading cycle's start and end. The work is the energy consumed by the boom and bucket actuators and the forward drive. It includes the energy required to fill the bucket, break out, raise the bucket, and accelerate the vehicle and soil. Much of the work is lost to frictional dissipation internally in the soil and at the bucket–soil interface. The physics-based simulator accounts for this dissipation. Energy dissipation in the vehicle engine and hydraulics is not included and should be added if needed.

6.2. Data Collection

We collected a total of 10,718 samples, $\{h_n, a_n, h'_n, \mathcal{P}_n\}_{n=1}^{10,718}$, by repeating the loading scenario (Section 5.2) in the simulation. First, six different initial piles were prepared: two triangular, two conical, and two wedged with respective slope angles of 20° and 30°. These are illustrated in Figure 8. Perlin noise [30] was applied to the surfaces to increase the variability in pile shape. For each of the six initial piles, 30 consecutive loading cycles were simulated. Each loading cycle produced one data point $(h_n, a_n, h'_n, \mathcal{P}_n)$. Since the resulting pile state was used as the initial state in the next loading cycle, a variety of piles of different shapes was achieved. This process was repeated 60 times for each of the six seed piles.



Figure 8. Six initial seed piles are used, triangular, conical, and wedged, with slope angle θ . The dig location *x* and heading are randomized at each loading. The grayscale encodes the pile surface height.

In each simulation, a random dig location was selected, and the wheel loader was positioned at x_n , a distance of 5 m from the pile, and the heading was given a random disturbance of $\pm 10^\circ$; see Figure 8. Each loading used a set of action parameters, a_n , randomly sampled by Latin Hypercube Sampling (LHS) with the ranges $0.0 \le \delta_{bm}$, $\delta_{bk} \le 0.7$ and $0.0 \le k_{bm}$, $k_{bk} \le 5.0$. Note that the actual slope angle θ and incidence angle α varied according to the dig location because of the Perlin noise.

6.3. Local Heightmap Settings

We found that the performance and pile state predictor models benefited from using different sizes of the local heightmap and displacement δ in the dig direction. We used a 3.6 m sided quadratic heightmap discretized by 36×36 grid cells for the performance predictor. The size leaves a margin of 0.5 m on the sides of the bucket. Due to avalanching, the pile state predictor needed a larger heightmap to capture the state that the pile finally settles into after the breakout. For this, we used a heightmap with 5.2 m side lengths, discretized in 52×52 cells, and displaced the heightmap 1.0 m towards the approach direction.

6.4. Model Training

This section describes the training processes of the performance predictor (Section 6.4.1) and the pile state predictor model (Section 6.4.2). The models were implemented using PyTorch.

6.4.1. Performance Predictor Model

The models were trained until convergence with a mean squared error (MSE) loss using the Adam optimizer using a learning rate of 10^{-5} . Hyperparameter tuning was performed via grid search over the number of hidden layers (1, 2, or 3), the number of units $(2^3, 2^4, \ldots, 2^{12})$ in each hidden layer with 0.1 dropout rate, as well as the activation function in fully connected and convolutional layers (Leaky ReLU or Swish [31]). The dataset was first min-max normalized. The training and validation set (split 90/10) was sequentially increased in size from 100 to 9646 samples, while the test set size was fixed at 1072 samples.

6.4.2. Pile State Predictor Models

The pile state predictor model was trained in a two-stage process. First, we trained the VAE to perform heightmap reconstruction. Before training, the heightmaps were re-scaled by subtracting the average slope and applying min-max scaling to them individually:

$$h_n := h_n - \frac{1}{N} \sum_{m=1}^N h_m$$
(13)

$$h_n := \frac{h_n - \min(h_n)}{\max(h_n) - \min h_n}.$$
(14)

The standardization in Equation (13) centers each height distribution around 0 across the entire data set before applying the min-max scaling in Equation (14). We found this helps to preserve the pile–ground boundary in our reconstructions. The normalization in Equation (14) makes the VAE agnostic to scale, which we found increases its overall performance. The VAE was trained using the Adam optimizer with a learning rate of 10^{-3} . We used a weighted sum of the element-wise MSE and the Kullback–Leibler Divergence (KLD), $\mathcal{L}_{VAE} = \mathcal{L}_{MSE} + 0.1 \mathcal{L}_{KLD}$ for the loss function.

In the second training phase, we used the encoder trained in the first phase to encode sampled pile state pairs (h, h') into latent pairs (z, z'). We then trained an MLP on the latent pairs with a 0.1 dropout rate for the hidden layers and used the Adam optimizer with a learning rate 10^{-5} .

The full dataset was used with a split ratio of 80/10/10 in training, validation, and test data for the pile state predictor. Since the wheel loader geometry and actions are left-right symmetric, we applied random reflections to the heightmaps to augment the dataset.

7. Results

The models for predicting the loading performance and resulting pile state were evaluated separately on the hold-out test data. The best models were selected and used in combination to test their ability to predict the outcome of sequential loadings. An overview of the full dataset is given in Appendix A.

7.1. Performance Predictor Model

In total, we developed 480 model instances during the hyperparameter tuning for both the high- and low-dimensional performance predictors, ψ^{high} and ψ^{low} . The models were evaluated using three metrics: mean relative error (MRE), training time, and inference speed. The MRE is relative to the simulation ground truth and is calculated from the average of five distinct training results. The training time was counted as the time per epoch (number of epochs ranging up to about 2000), calculating the median over the entire training. The inference speed is measured as the model execution time, taking the average of 1000 model executions. We selected the MLPs with two hidden layers with 2048 units as models of special interest. These are denoted $\psi_{\diamond}^{\text{high}}$ and $\psi_{\star}^{\text{high}}$, where the difference is the use of Swish and Leaky ReLU for activation functions. These models have roughly 10^7 model parameters, including the convolutional layers. The loading performance MRE, listed in Table 1, ranges between 3.5 and 7% with insignificant dependence on the activation function. For the low-dimensional model, the selected models of special interest, $\psi_{\circ}^{\text{low}}$ and $\psi_{\star}^{\text{low}}$, are smaller with two hidden layers and 512 units per layer. This amounts to roughly 5×10^5 parameters (no convolutional network is involved). It was trained on a smaller dataset with only 3000 samples, as adding additional samples did not significantly improve the performance. The loading performance MRE for the low-dimensional model saturates in the range of about 5.5–8.5%, in other words, with 20–70% larger errors than for the high-dimensional model. The training time and inference time are found in Table 1. The low-dimensional model trains about seven times faster, and inference runs three times faster than the selected high-dimensional model. The computational overhead of using Swish over Leaky ReLU is marginal. The high- and low-dimensional models have a memory footprint of 16.7 Mb and 0.027 MB, respectively. The effect of hyperparameters on the model performance is found in Appendix B.

Table 1. Key properties of the selected performance predictor models.

	$\psi^{ extsf{h}}$	igh	$\psi^{ m low}$			
	L ReLU	Swish	L ReLU	Swish		
mass MRE [%]	4.47	5.70	7.77	7.58		
time MRE [%]	3.61	4.43	5.66	5.41		
work MRE [%]	5.90	6.98	8.51	8.47		
training time [s]	1.9	1.92	0.27	0.27		
inference time [ms]	1.13	1.27	0.31	0.36		

It is interesting to see how the errors are distributed in the space of the predictions (Figure 9). We observe that the relative error (RE) remains comparatively small for high

load masses, while it increases for the smallest load masses. This suggests that the model is more reliable for high-performing loading actions, which is the general goal, than for low-performing ones.



Figure 9. The distribution of relative errors for the performance predictor visualized by scatter plots and moving averages (curves). The gray histograms show the data distribution.

It is also interesting to understand why the model fails occasionally. Therefore, we identified the test samples with the ten worst load mass predictions. These are displayed in Figure 10 with the local heightmaps. The common factors are (a) that the loaded mass is small in the ground truth samples while overestimated by the model and (b) that the heightmap is skewed, with most mass distributed on either the left or right side. It is understandable that the low-dimensional model, with only four parameters characterizing the pile state, has more difficulty with these piles as it cannot distinguish between uniform and irregular pile surfaces. That makes accurate load mass prediction difficult for complex pile surfaces.



Figure 10. The ten worst load mass predictions of $\psi_{\diamond}^{\text{low}}$. The listed values are the errors for the low- and high-dimensional models and the ground truth load mass. The heightmap of the initial pile state is shown with the grid cells color-coded by the decrease (gold) or increase (cyan) after loading. Some of the heightfields have been mirrored for illustrative purposes. The arrows indicate the dig directions.

7.2. Pile State Predictor Model

The developed pile state predictor model is evaluated by the mean absolute error (MAE) and mean relative error (MRE) of a prediction \hat{h}' compared to the simulated ground truth heightmap h'. MAE and MRE are calculated in terms of the volume difference of the enclosing surfaces. In detail, the MRE is computed as

$$MRE = \frac{1}{N} \sum_{n}^{N} \varepsilon_{n}^{rel} = \frac{1}{N} \sum_{n}^{N} \frac{\sum_{i,j} |\hat{h}_{ij}^{m} - h_{ij}^{\prime n}| \Delta l^{2}}{\sum_{i,j} |h_{ji}^{\prime n}| \Delta l^{2}},$$
(15)

where the volume of each cell, indexed by *ij*, is calculated between the surface and the extended ground plane in the grid. The MAE is calculated the same way, excluding the volume normalization. The results are summarized in Table 2. The MAE of misplaced volume constitutes roughly 25% of the bucket capacity but only 3% of the volume under the local heightmap. The post-processing step (Section 4.5.1) reduces the MAE from 0.84 to 0.75 m³, with insignificant overhead in inference time. The memory footprint of the pile state predictor model (Section 4.5) was 10.7 MB, with 6.1 MB and 4.6 MB for the VAE and MLP parts, respectively.

р-р	MAE [m ³]	MRE [%]	Inference Time [ms]			
on	0.75	3.04	4.51			
off	0.84	3.41	4.48			

Table 2. Test results of the pile state predictor model, with and without the post-processing (p-p).

The improvement by the post-processing can also be seen from the distribution of the errors in Figure 11.

Figure 12 visualizes five selected test predictions, marked by diamonds in Figure 11b. Overall, the predicted pile surfaces look reasonable compared to the ground truth. We find that the model performance tends to be better when the initial pile has a smooth, regular shape, which is to be expected as the VAE decoder tends to produce smooth outputs. Examples of the smoothness issue can be seen in the reconstruction column in Figure 12.



Figure 11. The distribution of the prediction test error (**a**) and AE correlation plot with and without post-processing (p-p) (**b**). The five points of interest, marked by diamonds (\diamond) in (**b**), are visualized in Figure 12.



(c) test data #388, AE (p-p on):0.68, AE (p-p off):0.88.

Figure 12. Cont.



(e) test data #1034, AE (p-p on):0.98, AE (p-p off):1.45.

Figure 12. Five selected samples of initial pile state *h* and the corresponding resulting ground truth pile state *h'* and model predictions \hat{h}' , with and without post-processing. The VAE reconstruction capability is included in the second column. The color codes the height. Samples (**a**) have small errors in the test data set, (**b**) have large errors, and (**c**) are close to the median. In samples (**d**,**e**), the post-processing increases and decreases the errors, respectively. The surface under the pile shows the absolute difference compared to the ground truth final pile state, except that only in the second column is it compared to the initial pile state. Here, color codes the difference in height.

Post-processing is most beneficial in cases where the initial state is irregular and when there is little change at the boundary (Figure 12e). When the initial pile is smooth with a uniform slope close to the angle of repose, loading will induce avalanches that affect the pile state at the rear boundary. This explains why the AE is sometimes increased by post-processing (Figure 12d).

7.3. Sequential Loading Predictions

This section compares the evolution of 40 sequential loading cycles using the high and low-dimensional world models to the simulation ground truth. The simulation and the predictor models start with the same initial pile state, H_1 , and use identical sequences of the dig location, heading, and action parameters, $\{x_n, t_n, a_n\}_{n=1}^{40}$. The initial pile is the one shown in Figure 1. The selected headings and dig locations are shown in Figure 13. The actions were picked randomly from the training data set. The simulated sequential loading can be seen in Supplementary Video S2. The performance predictor $\psi_{\diamond}^{\text{high}}$ is used together with the pile state predictor ϕ for high-dimensional model predictions. The lowdimensional performance predictor $\psi_{\diamond}^{\text{low}}$ is combined with the cellular automata model to predict the next pile state.

The evolution of the pile shape is shown in Figure 13 and in Supplementary Video S2. The accumulated load mass, time, work, and residual pile volume, evolving with the number of loadings *n*, are summarized in Table 3. The error accumulation over time is shown in Figure 14. While the pile state prediction error for the low-dimensional model grows at a nearly constant rate, the growth rate of the high-dimensional model is more irregular, starting at a lower rate but growing more rapidly after 15 sequential loadings. The error in loaded mass grows nearly linearly and roughly at the same rate for both models. Also, the errors in loading time and work grow mainly linearly. Both models follow the same trend.

The high-dimensional model uses Φ for the local pile state prediction, making no change outside the local heightmap, as discussed in Section 4.5. The accumulated error eventually leads to a predicted pile state with a slope steeper than the angle of repose (Figure 13).



Figure 13. Pile state evolution after every five sequential loadings for (**a**) ground truth simulation, (**b**) high-dimensional model, and (**c**) low-dimensional model combined with cellular automata. Identical dig poses (indicated by arrows) and action parameters are used in the three cases. The color codes the height.

	$\widehat{\mathcal{P}}_{1:N}$								\widehat{H}'_N			
	Load Mass [tonne]			Loading Time [s]		Work [MJ]			Pile Volume [m ³]			
N	GT	Ψ^{high}	Ψ ^{low}	GT	Ψ^{high}	Ψ ^{low}	GT	Ψ^{high}	Ψ ^{low}	GT	Φ	cell.aut.
5	14.7	15.7	16.1	55.8	53.2	54.5	2.0	2.0	2.0	768.6	765.9	767.7
10	28.2	31.5	31.3	106.1	101.0	101.9	4.1	3.9	3.8	760.7	755.9	758.8
15	40.8	45.3	47.0	161.1	155.4	154.5	6.5	6.1	6.0	753.3	745.6	749.7
20	56.8	62.0	64.4	218.0	210.8	208.1	8.9	8.5	8.3	744.0	728.7	739.5
25	71.0	75.4	80.6	273.7	265.7	260.2	11.1	10.6	10.4	735.7	714.7	730.1
30	86.1	90.5	96.6	327.3	318.7	308.9	13.2	12.8	12.4	726.9	700.1	720.7
35	101.4	106.3	114.3	389.4	385.1	369.9	15.7	15.3	14.6	718.0	679.1	710.4
40	116.0	118.6	130.9	443.0	441.9	418.8	17.9	17.4	16.6	709.5	664.2	700.6

Table 3. The accumulated load mass, time, work, and residual pile volume during sequential loading. The number of loading cycles is denoted by *n*, and GT stands for ground truth.

The high- and low-dimensional models' accumulated inference time was also measured over the 40 loading cycles. The high-dimensional model took 0.3 s in total. The lowdimensional model took 0.01 s. When the cellular automata are included, this amounts to 44.3 s, but it should be noted that the implementation of the cellular automata was not optimized or adapted for GPU computing.

To summarize, the high-dimensional model was better for shorter time horizons, but the low-dimensional model was more stable over longer horizons. The stability might be the same if the high-dimensional model used the cellular automata in place of the pile state predictor, but this would come with an additional computational cost.



Figure 14. The evolution of the prediction errors during sequential loading.

7.4. Diggability Map

The predictor model can be used to search for the most favorable dig location around the pile. We demonstrate this by creating *diggability maps*, where the model has been inferred around the entire edge of a pile. Figure 15 shows a map created by using $\Psi_{\diamond}^{\text{high}}$ to predict the performance using a fixed action a = [0.68, 4.51, 0.17, 4.46] at 150 dig locations along the local normal direction. The map highlights the regions with higher pile angles where the loaded volume would be relatively high and would be dug with lower energy



cost. Low-performance regions are visible where the pile bulges. The diggability maps can be further improved by searching also for the locally optimal heading and loading action.

Figure 15. Diggability maps that show the predicted loaded mass, time, and work at 150 locations with normal headings, indicated by the lines around the edge. The color coding is explained by the histograms for each of the predicted quantities.

8. Discussion

8.1. Feasibility

The relatively small errors of single loading prediction, 3–8.5%, are encouraging. There are, however, several practical limitations with the the method in its current form. The model is particular to the simulator's specific wheel loader and soil, which was homogeneous gravel. Consequently, a new model needs to be trained for each combination of wheel loader model and soil type unless this is made part of the model input. There is a greater need for automation solutions for more complex soils, such as blasted rock or highly cohesive media, than for homogeneous gravel. Although these materials can be numerically simulated, the problem of learning predictor models may be harder and require different data, as discussed for the loading process in [7]. For blasted rock, the local rock fragmentation probably needs to be included in the model input. On the other hand, the level of accuracy needed ultimately depends on the specific application and the prediction horizon required.

It is an attractive option to train on actual field data (rather than simulated data) to eliminate any simulation bias. That would entail collecting data from several thousand loadings with variations in pile shape and action parameters. This would be demanding but not unrealistic given the number of vehicles (of the same model) that are in operation worldwide. If all were equipped with the same bucket-filling controller, it would ultimately be a question of instrumenting these machines and sites to scan the pile shape before and after loading and tracking the vehicle's motion and force measurements.

Long-horizon planning of sequential loading requires long rollouts of repeated model inference. The error accumulation in the pile state predictor model (Figure 14) might then become an obstacle. Although the prediction accuracy at each loading cycle is improved by the post-processing (Figure 11), it does not involve avalanching and the pile may therefore evolve steep slopes during long rollouts (Figure 12). If that is the case, the pile state predictor can be replaced with mass removal and cellular automata, as described in Section 4.6, making the cellular automata the computational bottleneck. In our implementation, the cel-

lular automata were about 200 times slower than model inference, about 1 s versus 5 ms per loading. An optimized implementation can be an order of magnitude faster at least. An alternative is to use the cellular automata as a post-processor, forcing the predicted pile state to be consistent with the soil's angle of repose. If the post-processing is only occasionally needed, instead of after each pile state prediction, the computational overhead may be marginal.

8.2. Applications

We envision that the prediction models can be used in several different ways. They can be used to select the next loading action of an autonomous wheel loader or to plan the coordinated movement of both wheel loaders and haul trucks in a way that is optimal for multiple loading and hauling vehicles with the multi-objective goal of executing individual tasks efficiently while not obstructing the work of the other machines. This assumes adding a model for the V-cycle maneuver and emptying of the bucket into the load receiver. If the problem of optimal sequential planning is computationally intractable, the prediction models can be useful in developing good planning policies, for instance, using model-based reinforcement learning. Optimal dig location for repeated loading considering the characteristics of the current pile shape has been explored previously [32,33]. In contrast, our model supports optimization with respect also to future pile states. That means that it can be used to realize a specific pile shape that, for some reason, is beneficial for work site productivity and safety. If it is desirable to embed the model on a vehicle, it is worth noting that no extraordinary hardware is needed. The memory requirement, for instance, is less than 100 MB.

8.3. Implications and Future Research

This study demonstrates the feasibility of learning models for predicting the outcome of loading actions, including the future states of the pile. When these models are sufficiently precise and robust, it becomes possible to optimize wheel loading operations with respect to time and fuel consumption over long time horizons. Such a system can potentially achieve an efficiency that exceeds that of skilled human operators. This would unlock new opportunities for reducing fuel and emissions in mining and construction. This implies several research questions that first need to be addressed. Firstly, the stability of the model during long rollouts needs to be further improved. Secondly, models for more complex materials must be trained, e.g., including cohesive and fragmented rock. Finally, methods for solving the long-horizon optimization problem efficiently need to be researched.

9. Conclusions

This paper shows the feasibility of learning wheel loader world models that predict the outcome of single loading cycles given the local shape of the pile and the choice of control parameters for automatic bucket-filling. The proposed models can be used for automatic planning and control to maximize the net performance of a sequence of loading cycles predicted through repeated model inference. Topics left to explore in future work include handling more complex materials and how the proposed methods can be used for optimal planning. The latter would also provide better insight into model accuracy and inference speed requirements.

Supplementary Materials: The following supplemental videos are available at: https://www. mdpi.com/article/10.3390/automation5030016/s1, Video S1: The simulation of sequential loading with random action parameters at an initial pile state; Video S2: Comparison of sequential loading prediction between high- and low-dimensional models.

Author Contributions: Conceptualization and methodology, K.A., A.F., E.W. and M.S.; software and validation, K.A. and A.F.; formal analysis, K.A.; investigation, K.A. and M.S.; resources, K.A.; data curation, K.A.; writing—original draft preparation, K.A.; writing—review and editing, K.A., A.F., E.W. and M.S.; visualization, K.A., A.F. and M.S.; supervision and project administration, E.W. and

M.S.; funding acquisition, K.A. and M.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research was in part funded by Komatsu Ltd.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data presented in this study are available on reasonable request from the corresponding author.

Acknowledgments: This research was supported in part by Komatsu Ltd., Algoryx Simulation AB, the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation, and Swedish National Infrastructure for Computing at High-Performance Computing Center North (HPC2N).

Conflicts of Interest: Martin Servin is employed at Algoryx Simulation AB. Koji Aoshima is employed at Komatsu Ltd. The authors declare that this study was in part funded by Komatsu Ltd. The funder was not involved in the study design; the collection, analysis, and interpretation of data; the writing of this article; or the decision to submit it for publication.

Appendix A. Overview of the Dataset

First, the general characteristics of the collected dataset were investigated. As expected, we observed a wide spread in bucket-tip trajectories because of variability in loading action parameters and pile states (Figure A1). The loading performance, shown in Figure A2, is also well distributed in the intervals of 1–4.5 tons, 7–25 s, and 200–1100 kJ. As expected from a previous study [22], productivity (loaded mass per time unit) is positively correlated with the slope angle and negatively with the incidence angle.



Figure A1. All bucket-tip trajectories in the collected dataset, randomly colored.



Figure A2. The distribution of the collected performance measurements and how they are correlated with the characteristic pile slope θ , incidence angle α , and curvature κ_x and κ_z . The line is the performance moving average of each characteristics variable.

Appendix B. Model Performance Dependency on Hyperparameters

We trained and evaluated the model performance running on a desktop computer with Intel(R) Core(TM) i7-8700K, 3.70 GHz, 32 GB RAM on a Windows 64 bit system and NVIDIA GeForce RTX 2070 SUPER. The effect on the model performance by changing the amount of training data and the number of model parameters in the MLP is shown in Figure A3. As can be expected, the MRE decreases with increasing size of the dataset and the model, but the accuracy eventually levels out. The best high-dimensional model was achieved when using the full training dataset (9646 samples). The precise MLP architecture was less important.



Figure A3. The trend of generalization error (MRE), inference time, and training time (time/epoch) with the number of training/validation data (#training data) and parameters (#params). The parameters are changed by the number of units/layer and the number of hidden layers (H). The activation function used is either Leaky ReLU (L) or Swish (S). Note that the number of units is fixed at 512 for Ψ^{low} and 2048 for Ψ^{high} in the figures of #training data vs. MRE and time/epoch. The number of training/validation data is fixed at 3000 for Ψ^{low} and 9646 for Ψ^{high} in the other figures. The filled symbols \diamond and \star identify the models of Swish and Leaky ReLU as the extreme examples of higher inference speed Ψ^{low} and higher accuracy model Ψ^{high} . Table 1 shows these specific results.

Appendix C. Model Differentiability

The choice of activation function has a marginal effect on the model's accuracy, but the quality of gradients also needs to be investigated. To this end, we parameterized a line through the action space using $a = a_0 + (a_1 - a_0)s$, with $a_0 = [0.7, 0.0, 0.0, 5.0]^T$, $a_1 = [0.0, 5.0, 0.7, 0.0]^T$, and $s \in [0, 1]$. The gradients were calculated by algorithmic differentiation using Pytorch Autograd, with respect to the actions along a(s) using the chain rule

$$\frac{\partial \mathcal{P}}{\partial s} = \frac{\partial \mathcal{P}}{\partial a} \frac{\partial a}{\partial s}.$$

Figure A4 shows the result for the selected models. The Swish model produces smoother derivatives than the Leaky ReLU model, which can be expected since Leaky ReLU is not everywhere differentiable, although both appear useful for root-seeking.



Figure A4. The function values (**left**) and the Autograd gradients (**right**) of Leaky ReLU and Swish models with respect to the action *a* along a direction in the action space. The prediction values are normalized.

References

- Dadhich, S.; Sandin, F.; Bodin, U.; Andersson, U.; Martinsson, T. Adaptation of a wheel loader automatic bucket filling neural network using reinforcement learning. In Proceedings of the 2020 International Joint Conference on Neural Networks (IJCNN), Glasgow, UK, 19–24 July 2020; pp. 1–9. [CrossRef]
- Azulay, O.; Shapiro, A. Wheel Loader Scooping Controller Using Deep Reinforcement Learning. *IEEE Access* 2021, 9, 24145–24154. [CrossRef]
- Fernando, H.; Marshall, J. What lies beneath: Material classification for autonomous excavators using proprioceptive force sensing and machine learning. *Autom. Constr.* 2020, 119, 103374. [CrossRef]
- 4. Backman, S.; Lindmark, D.; Bodin, K.; Servin, M.; Mörk, J.; Löfgren, H. Continuous Control of an Underground Loader Using Deep Reinforcement Learning. *Machines* **2021**, *9*, 216. [CrossRef]
- 5. Eriksson, D.; Ghabcheloo, R. Comparison of machine learning methods for automatic bucket filling: An imitation learning approach. *Autom. Constr.* **2023**, *150*, 104843. [CrossRef]
- Halbach, E.; Kämäräinen, J.; Ghabcheloo, R. Neural Network Pile Loading Controller Trained by Demonstration. In Proceedings of the 2019 International Conference on Robotics and Automation (ICRA), Montreal, QC, Canada, 20–24 May 2019; pp. 980–986.
 [CrossRef]
- Borngrund, C.; Sandin, F.; Bodin, U. Deep-learning-based vision for earth-moving automation. *Autom. Constr.* 2022, 133, 104013. [CrossRef]
- 8. Singh, S.; Simmons, R. Task Planning For Robotic Excavation. In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, Raleigh, NC, USA, 7–10 July 1992; Volume 2, pp. 1284–1291. [CrossRef]
- Hemami, A.; Hassani, F. An Overview of Autonomous Loading of Bulk Material. In Proceedings of the 26th ISARC, Austin, TX, USA, 24–27 June 2009; pp. 405–411. [CrossRef]
- 10. Filla, R.; Frank, B. Towards finding the optimal bucket filling strategy through simulation. In Proceedings of the 15th Scandinavian International Conference on Fluid Power, Linköping, Sweden, 7–9 June 2017; Number 144, pp. 402–417. [CrossRef]
- Aoshima, K.; Servin, M.; Wadbro, E. Simulation-Based Optimization of High-Performance Wheel Loading. In Proceedings of the 38th International Symposium on Automation and Robotics in Construction (ISARC), Dubai, United Arab Emirates, 2–4 November 2021. [CrossRef]
- 12. Sotiropoulos, F.E.; Asada, H.H. Dynamic Modeling of Bucket-Soil Interactions Using Koopman-DFL Lifting Linearization for Model Predictive Contouring Control of Autonomous Excavators. *IEEE Robot. Autom. Lett.* 2022, 7, 151–158. [CrossRef]
- 13. Sotiropoulos, F.E.; Asada, H.H. Autonomous Excavation of Rocks Using a Gaussian Process Model and Unscented Kalman Filter. *IEEE Robot. Autom. Lett.* **2020**, *5*, 2491–2497. [CrossRef]
- 14. Saku, Y.; Aizawa, M.; Ooi, T.; Ishigami, G. Spatio-temporal prediction of soil deformation in bucket excavation using machine learning. *Adv. Robot.* 2021, *35*, 1404–1417. [CrossRef]
- Wagner, W.J.; Driggs-Campbell, K.; Soylemezoglu, A. Model Learning and Predictive Control for Autonomous Obstacle Reduction via Bulldozing. In Proceedings of the 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Kyoto, Japan, 23–27 October 2022; pp. 6531–6538. [CrossRef]
- 16. Lindmark, D.M.; Servin, M. Computational exploration of robotic rock loading. Robot. Auton. Syst. 2018, 106, 117–129. [CrossRef]
- Wang, S.; Yu, S.; Hou, L.; Wu, B.; Wu, Y. Prediction of Bucket Fill Factor of Loader Based on Three-Dimensional Information of Material Surface. *Electronics* 2022, 11, 2841. [CrossRef]
- Dadhich, S.; Bodin, U.; Andersson, U. Key challenges in automation of earth-moving machines. *Autom. Constr.* 2016, 68, 212–222. [CrossRef]
- 19. Aoshima, K.; Servin, M. Examining the simulation-to-reality gap of a wheel loader digging in deformable terrain. *arXiv* **2023**, arXiv:2310.05765.
- 20. Servin, M.; Berglund, T.; Nystedt, S. A multiscale model of terrain dynamics for real-time earthmoving simulation. *Adv. Model. Simul. Eng. Sci.* **2021**, *8*, 11. [CrossRef]
- Dobson, A.; Marshall, J.; Larsson, J. Admittance Control for Robotic Loading: Design and Experiments with a 1-Tonne Loader and a 14-Tonne Load-Haul-Dump Machine. J. Field Robot. 2017, 34, 123–150. [CrossRef]

- 22. Singh, S.P.; Narendrula, R. Factors affecting the productivity of loaders in surface mines. *Int. J. Min. Reclam. Environ.* 2006, 20, 20–32. [CrossRef]
- 23. Singh, S.; Cannon, H. Multi-resolution planning for earthmoving. In Proceedings of the 1998 IEEE International Conference on Robotics and Automation (Cat. No.98CH36146), Leuven, Belgium, 20 May 1998; Volume 1, pp. 121–126. [CrossRef]
- Magnusson, M.; Almqvist, H. Consistent Pile-shape quantification for autonomous wheel loaders. In Proceedings of the 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems, San Francisco, CA, USA, 25–30 September 2011; pp. 4078–4083. [CrossRef]
- Kuchur, N. Landscape Generator (Github Repository). 2021. Available online: https://github.com/nikitakuchur/landscapegenerator (accessed on 1 June 2021).
- 26. Kingma, D.P.; Welling, M. An introduction to variational autoencoders. Found. Trends® Mach. Learn. 2019, 12, 307–392. [CrossRef]
- Pla-Castells, M.; García, I.; Martínez, R.J. Approximation of continuous media models for granular systems using cellular automata. In *Cellular Automata, Proceedings of the 6th International Conference on Cellular Automata for Research and Industry, Amsterdam, The Netherlands, 25–28.October 2004*; Springer: Berlin/Heidelberg, Germany, 2004; pp. 230–237. [CrossRef]
 Kompeters Ltd. W4220 7: Kompeters Ltd. Talvas, Januar, 2017.
- 28. Komatsu Ltd. WA320-7; Komatsu Ltd.: Tokyo, Japan, 2017.
- 29. Algoryx Simulations. AGX Dynamics; Algoryx Simulations: Umeå, Sweden, 2023.
- 30. Perlin, K. An Image Synthesizer. Acm Siggraph Comput. Graph. 1985, 19, 287–296. [CrossRef]
- 31. Ramachandran, P.; Zoph, B.; Le, Q.V. Searching for Activation Functions. arXiv 2017, arXiv.1710.05941. [CrossRef]
- Magnusson, M.; Kucner, T.; Lilienthal, A.J. Quantitative evaluation of coarse-to-fine loading strategies for material rehandling. In Proceedings of the 2015 IEEE International Conference on Automation Science and Engineering (CASE), Gothenburg, Sweden, 24–28 August 2015; pp. 450–455. [CrossRef]
- Chen, G.; Wang, Y.; Li, X.; Bi, Q.; Li, X. Shovel Point Optimization for Unmanned Loader Based on Pile Reconstruction. Computer-Aided Civil and Infrastructure Engineering, n/a. Available online: https://onlinelibrary.wiley.com/doi/pdf/10.1111/mice.13190 (accessed on 1 January 2022). [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.