
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Tuomainen, Neea; Blanco-Mulero, David; Kyrki, Ville

Manipulation of granular materials by learning particle interactions

Published in:
IEEE Robotics and Automation Letters

DOI:
[10.1109/LRA.2022.3158382](https://doi.org/10.1109/LRA.2022.3158382)

Published: 01/04/2022

Document Version
Publisher's PDF, also known as Version of record

Published under the following license:
CC BY

Please cite the original version:
Tuomainen, N., Blanco-Mulero, D., & Kyrki, V. (2022). Manipulation of granular materials by learning particle interactions. *IEEE Robotics and Automation Letters*, 7(2), 5663-5670. Article 2.
<https://doi.org/10.1109/LRA.2022.3158382>

Manipulation of Granular Materials by Learning Particle Interactions

Neea Tuomainen , David Blanco-Mulero , and Ville Kyrki , *Senior Member, IEEE*

Abstract—Manipulation of granular materials such as sand or rice remains an unsolved problem due to challenges such as the difficulty of defining their configuration or modeling the materials and their particles interactions. Current approaches tend to simplify the material dynamics and omit the interactions between the particles. In this letter, we propose to use a graph-based representation to model the interaction dynamics of the material and rigid bodies manipulating it. This allows the planning of manipulation trajectories to reach a desired configuration of the material. We use a graph neural network (GNN) to model the particle interactions via message-passing. To plan manipulation trajectories, we propose to minimise the Wasserstein distance between a predicted distribution of granular particles and their desired configuration. We demonstrate that the proposed method is able to pour granular materials into the desired configuration both in simulated and real scenarios.

Index Terms—Deep learning in grasping and manipulation, machine learning for robot control, manipulation planning.

I. INTRODUCTION

THE concept of granular materials encompasses materials such as ground coffee, uncooked rice, and sand to name a few. These materials are difficult to manipulate due to their particle nature [1], as the particles of the material interact with each other potentially grouping into individual shapes. Recently, several attempts have been made to solve various tasks involving granular materials ranging from transporting the material [2], shaping [3], [4], to common daily tasks such as gathering, spreading or flipping [5]. A common approach is to use visual feedback to learn to manipulate the material. However, there is a common caveat on these approaches, the interactions of the material particles are not captured. Thus, these methods are not able to plan manipulation of granular materials into a desired configuration when there are complex interactions including e.g. dropping the material and the material interacting with rigid bodies. This is because learning the interactions from visual-feedback is not feasible as the material is not fully observable.

Manuscript received October 22, 2021; accepted February 23, 2022. Date of publication March 10, 2022; date of current version March 25, 2022. This letter was recommended for publication by Associate Editor Pablo Gil and Editor Markus Vincze upon evaluation of the reviewers' comments. This work was supported by the Academy of Finland under Grants 317020 and 328399. (Neea Tuomainen and David Blanco-Mulero are co-first authors.) (Corresponding author: David Blanco-Mulero.)

The authors are with the Department of Electrical Engineering and Automation (EEA), Aalto University, 02150 Espoo, Finland (e-mail: neea.tuomainen@aalto.fi; david.blancomulero@aalto.fi; ville.kyrki@aalto.fi).

This letter has supplementary downloadable material available at <https://doi.org/10.1109/LRA.2022.3158382>, provided by the authors.

Digital Object Identifier 10.1109/LRA.2022.3158382




Fig. 1. Granular manipulation using the GNN (left) and physical system (right). The granular material is manipulated using a rigid-body (cup) following the optimal trajectory T^* to pour the material in the particles target distribution $P_{g,G}$.

Instead of using models learned from direct observations, simulations of the physical interaction can be used. However, high-fidelity simulators are considerably expensive computationally. Instead, we can learn a computational surrogate model from these simulators to speed up the manipulation planning. Recently, graph neural networks (GNNs) have been proposed for modeling interactions between particles [6]–[9]. Some of these works have explored modeling the dynamics of complex materials such as granular [10] or deformable ones [11], [12]. The GNN based approaches have shown good accuracy and generalisation capabilities, and thus present a promising surrogate model for granular manipulation planning.

In this work, we address the problem of manipulating granular materials into a desired configuration by interacting with a rigid-object (see Fig. 1). In order to capture the interactions of the granular material with the rigid-body, we take on a model-based approach and learn a surrogate model from simulated data based on the idea of Graph Networks-based Simulators (GNS) [10]. We propose to solve the problem of manipulating the material into a desired configuration as a trajectory planning problem, where we use the GNS to rollout the dynamics of the system. The trajectory of the rigid-body is planned by minimising the Wasserstein distance between the simulated final distribution of granular particles and its target.

We evaluate two different aspects of the method. First, we perform an ablation study of the graph attributes and GNN architecture, and provide evidence that the chosen attributes and architecture are able to capture the rigid-body and granular material interactions. Secondly, we evaluate the proposed trajectory

planning method over different target distributions of the granular particles, both using the GNN rollout and a real robot set-up (Fig. 1). As a test-case task we perform pouring of a granular material into a desired configuration. Our results demonstrate that the proposed method effectively pours the material into different desired distributions. Furthermore, we demonstrate that the optimal trajectories in the real robot approximate the expected result using the GNN dynamics model.

The main contributions of this letter are:

- 1) A GNN based approach for modeling the interaction of granular materials and rigid bodies, and evaluation of the model parameters required to effectively predict the interactions.
- 2) A solution for trajectory planning based on the Wasserstein distance between the distribution of the manipulated granular material and its target distribution.
- 3) A demonstration of the framework for pouring granular materials into a desired configuration both in simulated and real scenarios.

II. RELATED WORK

A. Granular Manipulation

The manipulation of granular materials has been studied in different tasks such as robotic excavation [13], scooping [14] and pouring [15]. Recent works have focused on learning to manipulate granular materials using feedback-based methods where data is collected in the real-world [2]–[4], [15], [16]. In [2], a ConvNet model is learned to choose the required actions for transporting beans to a goal shape. Similarly, a learning-based approach was proposed by [16] to grasp granular foods. These methods learn a surrogate model from height or density maps, which does not capture the interaction of the granular particles. Furthermore, collecting data in the real-world is prohibitive and the feedback from the real data does not capture the inherent dynamics of the material. More recently, [5] proposed to learn a Reinforcement Learning policy to manipulate in simulation amorphous materials. However, the policy takes as input the density and height map, facing the same limitation of previous approaches. In contrast, [17] proposed a framework to infer the properties of granular materials using real-world data for fine tuning a simulator to perform robotic tasks.

In the present work, we propose to solve granular manipulation directly in simulation. This alleviates the burden of gathering data in the real-world of previous approaches. In addition, the proposed GNN approach enables learning the granular material interactions, as opposed to the aforementioned surrogate models. Furthermore, our method defines the desired configuration of granular particles as an experimental distribution. This enables planning for target particle configurations potentially more complex than those given by height or density maps.

B. Liquid Manipulation

Similar approaches to those in granular manipulation have been proposed for pouring liquids [18]–[20]. These either used visual feedback to detect the liquid [18], [19] or learned to

manipulate from human demonstrations [20]. While dynamics may be learned on from real examples, using real-world data is limited by data collection, as with granular materials, which can be solved by using simulated dynamics. Some of the works in liquid manipulation have considered the dynamics of the material either using high-complexity simulations [21] or simplified dynamics models [22]–[24]. Considering the dynamics of the system allows trajectory planning in more complex manipulation tasks. However, simplified dynamics models might lead to inaccurate trajectories, whereas using high-quality simulations might be costly for complex materials.

C. Learning Particle Interactions Using GNNs

An alternative is to learn the dynamics of the system from high-fidelity simulators using data-driven models [9]–[11]. These can learn from simulation techniques and present the benefit of lower computational cost than the simulators they learn from. Particularly, some GNN approaches have shown outstanding results on learning to simulate complex materials [10], [11]. Sanchez-Gonzalez *et al.* [10] proposed the Graph Networks-based Simulators (GNS) framework, which demonstrated accurate forecasting for different materials such as sand or fluids. However, their work does not involve modeling the interactions when the materials are manipulated. Recently, the GNS framework was used to learn the dynamics of manipulated cloths [12], which suggests it is a good candidate for learning to manipulate granular materials. The closer work to our aim is DPI-Nets [9], as it combines a GNN model for learning material dynamics with model predictive control for controlling a fluid. However, they do not consider the interaction dynamics of pouring the material and the fluid does not have the mechanical properties of granular materials.

III. METHOD

In this letter, we address the problem of manipulation of granular materials that are shaped into a desired configuration by interacting with them using a rigid-object. We distinguish between two types of materials: the rigid-body that manipulates the granular material, and the granular material itself. We model both materials as particles, which has proven as a good representation for modelling the interaction between different materials [7], [9], [10]. Rigid bodies are modelled as a set of rigid-body particles $\mathbf{p}_r \in P_r$ controllable, where the true position is assumed known. Granular material particles $\mathbf{p}_g \in P_g$ are affected by gravity, their interaction with each other and with rigid-body particles, and the boundaries of the manipulation scene. The granular particles cannot be directly controlled, but we can affect their state through their interaction with controlled rigid-body particles.

A. Graph-Based Representation

Our method makes use of a graph representation where each particle \mathbf{p}_i is represented by a node in a graph G . The nodes are connected via edges if the particles they represent interact with




Fig. 2. The initial graph $G_i = (V_i, E_i)$ is created from the particles of the granular material P_g and the rigid-body P_r . The particles are manipulated to match the goal particle distribution $P_{g,G}$. The graph G_i is mapped through the GNN model encoder \mathbf{f} , processor \mathbf{g} and decoder \mathbf{h} to predict the acceleration $\ddot{\mathbf{x}}_i \in A$ of the granular material particles.

each other. We assume an interaction exists when two particles are in contact, that is, their distance is lower than a connectivity radius R (see Section IV-B).

We define a graph as $G = (V, E)$, where $\mathbf{v}_i \in V$ is a node attribute vector and $\mathbf{e}_{ij} \in E$ represents an edge attribute vector, where i is the index of the sender node and j is the index of the receiver node. Here, we represent the node set as a matrix $V \in \mathbb{R}^{N \times D}$, where the i -th column represents the node attribute for node i , N represents the number of nodes and D the node attributes. Similarly, we represent the edge set as a matrix $E \in \mathbb{R}^{M \times F}$, where M is the number of edges and F the number of edge attributes. We define the node attribute vector as $\mathbf{v}_i = [\dot{\mathbf{x}}_{t-C}, \dots, \dot{\mathbf{x}}_{t-1}, \mathbf{b}_t, m, \mathbf{c}_t]$, where $\dot{\mathbf{x}}_{t-C}, \dots, \dot{\mathbf{x}}_{t-1}$ are the C previous velocities of the particle, \mathbf{b}_t represents the relative distance to the manipulation scene boundaries, m denotes whether the particle is a rigid-body or a granular material particle, and \mathbf{c}_t represents the control input applied to the particle. The control input is defined as the current velocity of the particle $\dot{\mathbf{x}}_t$ for rigid-body particles, and as a zero vector for granular material particles. The edge attribute vector is $\mathbf{e}_{ij} = [s_{ij}, d_{ij}]$, where s_{ij} represents the relative displacement between the particles $s_{ij} = \mathbf{p}_i - \mathbf{p}_j$, and d_{ij} represents their relative distance.

B. Graph Neural Network Model

We learn the dynamics of the granular material particles and their interaction with the rigid-body using a GNN model. The GNN model is used to predict the acceleration of the granular material particles and their trajectories assuming their initial position as well as the position of the rigid-body particles are known. Our network architecture is based on the GNS framework [10]. The model is divided into three sequential parts: (a) an encoder $\mathbf{f} : G_i \mapsto G_l$, (b) a processor $\mathbf{g} : G_l \mapsto G_K$, and (c) a decoder $\mathbf{h} : G_K \mapsto A$. The encoder maps the input graph G_i into a latent graph G_l . The processor propagates the information through K latent graphs, outputting the final latent graph G_K . Finally, the decoder produces as output the acceleration of each particle represented by a node using the node attributes of the final latent graph G_K . Each sequential part is built as follows:

1) *Encoder*: The encoder takes an input graph $G_i = (V_i, E_i)$ and encodes the node attributes $V_i \in \mathbb{R}^{N \times D_i}$ and edge attributes $E_i \in \mathbb{R}^{M \times F_i}$ into a latent graph $G_l = (V_l, E_l)$, where $V_l \in \mathbb{R}^{N \times D_l}$ and $E_l \in \mathbb{R}^{M \times F_l}$. The latent node and edges are given

by

$$\begin{aligned} V_l &= \varphi_v(V_i), \\ E_l &= \varphi_e(E_i), \end{aligned} \quad (1)$$

where φ_v and φ_e are multi-layer perceptrons (MLPs).

2) *Processor*: The processor consists of K interaction networks (IN) [6] such that the output attributes of one network are summed to its input attributes to produce a new latent graph. The new latent graph is given as input to the next network to propagate information through the graph. The output of the processor is the final latent graph $G_K = (V_K, E_K)$. Each IN is defined as

$$\begin{aligned} \mathbf{e}'_{ij} &= \phi_e(\mathbf{e}_{ij}, \mathbf{v}_i, \mathbf{v}_j), & \bar{\mathbf{e}}'_j &= \sum_{\mathbf{e}_{ij} \in E_j} \mathbf{e}_{ij}, \\ \mathbf{v}'_i &= \phi_v(\mathbf{v}_i, \bar{\mathbf{e}}'_i), \end{aligned} \quad (2)$$

where $E_j = \{\mathbf{e}_{ij} \mid i = 1, \dots, N\}$ is the set of edges whose receiver node is the node with index j . The functions ϕ_e and ϕ_v are MLPs such that each IN has the same architecture but different weights.

3) *Decoder*: The decoder produces the output for each node using the nodes of the graph produced by the final IN of the processor by

$$\ddot{\mathbf{x}}_i = \varphi_a(\mathbf{v}_i), \quad \mathbf{v}_i \in V_K, \quad (3)$$

where $\ddot{\mathbf{x}}_i \in A$ is the acceleration of the particle and φ_a is an MLP.

The predicted accelerations are used to update the next state and then compute the trajectory. The next state is computed using the semi-implicit Euler integration via

$$\begin{aligned} \dot{\mathbf{x}}_{i,t} &= \dot{\mathbf{x}}_{i,t-1} + \Delta t \ddot{\mathbf{x}}_{i,t-1}, \\ \mathbf{x}_{i,t} &= \mathbf{x}_{i,t-1} + \Delta t \dot{\mathbf{x}}_{i,t}. \end{aligned} \quad (4)$$

C. Trajectory Planning for Granular Manipulation

Our objective is to control the rigid-body particles P_r by following a trajectory \mathcal{T} that leads the granular particles to settle into a desired goal position $P_{g,G}$. We consider a planning horizon H , where the granular particles start from an initial position $P_{g,0}$, located inside the rigid-body. The rigid-body can be controlled through two set of actions: rotation and translation $\mathbf{u}_t = (\mathbf{R}_t, \mathbf{x}_t)$. The optimal trajectory is defined by Q via-points

$\mathcal{T}^* = (\mathbf{u}_0, \dots, \mathbf{u}_Q)$. The trajectory is then interpolated to get the actions in the planning horizon H that minimise the distance between the end position of the particles and the desired goal position:

$$\mathbf{u}_0, \dots, \mathbf{u}_Q = \arg \min d(P_{g,G}, P_{g,H}), \quad (5)$$

where $d(\cdot, \cdot)$ is the distance between the desired and predicted distributions of the particles. We choose this approach as it is unnecessary to transport each particle to the exact position of that same particle in the target distribution, as long as the distributions match. In contrast to existing approaches in particle manipulation that use the Chamfer distance, which measures the distance of every particle to its nearest neighbour [9], [25], we propose to use the Wasserstein distance, which can be understood as the minimum cost of moving the particles from one distribution to the other. We expect this to express better the semantics of material transport. Thus, the trajectory planning is formulated as a discrete optimal transport (OT) problem, where we minimise the quadratic Wasserstein distance:

$$d(P_{g,G}, P_{g,H}) = \left(\sum_{i=0}^N \|X_{i,G} - X_{i,H}\|^2 \right)^{1/2}, \quad (6)$$

where $X_{i,G} \in P_{g,G}$ and $X_{i,H} \in P_{g,H}$ are the samples of each empirical distribution. One of the issues when computing the Wasserstein distance for a large number of particles is the computational cost. In order to utilise the Wasserstein distance with large datasets we use the Sinkhorn loss $S_\varepsilon(P_{g,G}, P_{g,H})$, which is an approximation of OT and reduces the complexity as presented by [26].

To solve the optimisation problem, we adopt a population-based black-box optimiser, the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [27], and treat the trajectory planning as a constrained optimisation problem. As constraints, we use the manipulator end-effector limits and the velocity limits used for training the model, similar to [28]. The constraints guide the optimisation to consider trajectories within the training data of the model. The optimisation problem is then

$$\begin{aligned} \min_{\mathbf{u}_0, \dots, \mathbf{u}_Q} \quad & J(P_{g,G}, P_{g,H}, \mathcal{T}) \\ \text{subject to} \quad & |\mathbf{u}_q - \mathbf{u}_{q-1}| \leq \Delta \mathbf{u}_{\max}, \\ & \mathbf{u}_{\min} \leq \mathbf{u}_q \leq \mathbf{u}_{\max}, \end{aligned} \quad (7)$$

for all $q = 0, \dots, Q$, where $J(\cdot)$ is the cost function and \mathbf{u}_{\min} , \mathbf{u}_{\max} and $\Delta \mathbf{u}_{\max}$ are the boundaries and velocity limits respectively. The cost function, in addition to minimising the distance between experimental distributions, includes a term that penalises the acceleration of the rigid-body to reduces abrupt motions, such that

$$\begin{aligned} J(P_{g,G}, P_{g,H}, \mathcal{T}) = & \alpha S_\varepsilon(P_{g,G}, P_{g,H}) \\ & + \beta \sum_{t=2}^H \|\mathbf{u}_t - 2\mathbf{u}_{t-1} + \mathbf{u}_{t-2}\|^2 \end{aligned} \quad (8)$$

where α and β are regularisation constants.

Algorithm 1: Trajectory Planning.

Input: $P_{g,0}, P_{g,G}, \mathcal{T}$

Output: \mathcal{T}^*

```

1: for  $i = 0 \rightarrow T$  do
2:   for  $t = 0 \rightarrow H$  do
3:      $\mathbf{u}_t^* \leftarrow \mathcal{T}_t$ 
4:      $G_{i,t} \leftarrow \text{create\_graph}(P_{g,t}, \mathbf{u}_t^*) \triangleright \text{Section III-A}$ 
5:      $P_{g,t+1} \leftarrow \text{GNN}(G_{i,t})$ 
6:   end for
7:    $\mathcal{T} = \arg \min J(P_{g,G}, P_{g,H}, \mathcal{T}) \triangleright \text{Using (7)}$ 
8: end for
9:  $\mathcal{T}^* \leftarrow \mathcal{T}$ 

```

The optimal trajectory is found by following Algorithm 1. The algorithm takes as input an initial trajectory \mathcal{T} as well as the initial and target particle configurations $P_{g,0}$ and $P_{g,G}$. Then, the next position of the rigid-body particles $P_{r,t}$ is computed after applying the actions \mathbf{u}_t . From this, the control inputs \mathbf{c}_t are computed as the difference of the rigid-body position $P_{r,t}$ and $P_{r,t-1}$. Next, the graph $G_{i,t}$ is created as explained in Section III-A. The next position of granular particles $P_{g,t+1}$ is updated on line 5 by following the semi-implicit Euler formulation (4). Once the rollout has finished, the cost $J(\cdot)$ is computed. We perform the optimisation T times, where CMA-ES generates a new population until the maximum number of iterations is reached.

IV. EXPERIMENTAL RESULTS

In this section, we first provide information about the physics engine simulation set-up as well as the GNN training and the trajectory planning routine. Then, we perform an ablation study of the graph attributes and GNN parameters, where we evaluate the accuracy of the models predictions. After that, we assess the trajectory planning routine accuracy for pouring the granular material into different goal positions. The goals of the experiments are twofold: 1) to assess the required parameters for predicting accurately the interaction of granular materials with rigid-bodies, 2) to evaluate whether the proposed method can pour granular material into a desired shape, both in simulation and in a physical system.

A. Simulation Set-Up

We use Taichi-MPM [29] to generate training data for the GNN model. We create a simulated scene where a cup pours granular material into a container, which acts as boundaries of the scene. The container is a $10 \times 20 \times 20$ cm box with an open top allowing the cup and granular material to be located above the container. The cup is modelled as a 3D mesh of $\varnothing 7 \times 10$ cm, which has been populated and replaced by rigid-body particles P_r . The cup is controlled via translation and rotation actions. The granular material is modelled as sand particles which are initialised inside the cup. The simulation consists of $H = 300$ time-steps, where the number of particles is 1945, out of which




Fig. 3. Boxplot of ablation study using forecast Wasserstein distance. Boxes represent values between 1st and 3rd quartiles, orange horizontal lines represent medians and dotted green lines represent means. Models marked with darker blue represent the GNN model with the proposed graph attributes and architecture.

70% represent the granular material and the rest the rigid-body particles.

B. GNN Training Details

The training dataset consists of 20 simulations where the cup is translated along the Y -axis and rotated along the X -axis. The translation and rotation axes are constrained so that the granular material is always poured inside the container. We provide details about the different trajectories used for generating the training data in the Appendix. In addition, we have included a simulation where the cup only translates and rotates with small Gaussian noise and a simulation of the granular material falling without a cup in the scene.

The connectivity radius used to construct the graph is $R = 0.015$. Before constructing the graph we add a random-walk noise with standard deviation $\sigma = 3 \cdot 10^{-4}$ to the particle positions to improve the acceleration prediction as shown by [10]. The velocities and accelerations are normalised with mean and standard deviation of the dataset. The distance to boundaries \mathbf{b} in node attributes and displacement \mathbf{s}_{ij} in edge attributes are normalised using the connectivity radius R , and \mathbf{b} is clipped to range $[-1, 1]$ as in [10]. The model is trained for 2000 epochs using the Adam optimiser [30]. For the first 500 epochs the learning rate is constant $lr = 10^{-4}$, after which we switch to an exponential learning rate with $\gamma = 0.997$. The training is done using minibatches of size 2. For each MLP in our model, we use two hidden layers where each hidden layer has a size of 128. Each MLP in the encoder and processor are followed by a Layer Normalization [31]. We have experimented using two different loss functions, the number of message-passing steps, including and excluding control inputs \mathbf{c}_t , as well as the length of velocity history $\{\dot{\mathbf{x}}_{t-C}, \dots, \dot{\mathbf{x}}_{t-1}\}$ in the node attributes. The first loss function evaluated, \mathcal{L}_g , is the L1-loss that considers only accelerations of granular material particles. The second one, \mathcal{L}_{g+r} , is the L1-loss of both granular material and rigid-body particle acceleration. Both losses are averaged over all particles in the minibatch. All GNN models were trained until convergence, which takes 12 days for the model with parameters $K = 10$ and $C = 5$ on a computer with a NVIDIA V100 GPU.

C. Trajectory Planning Details

We use four target distributions $P_{g,G}$ of the granular material for evaluating the optimal trajectory determined using Algorithm 1. The target distributions are generated from simulations

that are not included in the training data of the GNN model. The trajectory \mathcal{T} has $Q = 6$ via-points and it is interpolated using a piecewise cubic hermite interpolating polynomial to the rollout horizon length used in the GNN training and test models. For the CMA-ES optimiser we used an initial variance of $\sigma_{\mathcal{T}} = 1.5$ and a population size of 20. In order to speed up the optimisation process, we scaled the trajectory rotation by π and the translation by 0.11 so that both variables are within the same range. The algorithm is run for 150 iterations with five different random seeds, where each random seed took 10 hours on a computer with a V100. The constraints were defined as $\mathbf{u}_{min,max} = \{\pm 2.8973, \pm 0.1\}$ and $\Delta \mathbf{u}_{max} = \{2.1973, 0.02\}$, which matches the maximum values of the training data distribution. We used $\alpha = 1000$ and $\beta = 0.001$ as values for the regularisation terms of the cost function (8). The α parameter is scaled with a high value due to the order of magnitude of the Sinkhorn loss, whereas β is reduced to keep it as a second minimisation objective. The values for these parameters were determined experimentally. We used one of the sinusoidal trajectories used for training (see Appendix) as initial trajectory \mathcal{T} , where the final material distribution does not match the test cases target.

D. GNN Ablation Study

We perform an ablation study of the aforementioned parameters. To evaluate the accuracy of the predictions, we use the Wasserstein distance between the predicted granular material particles and the ground-truth simulation. The initial Wasserstein distance for simulations when the particles are relatively still in the cup and the forecast error has not yet accumulated is in the order of 10^{-5} . In our experimental set-up we can consider that a Wasserstein distance of 10^{-2} for the entire rollout presents distant distributions, whereas a value in the order of 10^{-4} accurately matches the rollout particle distributions. The tests are performed by starting from the same initial state onward to create a rollout prediction, where the rollout length is the same as the simulation. The test-set includes 8 simulations that are in the model training domain. The results are summarised in Fig. 3, where the box plot excludes minimum and maximum values, since they are similar in each tested model. In Fig. 3 the best performing model is highlighted in darker blue, which was used to rollout the dynamics in the trajectory planning.

1) *Training loss*: We experimented on two loss functions \mathcal{L}_g and \mathcal{L}_{g+r} . As shown in Fig. 3, the loss \mathcal{L}_g presents a slightly

lower median and mean. The most noticeable difference between the two models is their quartiles, which are lower for the loss \mathcal{L}_g . This provides evidence that including the acceleration of rigid-body particles in the loss is unnecessary as it affects adversely when learning granular particle behaviour.

2) *Control inputs*: We investigated whether including control inputs c_t as a node attribute improves the model accuracy. Considering the quartiles of the models in Fig. 3 we can observe lower loss for the model that includes control inputs. This is to be expected since it gives additional information about the movement of the rigid-body, which improves the predictions of the granular material particles.

3) *Message-passing steps*: We also experimented on the number of INs in the processor. [10] suggests that using a higher number of message-passing steps provides better results. We selected $K = 10$ as maximum and evaluated how models with smaller values perform. The range of values selected was $K \in \{1, 2, 4, 5, 6, 8, 10\}$. We can notice in Fig. 3 that the highest K value indeed does perform best. The message-passing steps affect on how far the particle information reaches in the graph and, thus, how many interactions affect the particles. With $K = 4$ the performance is not significantly worse than $K = 10$. This suggests that $K = 4$ message-passing steps is enough to capture the system behaviour. However, the third quartile of $K = 6$ and $K = 8$ are higher than for $K = 4$, which suggests otherwise.

4) *Node velocity history length*: Finally, we assessed how the number of previous velocities included in the node attribute affects the performance. The values we tested were $C \in \{1, 2, 3, 4, 5\}$. One would hypothesise that $C = 3$ would be sufficient to capture the dynamics, as granular materials can be approximated via a second-order system. However, the performance of $C = 4$ and $C = 5$ is significantly better than the other models. This suggests that knowing the three previous velocities is not enough to accurately predict rollout dynamics in our system.

Overall, the results show that including the control input and the \mathcal{L}_g loss function yield drastically lower values compared to the other variants, as the interquartile range improves by almost one order of magnitude. The results for the number of message-passing steps showed that with $K = 10$ the mean and interquartile range are lowest within the tested models. The node velocity history length $C = 5$ provided slightly higher mean and median than $C = 4$, but presented lowest interquartile range. For this reason we choose to use $C = 5$, $K = 10$, control inputs and \mathcal{L}_g as loss function for the GNN model in the trajectory planning. This model speeds up the computation of the rollout dynamics by 20 times the high-fidelity simulator.

E. Trajectory Planning Results

The Sinkhorn loss for each of the four test cases is shown in Table I. The initial loss $S_\varepsilon(P_{g,G}, P_{g,0})$ measures the distance between the distribution of the particles in the initial configuration of the cup and its target. We also provide the loss for the end distribution of the particles following the initial trajectory $P_{g,H}$. This is used as a reference to highlight the improvement on

TABLE I
SINKHORN LOSS BETWEEN THE TARGET SHAPE $P_{g,G}$, THE INITIAL MATERIAL DISTRIBUTION $P_{g,0}$, THE FINAL DISTRIBUTION FOLLOWING THE INITIAL TRAJECTORY $P_{g,H}$, AND THE OPTIMAL TRAJECTORY P_{g,H^*} , FOR EACH TEST SET OVER FIVE DIFFERENT RANDOM SEEDS

Test	$S_\varepsilon(P_{g,G}, P_{g,0})$	$S_\varepsilon(P_{g,G}, P_{g,H})$	$S_\varepsilon(P_{g,G}, P_{g,H^*})$
1	$3.90 \cdot 10^{-2}$	$1.43 \cdot 10^{-4}$	$(2.09 \pm 0.80) \cdot 10^{-6}$
2	$3.87 \cdot 10^{-2}$	$5.68 \cdot 10^{-3}$	$(1.16 \pm 0.36) \cdot 10^{-6}$
3	$3.93 \cdot 10^{-2}$	$2.29 \cdot 10^{-3}$	$(1.88 \pm 0.53) \cdot 10^{-4}$
4	$3.83 \cdot 10^{-2}$	$1.57 \cdot 10^{-3}$	$(2.69 \pm 0.78) \cdot 10^{-6}$




Fig. 4. **Qualitative results** for each test case showing the granular material final configuration (predicted) and goal configuration (desired) after following an optimal trajectory.

matching the target distribution using the proposed algorithm. We can notice that the order of magnitude improves by an order of two following the optimal trajectory P_{g,H^*} . Our experiments showed that two particles distributions are well matched when the Wasserstein distance is in the order of 10^{-6} , whereas distances larger than 10^{-4} are far from the target distribution. The results indicate that our method is able to consistently reduce the distance between the end and target particle distributions by an order of magnitude. The particles distribution after following the optimal trajectory as well as the desired distribution for each test case are shown in Fig. 4. The final distribution for the test cases one, two and four nicely overlays the desired configuration of the material. On the other hand, the third test case Sinkhorn loss is reduced by an order of magnitude from the initial loss (see Table I). However, the final configuration does not accurately match the target shape. We hypothesise that with a larger population of CMA-ES the Algorithm 1 would find a better solution.

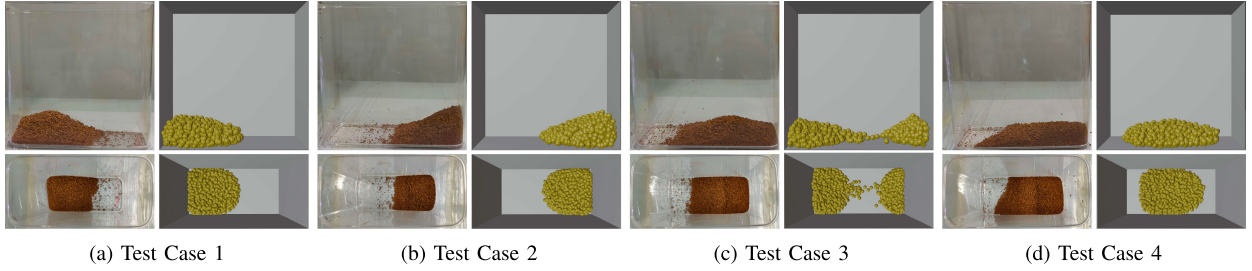


Fig. 5. **Qualitative results** of pouring ground coffee. The images show front and top views of the particles distribution after following the optimised trajectories on the real robot (**left**) and target distribution using the ground-truth simulation (**right**).

We noticed that our GNN model was unable to generalise to rotation and angular velocities greater beyond those seen during training, where the model would predict the material leaking from the rigid-body (see supplementary material). This limited the range of actions the trajectory planning algorithm could apply. Extensions of this work could focus on finding a model that generalises to any rigid-body action or training the model with a wider range of actions.

F. Granular Manipulation in Real-World

We perform experiments on a real set-up where we use a Franka Emika Panda robot to manipulate the cup. We selected ground coffee as the granular material to pour into the container. We executed the best four trajectories optimised in simulation using the real robot and qualitatively evaluated the end result of the ground coffee particles. Prior to running the trajectories, we transform the reference frame of the rigid-body into the robot end-effector reference frame. The robot is controlled via a Cartesian controller [32], where each trajectory point is interpolated from the simulation frequency (100 Hz) to the robot frequency (1KHz).

We show the qualitative results in Fig. 5. We can observe that the first and second test cases, (Fig. 5(a), (b)), match closely the desired distribution. In contrast, the fourth test case, Fig. 5(d), slightly differs from the target distribution. Furthermore, the third test case, Fig. 5(c), which required pouring in both left and right side of the container was the distribution furthest to the target. In addition, in the supplementary material we provide the real-world performance of the predefined trajectories used for generating the target distributions. These show that the distributions generated in the simulator do not exactly match the real-world, with some spilling of the material for the first test case, and a slight mismatch for the third and fourth test cases. This provides further evidence of the sim-to-real gap between the simulated and the real material.

As a summary, our results provide evidence that the proposed method is able to optimise the trajectory of the cup and pour the material in the desired configuration using the GNN rollout. In addition, the results on the real-world highlight that the proposed method is able to effectively pour the material in some of the presented test cases without gathering data from the real-world to train the model. However, some of the test cases were far from the results using the GNN rollout. This is at least partially caused by imperfect simulation. Particles sometimes leak through the

rigid body in the simulation, which is also illustrated in the supplementary material. We hypothesise that with a higher fidelity simulation that resembles better the granular material, such as the one proposed by [17], the proposed model would be able to match better the real material and pour precisely the granular material.

V. CONCLUSION

In this letter, we introduced a model-based approach for manipulating granular materials into a desired configuration by interacting with a rigid-object. The proposed method uses a GNN to learn the interactions of the granular material from simulations, as well as the interactions of the material with the rigid-body that manipulates it. We presented a trajectory planning routine that uses the GNN model to rollout the dynamics and optimises the trajectory by minimising the Wasserstein distance between the rollout result and the goal distribution of the poured material. We provided a study of the GNN model architecture as well as the graph attributes that can accurately predict the interactions between the material and the manipulated rigid-body. Our results show that the planning routine is able to find the required motions to pour the material in the desired configuration. We also demonstrated that the optimal trajectories used with the GNN rollouts are able to effectively pour the material in a real-world set-up.

One of the main limitations of using data from a simulator for planning is the simulation-to-real world gap. Naturally, using a high fidelity simulation is beneficial but the calibration of any simulation with a particular real world setting is difficult. Thus, it is an important avenue for future to study how simulation can be combined with limited data from the real world in order to achieve complex granular manipulation tasks across a variety of materials and tasks.

ACKNOWLEDGMENT

The authors would like to thank Fares J. Abu-Dakka and Gokhan Alcan of Aalto University for their help and support in this work. The authors would also like to acknowledge the computational resources provided by the Aalto Science-IT project.

APPENDIX ADDITIONAL TRAINING DETAILS

We defined four sets of trajectories equations to generate the data for training the Graph Neural Network:

- 1) The first trajectory rotates following a cosine function with zero-mean Gaussian noise. The rotation stops for several timesteps once the maximum rotation is reached and then rotates back to upright position. The translation is given by a zero-mean Gaussian noise which is added to the position at the previous timestep.
- 2) The second trajectory follows sinusoidal to translate and rotate the cup.
- 3) The third trajectory is given by a linear translation and rotation that tilt the cup a single time in both directions with constant velocity.
- 4) The last trajectory is given by a linear translation and rotation with Gaussian noise $\mathcal{N}(0, 10^{-3})$ and $\mathcal{N}(0, 1)$ respectively. The trajectory is defined as 300 time-steps where the first third is limited to translation, the second to rotation, and the last to going back to the origin position and orientation.

All the trajectories have randomised direction, maximum rotation angles and frequency for the sine and cosine trajectories, and maximum velocities for the linear trajectory. The specific details as well as the open source code can be found at.¹

REFERENCES

- [1] Z.-Y. Yin, P.-Y. Hicher, C. Dano, and Y.-F. Jin, "Modeling mechanical behavior of very coarse granular materials," *J. Eng. Mechanics*, vol. 143, no. 1, 2017, Art. no. C 4016006.
- [2] C. Schenck, J. Tompson, S. Levine, and D. Fox, "Learning robotic manipulation of granular media," in *Proc. Conf. Robot Learn.*, 2017, pp. 239–248.
- [3] A. Cherubini, J. Leitner, V. Ortenzi, and P. Corke, "Towards vision-based manipulation of plastic materials," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2018, pp. 485–490.
- [4] A. Cherubini, V. Ortenzi, A. Cosgun, R. Lee, and P. Corke, "Model-free vision-based shaping of deformable plastic materials," *Int. J. Robot. Res.*, vol. 39, no. 14, pp. 1739–1759, 2020.
- [5] Y. Zhang, W. Yu, C. K. Liu, C. Kemp, and G. Turk, "Learning to manipulate amorphous materials," *ACM Trans. Graph.*, vol. 39, no. 6, 2020, Art. no. 189.
- [6] Y. Li, Y. Liang, and A. Risteski, "Recovery guarantee of non-negative matrix factorization via alternating updates," in *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, Eds., vol. 29. Red Hook, NY, USA: Curran Associates, 2016.
- [7] D. Mrowca *et al.*, "Flexible neural representation for physics prediction," in *Proc. 32nd Int. Conf. Neural Inf. Process. Syst.*, Ser. NIPS'18. Red Hook, NY, USA: Curran Associates Inc., 2018, pp. 8813–8824.
- [8] P. W. Battaglia *et al.*, "Relational inductive biases, deep learning, and graph networks," 2018, *arXiv:1806.01261*.
- [9] Y. Li, J. Wu, R. Tedrake, J. B. Tenenbaum, and A. Torralba, "Learning particle dynamics for manipulating rigid bodies, deformable objects, and fluids," in *Proc. Int. Conf. Learn. Representations*, 2019. [Online]. Available: <https://openreview.net/forum?id=rJgbSn09Ym>
- [10] A. Sanchez-Gonzalez, J. Godwin, T. Pfaff, R. Ying, J. Leskovec, and P. Battaglia, "Learning to simulate complex physics with graph networks," in *Proc. 37th Int. Conf. Mach. Learn.*, (Proceedings of Machine Learning Research Series) H. D. III and A. Singh, Eds., vol. 119, 2020, pp. 8459–8468.
- [11] T. Pfaff, M. Fortunato, A. Sanchez-Gonzalez, and P. W. Battaglia, "Learning mesh-based simulation with graph networks," in *Proc. Int. Conf. Learn. Representations*, 2021. [Online]. Available: https://openreview.net/forum?id=roNqYL0_XP
- [12] X. Lin, Y. Wang, Z. Huang, and D. Held, "Learning visible connectivity dynamics for cloth smoothing," in *Proc. 5th Conf. Robot Learn.*, (Proceedings of Machine Learning Research Series) A. Faust, D. Hsu, and G. Neumann, Eds., vol. 164, 2022, pp. 256–266.
- [13] S. Singh, "Learning to predict resistive forces during robotic excavation," in *Proc. IEEE Int. Conf. Robot. Automat.*, vol. 2, 1995, pp. 2102–2107.
- [14] S. Sarata, H. Osumi, Y. Kawai, and F. Tomita, "Trajectory arrangement based on resistance force and shape of pile at scooping motion," in *Proc. IEEE Int. Conf. Robot. Automat.*, vol. 4, 2004, pp. 3488–3493.
- [15] S. Clarke, T. Rhodes, C. G. Atkeson, and O. Kroemer, "Learning audio feedback for estimating amount and flow of granular material," in *Proc. 2nd Conf. Robot Learn.*, vol. 87, Oct. 29–31, 2018, pp. 529–550.
- [16] K. Takahashi, W. Ko, A. Ummadisingu, and S.-I. Maeda, "Uncertainty-aware self-supervised target-mass grasping of granular foods," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2021, pp. 2620–2626.
- [17] C. Matl, Y. Narang, R. Bajcsy, F. Ramos, and D. Fox, "Inferring the material properties of granular media for robotic tasks," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2020, pp. 2770–2777.
- [18] C. Schenck and D. Fox, "Visual closed-loop control for pouring liquids," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2017, pp. 2629–2636.
- [19] C. Do and W. Burgard, "Accurate pouring with an autonomous robot using an RGB-D camera," in *Proc. Int. Conf. Intell. Auton. Syst.*, Springer, 2018, pp. 210–221.
- [20] Y. Huang, J. Wilches, and Y. Sun, "Robot gaining accurate pouring skills through self-supervised learning and generalization," *Robot. Auton. Syst.*, vol. 136, 2021, Art. no. 103692.
- [21] Z. Pan, C. Park, and D. Manocha, "Robot motion planning for pouring liquids," *Proc. Int. Conf. Automated Plan. Scheduling*, vol. 26, no. 1, 2016.
- [22] W. Aribowo, T. Yamashita, K. Terashima, and H. Kitagawa, "Input shaping control to suppress sloshing on liquid container transfer using multi-joint robot arm," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2010, pp. 3489–3494.
- [23] M. Reyhanoglu and J. H. Rubio, "Nonlinear modeling and control of slosh in liquid container transfer via a PPR robot," *Commun. Nonlinear Sci. Numer. Simul.*, vol. 18, no. 6, pp. 1481–1490, 2013.
- [24] Z. Pan and D. Manocha, "Motion planning for fluid manipulation using simplified dynamics," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2016, pp. 4224–4231.
- [25] S. Chen, X. Ma, Y. Lu, and D. Hsu, "Ab initio particle-based object manipulation," in *Proc. Robotics, Sci. Syst.*, 2021. [Online]. Available: <http://www.roboticsproceedings.org/rss17/p071.html>
- [26] J. Feydy, T. Séjourné, F.-X. Vialard, S.-I. Amari, A. Trounev, and G. Peyré, "Interpolating between optimal transport and MMD using Sinkhorn divergences," in *Proc. 22nd Int. Conf. Artif. Intell. Statist.*, 2019, pp. 2681–2690.
- [27] N. Hansen and A. Ostermeier, "Completely derandomized self-adaptation in evolution strategies," *Evol. Comput.*, vol. 9, no. 2, pp. 159–195, 2001.
- [28] F. J. Abu-Dakka, F. J. Valero, J. L. Suñer, and V. Mata, "A direct approach to solving trajectory planning problems using genetic algorithms with dynamics considerations in complex environments," *Robotica*, vol. 33, no. 3, pp. 669–683, 2015.
- [29] Y. Hu *et al.*, "A moving least squares material point method with displacement discontinuity and two-way rigid body coupling," *ACM Trans. Graph.*, vol. 37, no. 4, Jul. 2018, Art. no. 150.
- [30] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. 3rd Int. Conf. Learn. Representations*, Y. Bengio and Y. LeCun, Eds. San Diego, CA, USA, 2015. [Online]. Available: <https://openreview.net/forum?id=8gmWwjFyLj>
- [31] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," 2016, *arXiv:1607.06450*.
- [32] O. Khatib, "A unified approach for motion and force control of robot manipulators: The operational space formulation," *IEEE J. Robot. Autom.*, vol. 3, no. 1, pp. 43–53, Feb. 1987.

¹<https://sites.google.com/view/granular-gnn-manipulation>