

Differentiable physics-based system identification for robotic manipulation of elastoplastic materials

The International Journal of
Robotics Research
2025, Vol. 0(0) 1–30
© The Author(s) 2025



Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/02783649251334661
journals.sagepub.com/home/ijr



Xintong Yang¹ , Ze Ji¹ and Yu-Kun Lai²

Abstract

Robotic manipulation of volumetric elastoplastic deformable materials, from foods such as dough to construction materials like clay, is in its infancy, largely due to the difficulty of modelling and perception in a high-dimensional space. Simulating the dynamics of such materials is computationally expensive. It tends to suffer from inaccurately estimated physics parameters of the materials and the environment, impeding high-precision manipulation. Estimating such parameters from raw point clouds captured by optical cameras suffers further from heavy occlusions. To address this challenge, this work introduces a novel Differentiable Physics-based System Identification (DPSI) framework¹ that enables a robot arm to infer the physics parameters of elastoplastic materials and the environment using simple manipulation motions and incomplete 3D point clouds, aligning the simulation with the real world. Extensive experiments show that with only a single real-world interaction, the estimated parameters, Young's modulus, Poisson's ratio, yield stress and friction coefficients, can accurately simulate visually and physically realistic deformation behaviours induced by unseen and long-horizon manipulation motions. Additionally, the DPSI framework inherently provides physically intuitive interpretations for the parameters in contrast to black-box approaches such as deep neural networks.

Keywords

Elastoplastic material manipulation, system identification, differentiable physics, deformable material manipulation

Received 12 August 2024; Revised 5 February 2025; Accepted 17 February 2025

1. Introduction

Despite the recognised importance of robotic manipulation of deformable materials, this topic remains underexplored, particularly when it comes to high-precision manipulation of volumetric elastoplastic materials. A primary challenge in this area arises from the materials' infinite degrees of freedom (DoFs), leading to highly unpredictable deformation dynamics.

The intrinsic complexity of these dynamics inhibits the direct application of conventional robotic motion planning methods, which typically require explicit physics models for all concerned objects (Latombe, 2012). Learning approaches, such as reinforcement learning (RL), often involve training an agent to learn to interpret its perception and take actions through inefficient trial & error in a realistic physics-based simulation (Collins et al., 2021; Kroemer et al., 2021), which is both challenging and largely unavailable when it comes to deformable materials.

In contrast to the well-studied rigid body dynamics in robotics, where motions can be predicted and controlled

using well-defined equations of motion and deterministic models (Featherstone, 2014), deformable materials do not follow such straightforward patterns. Directly applying these methods is problematic because it is extremely difficult and often infeasible to accurately model and perceive real-world elastoplastic materials and measure the underlying physics parameters that govern their motions and deformations (Arriola-Rios et al., 2020; Yin et al., 2021). Additional challenges such as time-varying material properties (saturation, dryness, etc.) make the deformation dynamics even more unpredictable. As a result, achieving high-precision manipulation for such materials with motion planning or data-driven

¹School of Engineering, Cardiff University, Cardiff, UK

²School of Computer Science and Informatics, Cardiff University, Cardiff, UK

Corresponding author:

Ze Ji, School of Engineering, Cardiff University, Queen's Buildings, 14-17 The Parade, Cardiff CF24 3AA, UK.

Email: jiz1@cardiff.ac.uk

techniques is challenging due to the high computational cost and the lack of techniques to capture the dynamics accurately and efficiently.

To close this gap, this research proposes a Differentiable Physics-based System Identification (DPSI) framework for the robotic manipulation of volumetric elastoplastic materials. Our framework can efficiently estimate key physics parameters governing material deformation dynamics using minimal and simple manipulation motions (in less than 5 min with a good initial guess). The estimated physics parameters enable accurate material simulation for long-horizon predictions of real-world elastoplastic deformation behaviours.

The workflow of the proposed DPSI framework can be summarised as follows. As shown in Figure 1, a robot equipped with an in-hand 3D camera (Zivid) and three end-effectors is deployed to manipulate the elastoplastic object (e.g., play dough). Before manipulation, the robot takes multi-view point clouds of the object to minimise occlusions. These point clouds are used to create the initial particle system for the material point methods (MPM)-based simulation. The robot then performs a manipulation motion on the object and captures point clouds of the deformed state. Physics simulations of the same manipulation are run with the same initial state and motion, whose resultant particle states are compared to the real-world deformed state through variants of the Chamfer distance (CD) and the earth mover’s distance (EMD) loss. The parameters of the simulated physics models are updated after every simulation to minimise the loss. To facilitate fast optimisation, we use a differentiable simulator built on the TaiChi auto-differentiation mechanism, which allows automatic gradient computation from the losses and gradient-based optimisation for the physics parameters (Hu et al., 2019, 2020).

Our approach achieves unprecedented simulation-to-real alignment accuracy, characterised by the integration of the following novel features.

1.1. High-fidelity physics

Unlike previous methods that either employ non-physics-based models (e.g., neural networks) or highly simplified material geometry representations (e.g., sparse keypoints), we use high-fidelity physics-based simulation powered by the MPM (Jiang et al., 2016), which simulates materials as Lagrangian particles and keeps track of their positions and velocities. It achieves faster simulation by computing the motions, deformation gradients, and frictional contacts on a background Eulerian grid (Gao et al., 2017; Hu et al., 2018; Jiang et al., 2016; Stomakhin et al., 2013). MPM-based simulations provide highly efficient and realistic simulation with high physical plausibility by closely following real-world physics laws such as Newton’s laws and elastic and plastic energy conservation models.

1.2. Incomplete and noisy observations

Unlike existing works that rely on synthetic videos with complete sequences of perfect observations (Kaneko, 2024; Li et al., 2023; Murthy et al., 2020), our framework uses 3D point clouds to observe real-world object geometries. Capturing the full depth of an object during manipulation is impractical due to occlusions caused by the end-effector or environment. This means that only the point clouds before and after a manipulation motion are practical to obtain and informative enough to observe the full geometry of the deformed object in real-world experiments. In addition, real-world point clouds tend to suffer from inaccurately estimated camera matrices and sensory noises.

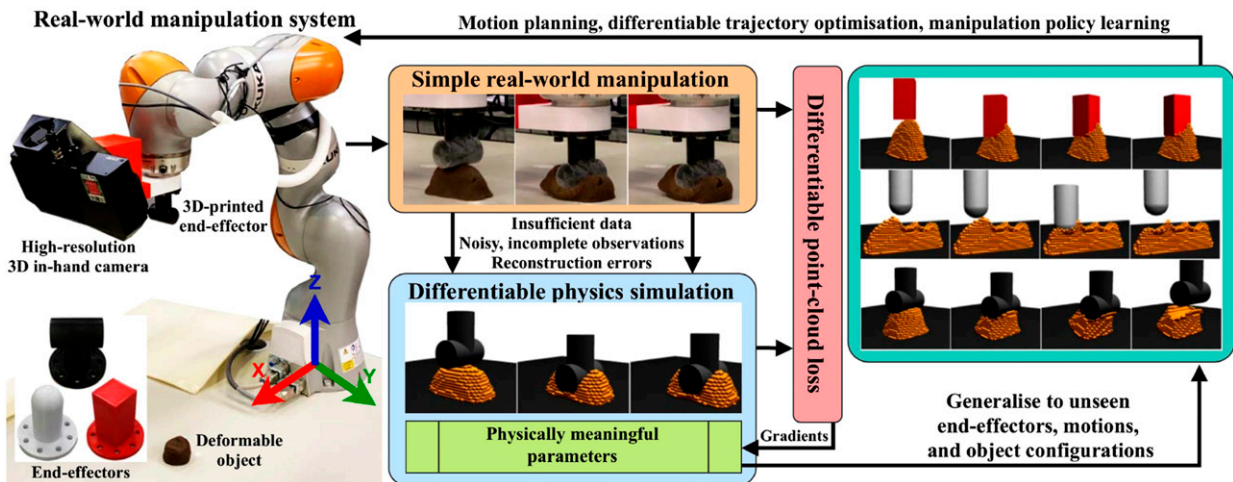


Figure 1. The proposed system identification framework enables a robot to interact with elastoplastic material via simple manipulation motions (orange box) and then identify the physics parameters of the real-world manipulation dynamics. The parameters are found using gradients computed, through differentiable simulation, from a differentiable point-cloud-based similarity function between the real and simulated observations of the manipulated material (blue and red boxes). These parameters then enable accurate simulations that allow the real-world grounding of motion planning, trajectory optimisation or policy learning techniques (cyan box).

1.3. Small data, short and simple motions

Using extensive and diverse manipulation motions to collect real-world deformation data is time-consuming and costly. Existing studies demand a significant number of real-world interactions or complete sequences of simulation videos to identify object deformations under various motions, yet still resulting in simulations with insufficient accuracy for real-world applications (Kaneko, 2024; Li et al., 2023; Lin et al., 2022; Shi et al., 2023). Our goal is to recover physics parameters that enable accurate predictions of long-horizon, unseen and complicated elastoplastic material manipulation dynamics, using minimal simple and short real-world interactions.

1.4. Joint parameter estimation

We aim to jointly estimate the physics parameters provided by physics models. Besides Newton’s laws, we employ the fixed corotated elastic energy model (Stomakhin et al., 2012), the von Mises plasticity model (Jones, 2009) and the dynamic friction model in our simulation. These lead to six key parameters: Young’s modulus E , Poisson’s ratio ν , yield stress σ_y , material density ρ and the friction coefficients of the table η_t and end-effectors η_m (assuming the three end-effectors share the same coefficient). The first four parameters primarily govern the deformation responses, while the last two handle frictional contacts. These parameters are heavily intertwined in governing the behaviours of the manipulated object and there is no intuitive solution to identify one of them without estimating the effects of other parameters. Therefore, we seek to identify these parameters simultaneously.

1.5. Differentiable physics

Identifying the physics parameters in their discretised spaces via search or evolutionary algorithms is computationally slow due to the exponentially growing number of possible combinations as the discretisation becomes finer. While gradient-based optimisation methods offer faster convergence toward the minimum, it is infeasible with most physics simulations because many computation steps are not differentiable and these simulators do not support derivative computations. In this work, we explore the feasibility of optimising system parameters using gradients computed by differentiating loss functions through a physics simulator written by DiffTaiChi, a programming language tailored for GPU-accelerated parallel computation and automatic differentiation (Hu et al., 2019, 2020). DiffTaiChi generates derivative functions for simulation steps via source code transformation that retains arithmetic intensity and parallelism. It uses a memory-efficient tape to record the order of computation kernels for forward simulation and traverses their derivative functions in the backward order to generate gradients through the computation graph. We build DPSI upon DiffTaiChi and explore the feasibility of directly

optimising several physics parameters jointly with gradients computed by differentiating point-cloud-based loss functions through the high-fidelity physics simulator.

Substantial experiments demonstrate that our main contribution, DPSI, can achieve highly accurate simulation-to-reality alignment for elastoplastic materials manipulated by unseen, long horizon and complex motions using minimal simple and short interactions, and noisy and incomplete observations. Results show that when multiple solutions and parameter uncertainty exist, DPSI can provide physically intuitive parameter interpretations that can guide further system identification, model improvement, and motion adaptation. Statistics on the computation costs also indicate promising practical deployment of the DPSI framework.

The rest of the article reviews related literature, presents formally our method and experiment results, and discusses limitations and future directions.

2. Related work

2.1. Deformable object manipulation

Both model-free and model-based approaches have been taken for manipulating deformable materials. Existing model-free methods often lack manipulation precision due to the absence of physics laws that describe the motion and deformation under complex contacts (Cherubini et al., 2020; McConachie et al., 2020; Shen et al., 2024; Shi et al., 2023, 2024). Physics-model-based methods, while more accurate, struggle with aligning the simulated dynamics with the real world and often rely on simplified geometric representations for higher computational efficiency, sacrificing manipulation precision (Navarro-Alarcon et al., 2016; Shetab-Bushehri et al., 2023; Yang et al., 2023). With or without a model, many of them use simulations that offer a cost-effective way to test manipulation methods and allow the collection of massive data for learning-based approaches (Arriola-Rios et al., 2020; Collins et al., 2021; Yin et al., 2021). However, inaccurate dynamics predictions and control precision are unacceptable for many tasks such as surgery, assembly and disassembly. Therefore, akin to the way humans efficiently learn about object and environment physics properties, this work proposes to actively identify system parameters that enable high-precision simulation of real-world manipulation dynamics for volumetric elastoplastic objects.

2.2. Deformable object modelling

Simulators are essential for advancing robotic manipulation, providing a fast, low-cost alternative to real-world testing (Collins et al., 2021; Featherstone, 2014). In recent years, arguably the most efficient and accurate simulation for 3D deformable objects is achieved by the material point method (MPM) (Jiang et al., 2016). Like most well-known methods for simulating continuum matter, such as position-based dynamics (PBD) and smoothed particle hydrodynamics

(SPH) (Yin et al., 2021), MPM represents the object as Lagrangian particles and keeps track of their positions and velocities. Unlike pure Lagrangian methods, MPM achieves faster simulation by computing the motions, deformation gradients, and frictional contacts on a background Eulerian grid (akin to the finite element method), governed by elastic and plastic energy functions and Newton’s laws in the form of partial differential equations (Gao et al., 2017; Hu et al., 2018; Jiang et al., 2016; Stomakhin et al., 2013). MPM has been proven superior to other methods in terms of efficiency and visual effects for objects that undergo large deformations, fractures, and self-collision (Hu et al., 2018). Despite its advantages, accurately identifying physics parameters for real-world robotic manipulation of elastoplastic objects remains an open challenge.

2.3. Deformable object system identification

Unlike rigid bodies (Chen et al., 2022; Heiden et al., 2022; Jaques et al., 2022), identifying the physics properties of deformable objects is more complex than for rigid bodies, primarily due to the difficulty in measuring key parameters like material properties and friction coefficients (Arriola-Rios et al., 2020; Yin et al., 2021). Previous research has mostly focused on linear and planar deformable objects, such as ropes and cloths (Caporali et al., 2024; Sundaresan et al., 2022; Yang et al., 2022).

Early works that sought to identify 3D volumetric material properties have much simpler assumptions and scenarios. For instance, one of the earlier works uses gradient-based optimisation to retrieve the stiffness of a spring system representing elastic deformable objects (Lloyd et al., 2007), while another work uses an exhaustive

search method to find the value of Poisson’s ratio for an elastic form object (Güler et al., 2017). They focus on single parameter identification for volumetric elastic deformable objects with reduced Dofs that under-represent the geometries and deformation behaviours of real-life objects.

Recent efforts like GradSim (Murthy et al., 2020) have used differentiable physics and rendering to identify five parameters for elastic objects with much higher DoFs from single-view simulation videos, demonstrating the feasibility of differentiable system identification with synthetic videos. Two following works, PAC-NeRF (Li et al., 2023) and LPO (Kaneko, 2024) propose to jointly reconstruct object geometries (shapes, positions and colours) and physical properties by using a voxel neural radiance field (Sun et al., 2022) that performs differentiable rendering and allows gradients to be back-propagated from the image space to the Eulerian grid.

Compared to these works that focused on simulations, our study tackles the more challenging task of system identification for real-world objects through robot interactions. Similarly, we also employ the material point method (Hu et al., 2018) and DiffTaiChi (Hu et al., 2020) for differentiable physics simulation and study the system identification task without known object geometries. However, we aim to align simulations with real-world dynamics with minimal robot interactions, using only incomplete, occluded, and noisy point cloud data.

3. Method

3.1. Overview

The proposed differentiable physics-based system identification (DPSI) framework, as shown in Figure 2, can be

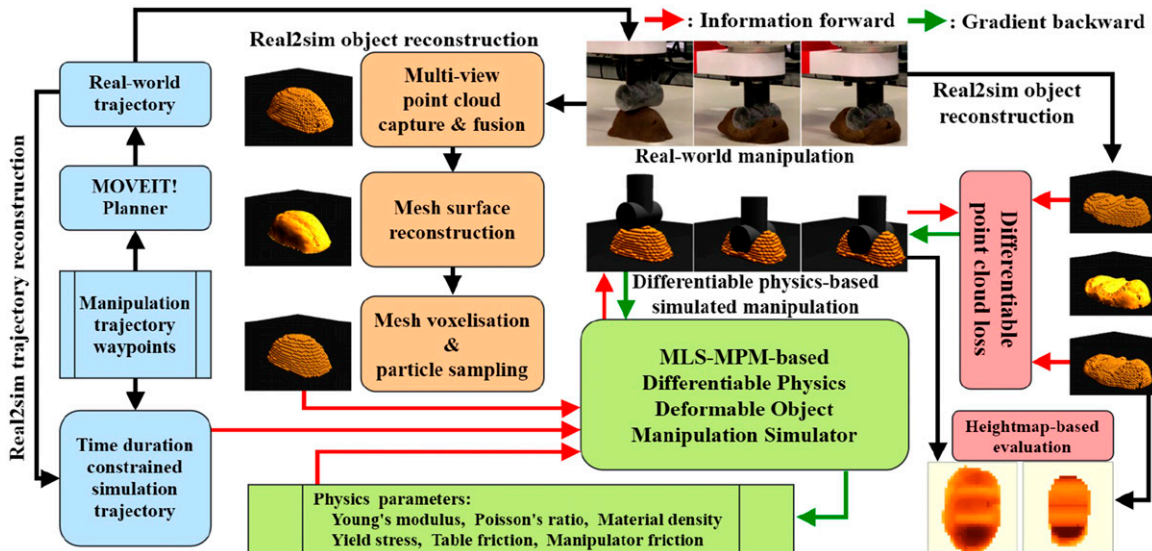


Figure 2. The overall workflow of the proposed differential physics-based system identification (DPSI) framework. Modules in each colour are elaborated in individual subsections of the Method section. Green: differentiable dynamics modelling and the physics parameters. Orange: real-to-sim object reconstruction. Blue: real-to-sim trajectory reconstruction. Pink: optimisation and evaluation loss functions.

divided into the following modules: differentiable dynamics modelling, real-to-sim object and trajectory reconstruction and the loss functions. This section starts with a formal problem description and elaborates on each of these modules.

3.2. Problem statement

We define the physics parameter identification problem for real-world deformable materials as follows. Denote \mathcal{X}_0^{real} and \mathcal{X}_T^{real} as the observations of a deformable object before and after a manipulation trajectory τ^{real} of time duration T is applied, denote \mathcal{S}_0^{real} as the observation of the real-world end-effector before the manipulation, the stochastic real-world forward dynamics can be written as follows:

$$\mathcal{X}_T^{real} \sim p^{real}(\mathcal{X}_T^{real} | \mathcal{X}_0^{real}, \theta^{real}, \tau^{real}, \mathcal{S}_0^{real})$$

where θ^{real} is the set of real-world parameters that govern the behaviours of the dynamical process. As it is unlikely to obtain p^{real} and θ^{real} explicitly, we use an approximate, deterministic and differentiable dynamics model, f^{sim} , to simulate such a dynamical process, written as follows:

$$\mathcal{X}_T^{sim} = f^{sim}(\mathcal{X}_0^{rec}, \theta^{sim}, \tau^{rec}, \mathcal{S}_0^{rec})$$

where θ^{sim} is the set of adjustable parameters that govern the behaviours of the simulated deformation and contact process, \mathcal{X}_0^{rec} is the reconstructed initial observation of the real object, \mathcal{X}_T^{sim} is the resultant observation of the reconstructed object after manipulated by the reconstructed motion τ^{rec} and \mathcal{S}_0^{rec} is the initial observation of the simulated end-effector. In our work, we assume the end-effectors are rigid bodies, meaning the observations of the real or simulated end-effectors are coordinates of their frames relative to the world frame, that is, $\mathcal{S} \in \mathbb{R}^3$, which can be readily retrieved from the robot platform or simulator. As we are interested in point cloud observations from the real world and particles in the simulation, all object observations will be 3D point sets, that is, $\mathcal{X} \in \mathbb{R}^{N \times 3}$ where N is the number of points or particles.

With these notations, the optimisation problem could be formulated as the minimisation of some distance function between the real-world and simulated observations after manipulation. As the parameters are optimised with a dataset of interaction experiences with the real objects, \mathcal{D} , the minimisation problem can be written as follows:

$$\min_{\theta^{sim}} \frac{1}{|\mathcal{D}|} \sum_{\mathcal{D}} d(\mathcal{X}_T^{real}, \mathcal{X}_T^{sim}) \quad (1)$$

The rest of this section will discuss the differentiable simulation dynamics f^{sim} , the reconstructed object state, \mathcal{X}^{rec} , the reconstructed end-effector trajectory, τ^{rec} , and the loss functions, $d(\mathcal{X}_T^{real}, \mathcal{X}_T^{sim})$.

3.3. Differentiable dynamics modelling

3.3.1. Material point method (MPM). This work employs the mean-least-square material point method (MLS-MPM) to simulate the deformable object manipulation dynamics. MPM is a meshless, hybrid computation scheme that enables efficient computation and preserves high physical fidelity for various materials, especially for elastoplastic materials that undergo large deformation, while the MLS-MPM arises from a novel weak form discretisation of the conservation equations and replaces the shape functions in the force computation with MLS approximators, leading to faster and more realistic simulation of sharp separation of particles and two-way coupling with rigid objects that traditional MPM cannot simulate.

We modified the standard procedures in one simulation step of the MLS-MPM as shown below to incorporate end-effector control following Xian et al. (2023). The readers are referred to the original paper for more details about MPM (Stomakhin et al., 2013) and MLS-MPM (Hu et al., 2018).

1. Compute particle deformation gradient using the MLS approximation equation.
2. Applied plasticity to recompute deformation gradients and particle stress using the elastic energy model
3. Particle to grid. Use the affine particle-in-cell transform (Jiang et al., 2015) to transfer the velocities and masses of the Lagrangian particles to the background Eulerian grid nodes. In our implementation, we assume equal volume and mass for the particles.
4. Update end-effector positions given control inputs.
5. Compute grid node momenta and velocities with gravity applied.
6. Signed-distance field (SDF)-based collision detection with rigid objects (end-effector and boundaries) and frictional contact computation.
7. Grid to particle. Use the affine particle-in-cell transform to transfer the velocities and affine coefficients from the grid nodes to the particles. Perform SDF-based collision detection and frictional contact computation again to minimise particle penetration.
8. Update particle positions with the new velocities.

3.3.2. Elastoplasticity. We assume strain elastoplasticity for the studied objects, meaning that the strain-stress relationship of the material is described using the deformation gradient that can be decomposed into elastic and plastic parts, $\mathbf{F} = \mathbf{F}^E \mathbf{F}^P$, and the first Piola–Kirchhoff stress, \mathbf{P} , as $\mathbf{P} = \partial \Psi / \partial \mathbf{F}^E$, where $\Psi(\mathbf{F}^E)$ is the elastic energy density function. In this study, we use the fixed corotated model (Stomakhin et al., 2012) as the energy function for its robustness under large deformation, although other constitutive models may be used for different needs. The fixed corotated elastic energy density function is

$$\Psi(\mathbf{F}^E) = \mu \|\mathbf{F}^E - \mathbf{R}^E\|_F^2 + \frac{\lambda}{2} (J^E - 1)^2 \quad (2)$$

where $J^E = \det(\mathbf{F}^E)$, and λ and μ are the Lamé parameters, such that

$$\lambda = \frac{E}{2(1+\nu)} \quad (3)$$

$$\mu = \frac{E\nu}{(1+\nu)(1-2\nu)} \quad (4)$$

where E is Young's modulus and ν is Poisson's ratio. Thus, the stress can be calculated as

$$\mathbf{P} = \frac{\partial \Psi}{\partial \mathbf{F}^E} = 2\mu(\mathbf{F}^E - \mathbf{R}^E)(\mathbf{F}^E)^T + \lambda(J^E - 1)J^E \quad (5)$$

where \mathbf{R}^E is the rotation matrix that comes from the polar decomposition of the deformation gradient, $\mathbf{F}^E = \mathbf{R}^E \mathbf{S}^E$, where \mathbf{S}^E is a symmetric matrix. In practice, they are computed from the singular value decomposition of \mathbf{F}^E as $\mathbf{F}^E = \mathbf{U}^E \Sigma^E (\mathbf{V}^E)^T = \mathbf{U}^E (\mathbf{V}^E)^T \mathbf{V}^E \Sigma^E (\mathbf{V}^E)^T$, and then $\mathbf{R}^E = \mathbf{U}^E (\mathbf{V}^E)^T$ and $\mathbf{S}^E = \mathbf{V}^E \Sigma^E (\mathbf{V}^E)^T$ (Baker, 2005).

For the n -th simulation step, we compute a trial elastic deformation gradient, $\tilde{\mathbf{F}}^{E,n}$. If no plastic deformation appears, meaning $\mathbf{F}^{P,n} = \mathbf{F}^{P,n-1}$, then we assign $\mathbf{F}^{E,n} = \tilde{\mathbf{F}}^{E,n}$. However, if the yield criterion is violated, $\tilde{\mathbf{F}}^{E,n}$ needs to be modified according to the plasticity model. This process of modifying $\tilde{\mathbf{F}}^{E,n}$ is called the return mapping process, which determines how the object would respond plastically.

In this work, we follow Gao et al. (2017) to use the von Mises model (Jones, 2009) to compute the return mapping, which takes on the associative plastic flow assumption. The projection process of the trial stress outside of the yield region can be described succinctly as $\tilde{\boldsymbol{\epsilon}}^{E,n} - \boldsymbol{\epsilon}^{E,n} = \delta_\gamma \tilde{\mathbf{s}}^{E,n} / \|\tilde{\mathbf{s}}^{E,n}\|$, where $\tilde{\boldsymbol{\epsilon}}^{E,n} = \log(\tilde{\mathbf{F}}^{E,n})$ and $\boldsymbol{\epsilon}^{E,n} = \log(\mathbf{F}^{E,n})$ are the trial and modified Hencky strains, $\tilde{\mathbf{s}}^{E,n} = \tilde{\boldsymbol{\epsilon}}^{E,n} - \text{tr}(\tilde{\boldsymbol{\epsilon}}^{E,n})/3 \mathbf{I}$ is the trial deviatoric stress, and $\delta_\gamma \tilde{\mathbf{s}}^{E,n} / \|\tilde{\mathbf{s}}^{E,n}\|$ is the modifying term based on the solution given in Gao et al. (2017). In particular, based on the von Mises yield criterion, $\sqrt{3}J_2 - \sigma_y \leq 0$, where σ_y is the yield stress and J_2 is the second deviatoric stress invariant, it was derived that $\delta_\gamma = \|\tilde{\mathbf{s}}^{E,n}\| - \sigma_y/2\mu$. With the fact that the singular vectors of the trial elastic deformation gradient do not change, the return mapping operation is done on the eigenvalues of the trial Hencky stress, thus on the eigenvalues of the trial elastic deformation gradient. Then, we can obtain the result of the return mapping by taking the exponential of the modified deviatoric stress: $\Sigma^{E,n} = e^{\boldsymbol{\epsilon}^{E,n}}$ and then the modified deformation gradient: $\mathbf{F}^{E,n} = \tilde{\mathbf{U}}^{E,n} \Sigma^E (\tilde{\mathbf{V}}^{E,n})^T$.

Note that the Hencky strain formation is only used in plastic response computation. At each simulation step, we compute the trial deformation gradient using the MLS approximation equation, apply plastic response, and then compute the new stress with equation (5).

3.3.3. Frictional contacts. Collision detection is done by checking the distance of the particles to the surface of the

rigid objects using pre-computed signed distance fields (SDFs). For each rigid object (table, and three end-effectors in our case), we employ the procedure described in Park et al. (2019) to generate SDFs for watertight meshes. We assume that the frictional contacts only happen in two cases: between the particles and the table, or between the particles and the end-effector. We also assume uniform friction coefficient distribution over the contact surface of the table and the end-effectors. When a particle or grid node and an object are in contact, we follow Stomakhin et al. (2013) to determine the velocity of the particle and grid nodes using dynamic friction with sticky impulse.

Specifically, for each particle, we calculate the local normal \vec{n} and the velocity v_{obj} of the rigid object at the particle position. Then, we project the particle velocity v into the object reference frame to have $v_{rel} = v - v_{obj}$. No collision is applied if the particle is separating from the object, that is, $v_n = v_{rel} \cdot \vec{n} \geq 0$. Define the tangential portion of the relative particle velocity as $v_t = v_{rel} - \vec{n} v_n$. For sticky impulse, we set the relative particle velocity v'_{rel} after the collision to 0 if $\|v_t\| \leq -\mu v_n$, where μ is the friction coefficient. If the sticky impulse is overcome, dynamics friction is applied, such that $v'_{rel} = v_t(1 + \mu v_n / \|v_t\|)$. After the collision, the new particle velocity in the world frame is calculated as $v' = v'_{rel} + v_{obj}$. We represent the friction coefficients for the table and the manipulators separately as μ_t and μ_m .

3.3.4. Differentiable programming. Several programming tools are available for creating differentiable simulations, such as PyTorch (Arnavaz et al., 2023; Sundaresan et al., 2022), DiffTaiChi (Hu et al., 2020), Jax (Schoenholz and Cubuk, 2020) and neural-network-based simulators (Heiden et al., 2021). We build our simulator based on DiffTaiChi due to its automated differentiation mechanism, GPU-accelerated parallel computation, fast computation kernel evaluation, intuitive Python APIs, rich community support and various promising applications (Huang et al., 2021; Lin et al., 2022; Xian et al., 2023).

In particular, for each numerical computation step, DiffTaiChi flattens its computation branches (e.g., boundary and collisions) and replaces mutable local variables with extra local storage variables, producing straight-line codes without mutable variables. It then uses standard source code transform to generate the derivative function based on the adjoint method. To compute gradients with a loss function, DiffTaiChi records the order of computation kernels and the scalar variables in the forward simulation direction, and then it computes the gradients for the concerned variables by evaluating the derivative functions in the reversed simulation direction (Hu et al., 2020). We build a simulator based on this programming language to allow automatic gradient computation for the physics parameters.

In summary, the elastic and plastic models describe the deformation behaviours of the object using Young's modulus E , Poisson's ratio ν and yield stress σ_y , the computation of the particle and grid velocities will involve another

parameter, the object density ρ , and the frictional contact processes are controlled by respective friction coefficients: table friction coefficient μ_t and manipulator friction coefficient μ_m . We optimise these six parameters to align the simulation to the real-world dynamics.

3.4. Real-to-sim object reconstruction

3.4.1. Real-world platform. To collect real-world data, we set up a deformable object manipulation system where a Kuka IIWA LBR 14 industrial arm (Kuka, 2024) is equipped with a Zivid One + medium camera (Zivid, 2024) for perceiving the object, and one of the three end-effectors, namely a rectangular cuboid, a cylinder roller and a bullet-shaped object for collecting interaction experiences with different contact geometries, as shown at the bottom-left of Figure 1. We use plasticine (non-hardening modelling clay) as the main material for our experiments and cloud slime and soil for generalisation experiments. The clay and cloud slime are purchased from Amazon.²

3.4.2. Real-world perception. We create a multi-view point-cloud capture and fusion process to obtain the real-world observation of the deformable object \mathcal{X}_0^{real} and \mathcal{X}_T^{real} . In particular, for each object state, the robot arm moves to six poses around the object for the camera to take a point cloud of the scene. The point clouds are fused and then cropped to contain only the points of the object.

These observations are noisy and incomplete in three senses. Firstly, due to camera calibration error, the fusion result always exhibits a $\sim\pm 3$ mm discrepancy. Secondly, due to joint limits, the robot arm cannot reach poses that allow the camera to capture the bottom part of the object. To simplify the problem, we assume that the angle between the object boundary and the table surface is equal to or greater than 90° , which allows the camera to capture as much as possible the bottom part of the object. We then project all points to its bottom to form a closed surface. Note that the initial configuration of the object can be manually shaped to satisfy this assumption of contact angle but the end configuration of the object after being manipulated is out of manual control, which tends to have more occlusions. Lastly, we only take the observations before and after applying the manipulation motions, without providing the intermediate observations during the manipulation, because it is impractical to do so when too much occlusion occurs during manipulation.

3.4.3. Reconstruction pipeline. To simulate the object as a set of particles, we design a pipeline to reconstruct the object particle system \mathcal{X}^{rec} from the given fusion point cloud, as shown by the orange boxes in Figure 2. Firstly, we down-sample the processed fusion point cloud and create a watertight mesh using the ball-pivoting (BP) algorithm (Bernardini et al., 1999). The artificial bottom surface points prevent the BP algorithm from creating a non-flat bottom mesh surface. To sample particles, we create a set of voxels

that fill in the reconstructed mesh surface uniformly given a resolution (1080×1080). Lastly, we fill particles into the voxels with a given density (4×10^6 particles/ m^3). We assume that the particles are uniformly distributed within the mesh.

For the end-effectors, we assume they are rigid bodies. In simulation, we keep track of the coordinate of its frame and a pre-computed SDF for collision detection.

3.5. Real-to-sim trajectory reconstruction

3.5.1. Manipulation motions. As we are interested in recovering real-world dynamics with small data and short, simple motions, we design experiments in two contact complexity levels, for which we create two poking motions and two poking-shifting motions for optimisation and in-distribution validation, and three longer motions for out-of-distribution validation. These motions are summarised in Table 1 in terms of consecutive trajectory segments in the world frame.

The first contact level focuses on identifying the parameters that primarily govern the deformation of the object (Young’s modulus, Poisson’s ratio, yield stress and material density) using two poking motions, which press the object down by a certain distance. The second level further includes the friction coefficients, using two poking-shifting motions, which press down the object and make horizontal shifting movements. For out-of-distribution validation, three longer motions with more drastic contact processes are created. Each motion uses a different end-effector. The triple-poking motion is designed for the round end-effector to validate the long-term deformation prediction with small fraction influences. The flattening motion is designed to validate the long-term deformation and frictional contact prediction under large linear movements, using the cylinder end-effector. The pressing-180-rotating motion is designed to validate the long-term deformation and frictional contact prediction under large rotational movements using the rectangle end-effector.

All motions start from a configuration where the end-effector tip is positioned at the top centre of the object. ROS (Quigley et al., 2009) and the MoveIt!-based (Gömer et al., 2019) OMPL planner (Sucan et al., 2012) are used to plan real-world motion trajectories τ^{real} . For calculating each motion plan, we pass a series of waypoints to the MoveIt! planner by discretising each segment of the motion with an interval of 0.002 m or 5° .

A challenging phenomenon that occurs during our data collection process is that the object tends to stick to the end-effectors after contact is made. If we allow this to happen and assume that the object always drops down eventually, the optimisation process will be extremely difficult due to a large uncertain dynamics process of the object dropping from the air. For example, the object may bounce out of the workspace or even off the table. This further exacerbates the difficulty of optimising with only the start and end observations of the object. However, it is beyond the capability of

Table 1. Waypoint designs and statistics of the interaction motions. For collecting optimisation data, there are two motions for each of the two levels of contact complexity. For out-of-distribution validation, one motion per end-effector is designed. Directions of the waypoints are relative to the robot base frame, as shown in Figure 1. Each waypoint of the real trajectories takes an uneven time interval, while each waypoint in simulation takes exactly 0.01 s.

Optimisation motions (short and simple)			Duration (sec.)	Real Num. waypoints	Sim. Num. waypoints
Lv. 1	Poking-1	$-z$ 0.015 m, $+z$ 0.03 m	0.86	38	87
	Poking-2	$-z$ 0.02 m, $+z$ 0.03 m	0.93	40	94
Lv. 2	Poking-shifting-1	$-z$ 0.02 m, $-x$ 0.03 m, $+z$ 0.03 m	1.52	68	152
	Poking-shifting-2	$-z$ 0.02 m, $+x$ 0.03 m, $+z$ 0.03 m	1.50	54	153
Validation motions (unseen, longer, more complex contacts)					
Flattening (cylinder)		$-z$ 0.025 m, $+x$ 0.025 m, $+z$ 0.025 m, $-x$ 0.025 m, $-z$ 0.025 m, $-x$ 0.025 m, $+z$ 0.025 m	3.74	275	378
Triple-poking (round)		$+y$ 0.025 m, $-z$ 0.025 m, $+z$ 0.025 m, $-y$ 0.025 m, $-z$ 0.025 m, $+z$ 0.025 m, $-y$ 0.025 m, $-z$ 0.025 m, $+z$ 0.025 m	4.55	328	460
Poking-180-rotating (rectangle)		$-z$ 0.025 m, about $+z$ 180°, $+z$ 0.025 m	6.23	504	625

current perception hardware and, thus, a future research direction. To simplify the problem, we take a simple workaround by covering the end-effector with a thin layer of flour before each manipulation motion is executed. We sink the effector into the flour and flick it gently to ensure that it is only covered by a thin layer of flour. This greatly prevents the sticking phenomenon from happening. Note that the end-effector friction coefficient we are optimising for is then the one covered by flour instead of the original value.

3.5.2. Time duration constrained reconstruction. To simulate end-effector motions, the real-world trajectory generated by MoveIt! is inconvenient as it has uneven time differences between consecutive waypoints and the simulation can only handle a constant step size dt . Thus, for a motion segment between a pair of waypoints, we reconstruct its simulation counterpart to have constant velocity by dividing the travelled distance by the real-world time duration, which is provided by the motion planner, and then we discretise the segment with a constant dt . In this study, we set $dt = 0.01$ second for better simulation stability. A larger dt will result in a too-high compounding error during simulation stepping, while a too-small dt will demand too much computation. The statistics of the real and reconstructed motions are summarised in Table 1. It can be seen that the validation motions are much longer than the optimisation ones. The MoveIt! trajectories and the reconstructed ones are saved as .npy files and will be open to the public upon acceptance.

3.6. Loss functions

We use four loss functions to calculate the difference between the simulated and real object states. As we are dealing

with points and particles, it is natural to select point-based metrics. Therefore, we employ the two most common distance metrics for point sets, namely, the Chamfer distance (CD) and the earth mover’s distance (EMD). Given two point sets \mathcal{X}_0 and \mathcal{X}_1 , the CD and EMD are defined as follows.

$$d_{CD}(\mathcal{X}_0, \mathcal{X}_1) = \frac{1}{|\mathcal{X}_0|} \sum_{x \in \mathcal{X}_0} \min_{y \in \mathcal{X}_1} \|x - y\|_2 +$$

$$\frac{1}{|\mathcal{X}_1|} \sum_{y \in \mathcal{X}_1} \min_{x \in \mathcal{X}_0} \|x - y\|_2$$

$$d_{EMD}(\mathcal{X}_0, \mathcal{X}_1) = \frac{1}{|\mathcal{X}_0|} \sum_{x \in \mathcal{X}_0} \min_{\phi \in \phi^{EMD}: \mathcal{X}_0 \rightarrow \mathcal{X}_1} \|x - \phi(x)\|_2$$

where x and y denote the 3D coordinates of the points, ϕ^{EMD} denotes a one-to-one injective mapping that only exists when $|\mathcal{X}_0| \leq |\mathcal{X}_1|$. In practice, a linear assignment algorithm is used to calculate ϕ^{EMD} . We also find that the three average operations reduce the magnitudes of the gradients of the physics parameters to an order of 4 to 5, which is undesired for optimisation; therefore, we use the distances without averaging the values over the point sets.

To calculate the loss between the real and simulated end configurations, one can use the original fusion point cloud \mathcal{X}_T^{real} or the reconstructed particle systems \mathcal{X}_T^{rec} as the target observations. Thus, there are two variants for each loss: $d_{PCD-CD}(\mathcal{X}_T^{real}, \mathcal{X}_T^{sim})$, $d_{PRT-CD}(\mathcal{X}_T^{rec}, \mathcal{X}_T^{sim})$, $d_{PCD-EMD}(\mathcal{X}_T^{real}, \mathcal{X}_T^{sim})$, $d_{PRT-EMD}(\mathcal{X}_T^{rec}, \mathcal{X}_T^{sim})$. When reconstructing \mathcal{X}_T^{rec} , we decrease the particle density for filling the voxels slowly until $|\mathcal{X}_T^{rec}| \leq |\mathcal{X}_T^{sim}|$ is satisfied. When computing losses involving the real-world point clouds, we

downsample it with a voxel radius of 0.005 so that the number of points in it is about 1/6 of the number of particles in the reconstructed system. This is because a point cloud only contains points on the surface, while the reconstructed particle system is non-hollow.

Finally, although the CD and EMD losses are very common in calculating point set distance, they are not intuitive to visualise and understand. Also, as revealed by the experiment results shown in Section 4.3, the CD and EMD losses focus on different spatial aspects of two point sets, which make them biased for result analysis across all loss functions. To make results easier to analyse and compare in a more unbiased way, we further calculate the heightmaps of the fusion point cloud and the simulated particles, denoted as \mathcal{I}_T^{real} and \mathcal{I}_T^{sim} . Specifically, we discretise a $0.11 \times 0.11 \text{ m}^2$ square area centred at the centre of the object into a 32×32 grid. For each cell of the grid, we find the point or particle in the x-y range with the highest z value and assign that value as the height of that cell. To compare the distance between two heightmaps, we simply take their summed pixel difference as follows.

$$d_{Heightmap} = \sum_i \sum_j \|\mathcal{I}_{T,ij}^{real} - \mathcal{I}_{T,ij}^{sim}\|_2$$

In practice, we use the fusion point cloud without downsampling to compute \mathcal{I}_T^{real} . However, a particle density that is too low will result in void cells when generating \mathcal{I}_T^{sim} , and we found that a density value of 4×10^6 particles per cubic metre is good enough for generating 32×32 heightmaps. The heightmap distance is only used for evaluation purposes because we found it leads to very poor optimisation results when used as a loss function.

4. Results

This section presents the design of the experiment and the result analyses including the optimisation and generalisation performances of the DPSI framework in identifying the physical characteristics of elastoplastic matter through simple robot manipulation.

4.1. Experiment design

The performance of the DPSI framework is examined in three steps: loss landscape analysis, in-distribution performance analysis and out-of-distribution generalisation analysis. The first step visualises the four loss functions (PCD-CD, PRT-CD, PCD-EMD and PRT-EMD) to help understand how the physics parameters are related to the loss value distributions. The second step investigates the in-distribution performance of the proposed DPSI framework, that is, whether it can produce realistic simulations for unseen object configurations with the same motions used in

parameter identification. The loss landscape and in-distribution performance analyses are carried out at two levels of contact complexity. At each level, the performances are compared across four loss functions and five optimisation datasets (see data collection description below). The last step explores the out-of-distribution generalisability of the DPSI by inspecting the simulation accuracy of three unseen longer motions using identified parameters from the contact complexity level 2 experiment.

4.1.1. Data collection. For each optimisation motion (i.e., Poking-1, Poking-2, Poking-shifting-1, Poking-shifting-2), we collect datapoints by executing the motions on the objects with different end-effectors. Each datapoint contains the multi-viewpoint point clouds of the object before and after a manipulation motion. For each of the two complexity levels, five optimisation datasets are created to examine the data hungriness of the proposed DPSI framework.

With level (Lv.) 1 as an example, the first dataset is created with both motions (Poking-1 and Poking-2). Each motion is performed twice using three different end-effectors. Therefore, 6 datapoints are created for each motion, and a total of 12 datapoints are created for both motions. The second dataset for Lv. 1 is similar, except that only one datapoint is created from each motion-effector pair, resulting in three datapoints for each motion, and a total of six datapoints for both motions. The other three datasets for Lv. 1 are more straightforward, with only one datapoint collected for each dataset. The three datapoints are created using the three end-effectors individually, performing the second motion from Lv. 1 (i.e., Poking-2). We name these datasets as *12-mix*, *6-mix*, *1-rectangle*, *1-round* and *1-cylinder*, respectively. Data collection for Lv. 2 is similar, except that the two motions of Poking-1 and Poking-2 are replaced by Poking-shifting-1 and Poking-shifting-2, respectively. The datasets will be open-sourced upon the acceptance of the article.

For in-distribution validation, with each contact complexity level, we collect two extra datapoints from the second motion (i.e., Poking-2 and Poking-shifting-2) with all three end-effectors, resulting in 12 validation datapoints. For out-of-distribution validation, where we collect two datapoints with each of the three long-horizon motions, resulting in six datapoints.

For each datapoint, the real-world object is initially roughly shaped into a convex shape. We then acquire its point cloud by fusing the captured multi-view data. Next, we calculate the coordinates of the top centre of the object, where the end-effector tip will be moved to, and the motion will always be executed from there. We perform the same point cloud capturing and fusion process after a motion is executed. Each datapoint takes no more than 2 min to collect, with most of the time spent on capturing point clouds.

4.1.2. Optimisation and validation. At each contact level, for each pair of loss function and dataset, gradient descent is

carried out with the Adam algorithm (Kingma and Ba, 2015) with three random seeds for 100 iterations (gradient updates). For each datapoint, the simulation loads the initial particle configuration, simulates the motion and produces the resultant particle state, which is used to compute losses and gradients. Each optimisation iteration goes over all datapoints within a dataset and takes the average gradients to update the parameters. In-distribution validation is done after every gradient update with the validation dataset, simulating all datapoints and calculating the losses.

Table 2 displays the step sizes and the value ranges of the parameters of interest for gradient updates. To determine the appropriate step sizes, we randomly sample 20 sets of parameter values and compute the statistical means and standard deviations of the gradients for each parameter based on the *12-mix* dataset. We then randomly initialise the parameters, run the Adam optimiser for 100 gradient updates with uniformly randomly generated gradients using the statistical means and deviations with different step sizes, and select the step size that allows the randomised parameters to converge in about 70 gradient updates. The value ranges of the parameters are selected according to either the definitions of the physics models (Jones, 2009; Stomakhin et al., 2013) or values reported by various research on manual parameter identification experiments with similar real-world materials, such as soil and clay (StructX, 2014-2024; Waheed and Asmael, 2023). However, a too-high value of Young’s modulus E was found to cause numerical instability in the simulation with the particles exploding away. Thus, an empirical upper bound for E was determined by slowly increasing it until the simulation became unstable.

4.2. Loss landscape analysis

The loss landscapes, including point-cloud and particle chamfer distance (PCD CD, PRT CD), point-cloud and particle earth mover’s distance (PCD EMD, PRT EMD), are computed over three pairs of physics parameters: (E, ν) , (σ_y, ρ) and (μ_t, μ_m) . For each pair, we set other parameters to fixed values and compute the losses with the parameters of interest discretised into 30 intervals. Figure 3 displays the loss landscapes computed at the two contact complexity levels and five datasets. To facilitate distribution-level

Table 2. Step sizes and the ranges of values for the parameters: Young’s modulus E , Poisson’s ratio ν , yield stress σ_y , material density ρ and the friction coefficients of the table η_t and end-effector η_m .

Parameter	E (kPa)	ν	ρ (kg/m ³)	σ_y (kPa)	η_t	η_m
Step size	4000	0.01	10	500	0.01	0.01
Min value	10000	0.01	1000	1000	0.01	0.01
Max value	300000	0.48	2000	20000	2.0	2.0

comparison, the loss landscapes are centralised to have zero means (subtracting the mean of the matrix). The following observations can be drawn.

Firstly, we compare the landscapes in Figure 3 vertically to examine the sensitivity of the loss functions against different parameters. At contact complexity level 1, it shows that the CD and EMD losses exhibit quite similar value changing directions along the E , ν and ρ axes, while quite the opposite directions along the σ_y axis. This can be observed from the vertically reversed colour distributions between the first and last two rows of landscapes against σ_y and ρ . On the other hand, the losses at level 2 contact complexity exhibit similar distributional patterns along all parameter axes. In addition, at level-2 contact complexity, the PRT CD loss is more sensitive (exhibiting more drastic value changes) along the ν axis than other losses (see the second row of the level-2 landscapes against E and ν). In contrast, the values of the PRT CD and both EMD losses are more sensitive along the σ_y axis (the last three rows of the level-2 landscapes against σ_y and ρ). The level-2 landscapes against the friction coefficients (μ_t and μ_m) show that all losses distribute very similarly over these two parameters. Overall, the EMD losses seem to be more sensitive to all the parameters, as their loss landscapes tend to be less flat.

These observations indicate that, with the adopted von Mises plasticity model, whose plastic deformations are governed by the yield stress (σ_y) of the material, the CD and EMD losses may focus on different aspects of the point sets, as they tend to disagree with each other distributionally along the σ_y axis. Also, we may expect the EMD losses to perform better in parameter identification because they exhibit higher sensitivity against most parameters.

Secondly, by comparing the landscapes horizontally, Figure 3 shows the distributional patterns of most (not all) loss landscapes are quite similar across the five datasets. This means that the loss values, and thus the optimisation processes, are not sensitive to the number of datapoints used in computation. This observation allows us to expect the recovery of physics parameters with small data, even with a single interaction experience.

Thirdly, many of the loss distributions display large areas of flat regions, where the loss values are very similar. This indicates many optimisation saddle points and that multiple parameter value combinations may serve as plausible solutions, describing the same physical characteristics. Whether this is an issue from an inaccurate physics model or a general fact of the real-world dynamics remains to be determined.

4.3. In-distribution performances

This subsection investigates the in-distribution performance of the parameter identification task through gradient descent at two levels of contact complexity. For each level,

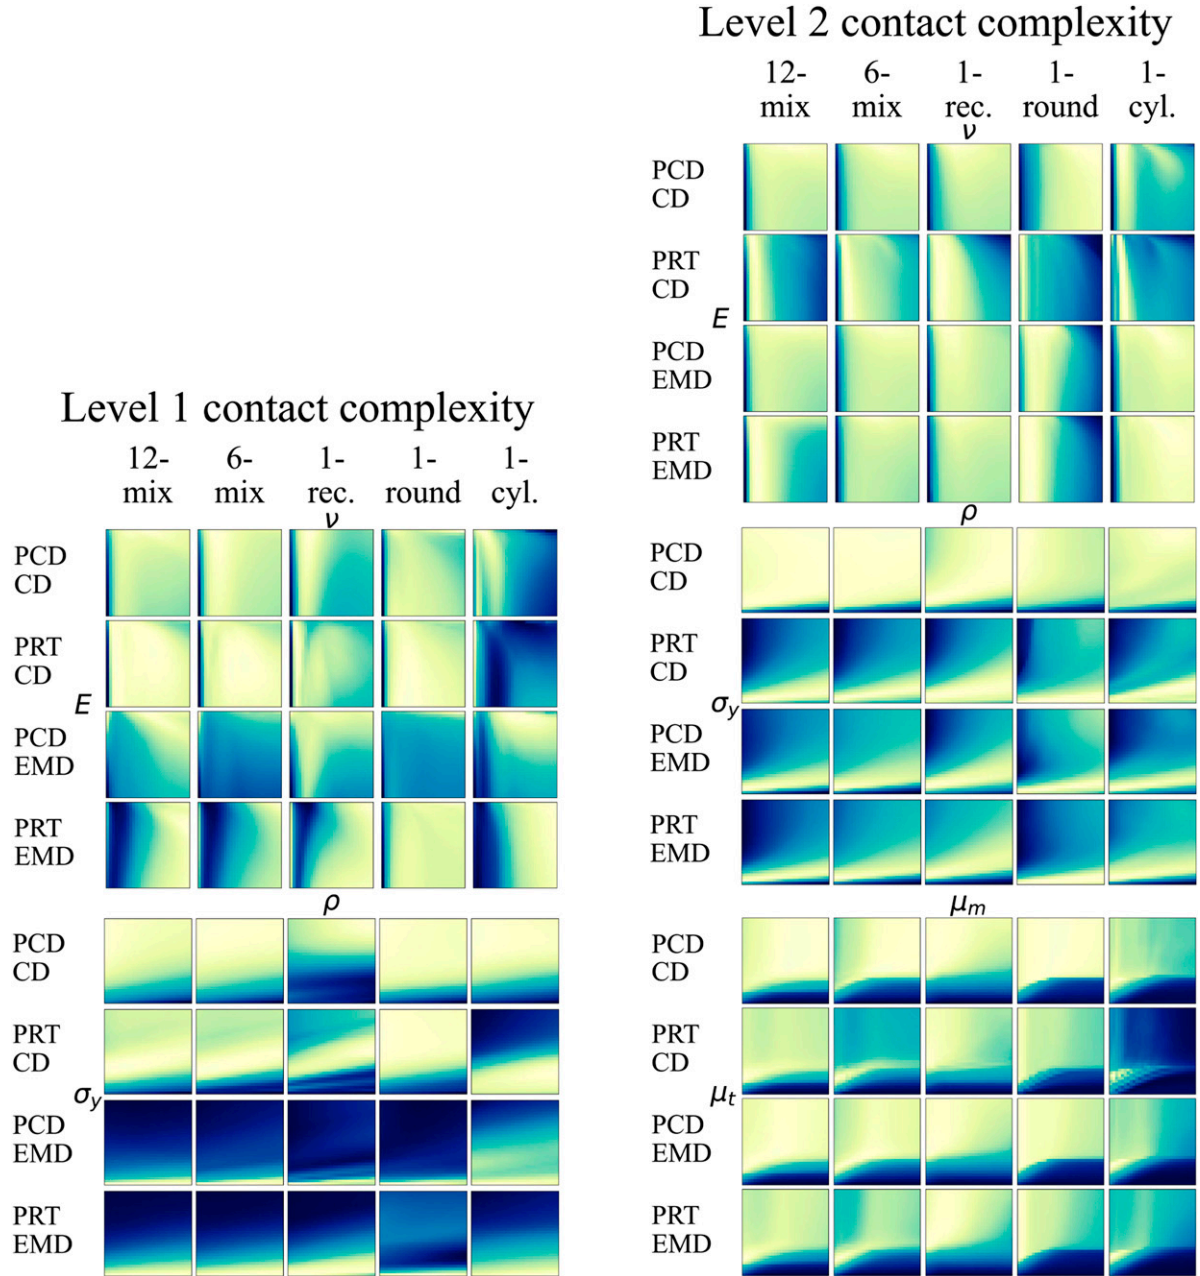


Figure 3. Loss landscapes (centralised to have zero means) at level-1 (left) and level-2 (right) contact complexity over pairs of physics parameters. Darker colours represent higher loss values. Parameters: Young’s modulus (E) against Poisson ratio (ν), yield stress (σ_y) against material density (ρ) for both levels, and table friction coefficient (μ_t) against manipulator friction coefficient (μ_m) for level 2. Other parameter values are fixed while computing the losses over one pair of parameters. Each column presents the losses computed with a different dataset. Each row presents the values of a different loss function. PCD: point cloud. PRT: particle system. CD: Chamfer distance. EMD: earth mover’s distance.

experiments are conducted with the four loss functions and five datasets. To evaluate the performances thoroughly, the four loss functions and a heightmap-based distance function are used to evaluate the differences between the real and simulated manipulation results using the in-distribution validation dataset. This enables the observation of the influences on other distance functions from minimising each objective. In particular, this subsection investigates the following questions.

- Can the loss functions be minimised (does optimisation converge)?
- Do the loss functions agree with each other in such system identification tasks?
- If local minima appear and multiple solutions exist, do they produce visually distinct manipulation results?
- How does the number of datapoints affect optimisation?

- Does DPSI produce parameter values that are physically realistic and interpretable? Are they close to the parameter values reported in relevant literature and studies?

4.3.1. Quantitative results. We start by analysing the quantitative results of the parameter identification task. The changes of the validation losses and parameter values over the course of optimisation at both contact complexity levels are shown in [Figures 4](#) and [5](#). The best parameter values corresponding to the lowest validation heightmap loss among the three random seeds are summarised in [Table 3](#).

First of all, we start by observing the tendency of convergence. The top five rows in [Figures 4](#) and [5](#) show that, at both contact complexity levels, most of the individual validation loss curves (dotted lines) tend to stabilise and converge, which indicates that DPSI can effectively converge to local minima. The last four rows in [Figure 4](#) and the last six rows in [Figure 5](#) reveal that the parameters converge to different solutions. [Table 3](#) also shows that multiple parameter solutions exist for a low validation heightmap distance at both levels of experiments. These observations mean the parameter identification task at both contact complexity levels does converge but has multiple local minima and solutions, aligned with the large flat regions observed from the loss landscapes.

However, the parameter distributions in [Figure 5](#) clearly show that some parameters converge to much smaller and distinct value regions than those found in the level-1 experiments. More specifically, the values of Young’s modulus (E), Poisson’s ratio (ν) and yield stress (σ_y) become more converged, while they are more dispersed in the level-1 results. Interestingly, the material density (ρ) remains quite dispersed at both contact complexity levels. This is not surprising because, compared to the level-1 motions, the higher contact complexity induced by the shifting parts of the level-2 motions naturally poses a stricter selection range for the physics parameter values.

Also, some parameter values are closer to empirical studies of soft/hard/saturated clay in the level-2 experiments. For example, Young’s modulus values are closer to the reported range of 5000 to 54000 kPa, the material density seems to be closer to 1400 kg/m³, and Poisson’s ratio is closer to the reported range of 0.4 to 0.5 ([Waheed and Asmael, 2023](#); [StructX, 2014–2024](#)). However, the recovered yield stress values are higher than the reported values for clays (210 to 600 kPa) ([Rehman et al., 2018](#)) because the reported values were found to make the material collapse in simulation. Based on the empirical values from [Waheed and Asmael \(2023\)](#); [StructX \(2014–2024\)](#) and the loss landscapes, we hand-pick a set of parameter values and calculate the validation losses at both levels. As denoted in [Table 3](#), the resultant performances are reasonable but not as good as the performances of the DPSI method.

Secondly, the top five rows in [Figure 4](#) reveal that, in most cases, the CD and EMD losses have a negative correlation. For example, the red curves of the first column

show that minimising the PRT EMD loss reduces both EMD losses and the heightmap distance, but increases both CD losses. Also, [Table 3](#) shows that optimising the CD losses tends to produce large yield stress values (σ_y), while optimising the EMD losses tends to produce small ones. Similar observations can be made from the level-2 results, but only the PCD CD validation losses tend to exhibit opposite tendencies compared to the other losses (observe the blue lines that deviate from others in [Figure 5](#)). This is aligned with the analyses of the loss landscapes at both contact complexity levels, where reversed distributional patterns of the CD and EMD losses were observed. Furthermore, the EMD losses tend to produce smaller validation heightmap distances at both levels as shown by the green and red lines in [Figures 4](#) and [5](#) and the average heightmap distances in [Table 3](#), indicating better sim2real alignment. This again is aligned with the analysis of the level-2 loss landscapes, where only the PCD CD loss distributes significantly differently along the σ_y axis.

Thirdly, the results show that it is possible to obtain comparable performances with just one datapoint at both contact complexity levels. This can be concluded by comparing the results horizontally: (1) most validation curves of the top five rows in [Figures 4](#) and [5](#) with 1-datapoint show highly similar tendencies with the 6-mix and 12-mix datasets, and (2) the last four rows in [Figure 4](#) and last six rows in [Figure 5](#) also show that the parameters found with different datasets mostly converge to similar value regions. Moreover, [Table 3](#) shows that the lowest validation heightmap distances at level-1 contact complexity and the third lowest at level-2 are achieved by optimising with the 1-rec. and 1-round datasets. These results are aligned with the analysis of the loss landscapes, where similar loss distributions are observed among different datasets.

4.3.2. Qualitative results. We now examine the manipulation processes simulated using the best physics parameters corresponding to the lowest validation heightmap losses presented in [Table 3](#). [Figure 6](#) shows the resultant particles and heightmaps after the manipulation of objects at two contact complexity levels with three end-effectors (two object configurations are included for each contact level and end-effector).

Firstly, all the simulated particles and heightmaps post-manipulation are highly similar to the ground truths across the loss functions and datasets. This shows that minimising any of the four loss functions with any of the five datasets can indeed reproduce visually plausible manipulation results close to the real-world system. It indicates that DPSI is not data-hungry and it is robust to the choice of common point-based loss objectives. These characteristics are highly preferred in robotic applications, as they lead to smaller data collection costs, simpler observation preprocessing, and simpler loss function engineering.

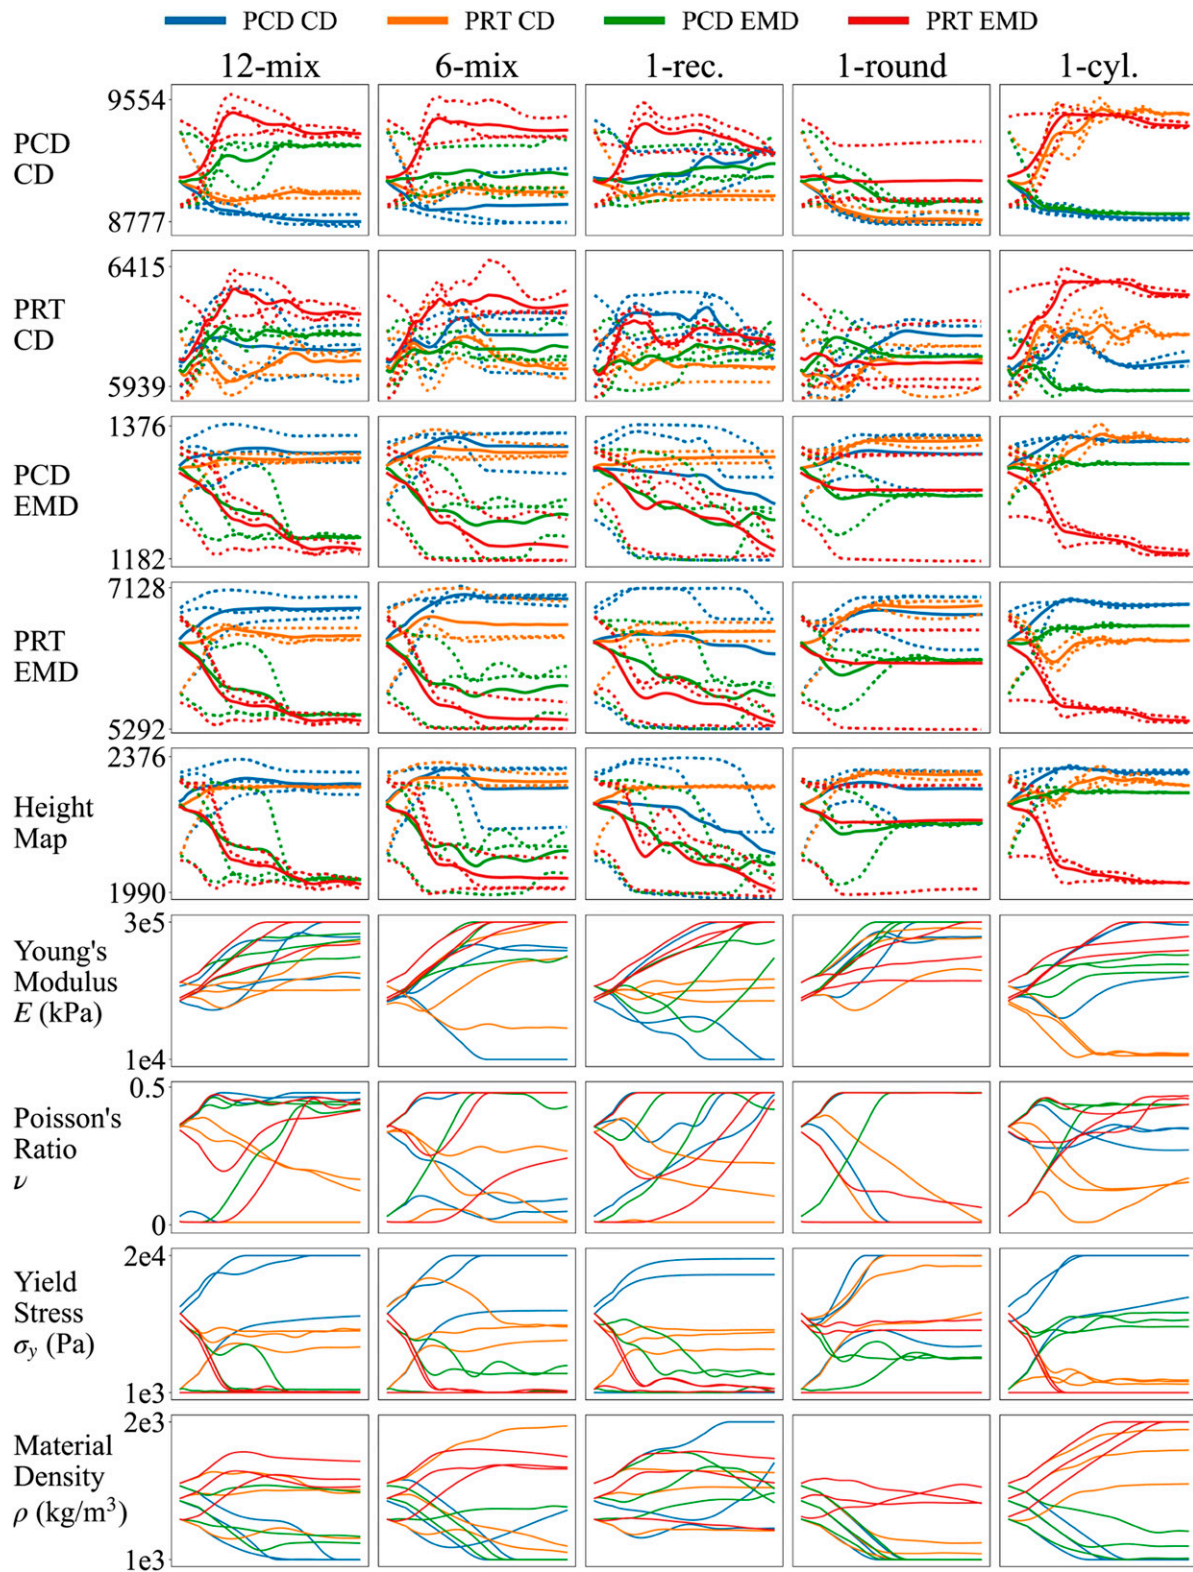


Figure 4. Validation losses (top five rows) and parameter values (last four rows) over 100 gradient updates at level-1 contact complexity. Each column presents the results of optimising with a different dataset. Each row shows the changes of an evaluation metric or a parameter, denoted on the left. In each figure, different colours indicate the results of minimising a different loss function, as labelled by the legend on the top. For the top five rows, each colour has three dotted lines corresponding to the results of three random seeds and a solid line corresponding to their means. For the last four rows, each colour has three solid lines corresponding to the results of three random seeds.

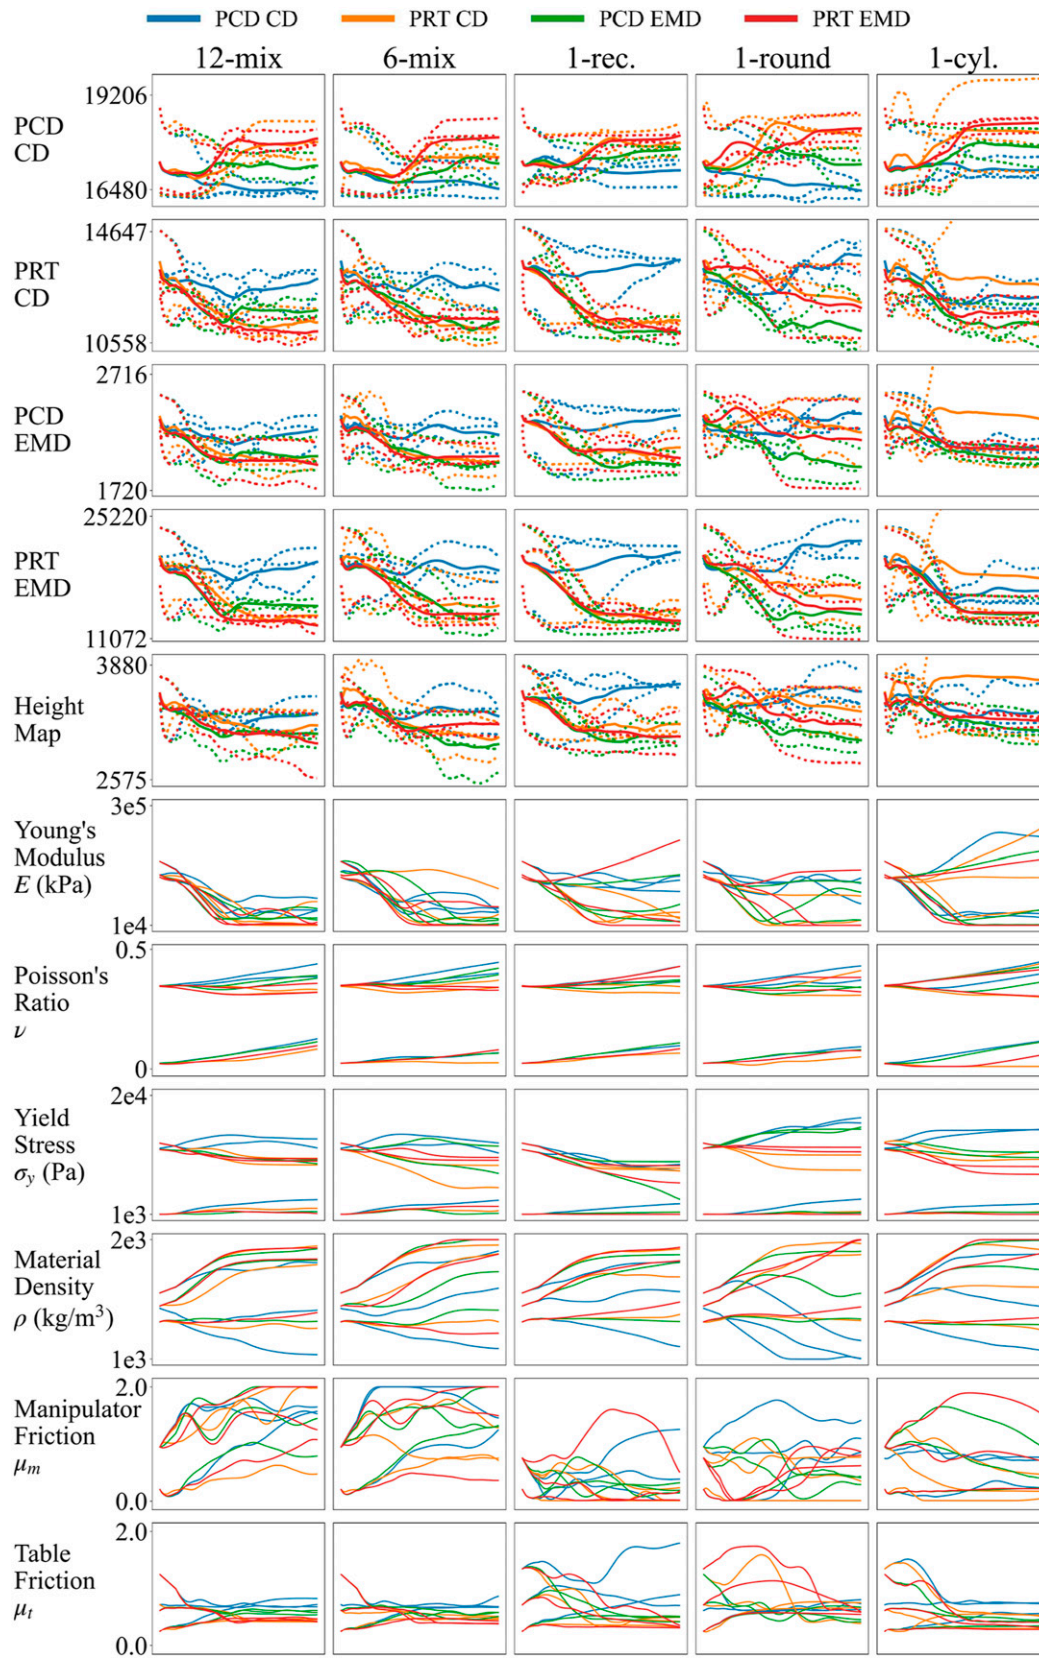


Figure 5. Validation losses (top five rows) and parameter values (last six rows) over 100 gradient updates at level-2 contact complexity. Each column presents the results of optimising with a different dataset. Each row shows the changes of an evaluation metric or a parameter, denoted on the left. In each figure, different colours indicate the results of minimising a different loss function, as labelled by the legend on the top. For the top five rows, each colour has three dotted lines corresponding to the results of three random seeds and a solid line corresponding to their means. For the last four rows, each colour has three solid lines corresponding to the results of three random seeds.

Table 3. The parameter values corresponding to the lowest validation heightmap distance in each optimisation case at both contact complexity levels.

Contact complexity	Loss	Metric/Parameter	Dataset					Avg. Height map (mm)
			12-mix	6-mix	1-rec.	1-round	1-cyl.	
Level 1	PCD CD	Min. height map (mm)	2267.44	2176.21	1973.81	2183.30	2327.01	2185.55
		Young's modulus (kPa)	300,000	10,000	300,000	300,000	294,991	
		Poisson's ratio	0.480	0.480	0.480	0.480	0.348	
		Yield stress (kPa)	20,000	12,367	1000	7527	20,000	
		Material density (kg/m ³)	1000	1380	1225	1000	1000	
	PRT CD	Min. height map (mm)	2287.21	2293.42	2285.98	2311.87	2290.71	2273.83
		Young's modulus (kPa)	189,900	300,000	179,373	196,393	20,335	
		Poisson's ratio	0.165	0.010	0.223	0.010	0.178	
		Yield stress (kPa)	7373	10,286	7016	12,358	2466	
		Material density (kg/m ³)	1157	1088	1210	1000	1797	
	PCD EMD	Min. height map (mm)	2024.78	2039.95	2040.51	2186.90	2272.50	2112.93
		Young's modulus (kPa)	265,420	237,337	270,940	300,000	193,790	
		Poisson's ratio	0.445	0.439	0.419	0.480	0.436	
		Yield stress (kPa)	1403	1036	1382	5839	10,138	
		Material density (kg/m ³)	1487	1379	1482	1000	1009	
PRT EMD	Min. height map (mm)	2000.36	1999.14	1980.35	1999.34	2016.10	1999.06	
	Young's modulus (kPa)	300,000	300,000	300,000	300,000	271,247		
	Poisson's ratio	0.453	0.480	0.480	0.478	0.442		
	Yield stress (kPa)	1000	1000	1059	1000	1000		
	Material density (kg/m ³)	1542	1670	1214	1518	2000		
*Validation height map (mm) with hand-picked parameters: 2122.59								
Level 2	PCD CD	Min. height map (mm)	3163.64	3091.73	3622.58	3423.00	3125.04	3285.20
		Young's modulus (kPa)	47,700	42,136	91,785	132,840	28,679	
		Poisson's ratio	0.444	0.451	0.364	0.095	0.403	
		Yield stress (kPa)	11,407	10,479	7839	15,590	14,514	
		Material density (kg/m ³)	1810	1919	1819	1000	1880	
		Table friction	0.725	0.696	1.816	0.696	0.715	
		Manipulator friction	1.495	1.448	1.264	1.037	0.704	
	PRT CD	Min. height map (mm)	3088.31	2865.14	3104.15	3031.56	2999.36	3017.70
		Young's modulus (kPa)	66,514	47,786	28,023	24,068	256,786	
		Poisson's ratio	0.331	0.376	0.317	0.309	0.451	
		Yield stress (kPa)	1904	1809	8363	8063	1537	
		Material density (kg/m ³)	1267	1284	1926	1965	1225	
		Table friction	0.473	0.488	0.414	0.419	0.487	
		Manipulator friction	0.494	0.604	0.249	0.293	0.084	
	PCD EMD	Min. height map (mm)	2933.23	2667.43	2959.43	2895.74	2979.04	2886.97
		Young's modulus (kPa)	48,459	14,959	135,207	113,474	12,017	
		Poisson's ratio	0.392	0.400	0.374	0.339	0.369	
		Yield stress (kPa)	1417	1176	1359	1340	10,921	
		Material density (kg/m ³)	1319	1409	1313	1331	1995	
		Table friction	0.544	0.525	0.508	0.453	0.434	
		Manipulator friction	0.805	1.308	0.194	0.381	0.878	
	PRT EMD	Min. height map (mm)	2598.28	3060.10	2847.69	2760.83	3172.49	2887.88
		Young's modulus (kPa)	27,779	10,000	222,856	144,827	10,000	
		Poisson's ratio	0.359	0.330	0.435	0.383	0.303	
		Yield stress (kPa)	1008	9635	1000	1000	7377	
		Material density (kg/m ³)	1396	2000	1492	1445	2000	
		Table friction	0.487	0.373	0.311	0.543	0.411	
Manipulator friction		1.013	1.478	0.015	0.625	1.433		
*Validation height map (mm) with hand-picked parameters: 2985.5688								

Bolded texts denote the three lowest losses at each level. It can be seen that the PRT EMD loss tends to produce the best results and DPSI can find parameters that outperform hand-picked parameters (from the geotechnical literature) with even just one interaction datapoint.

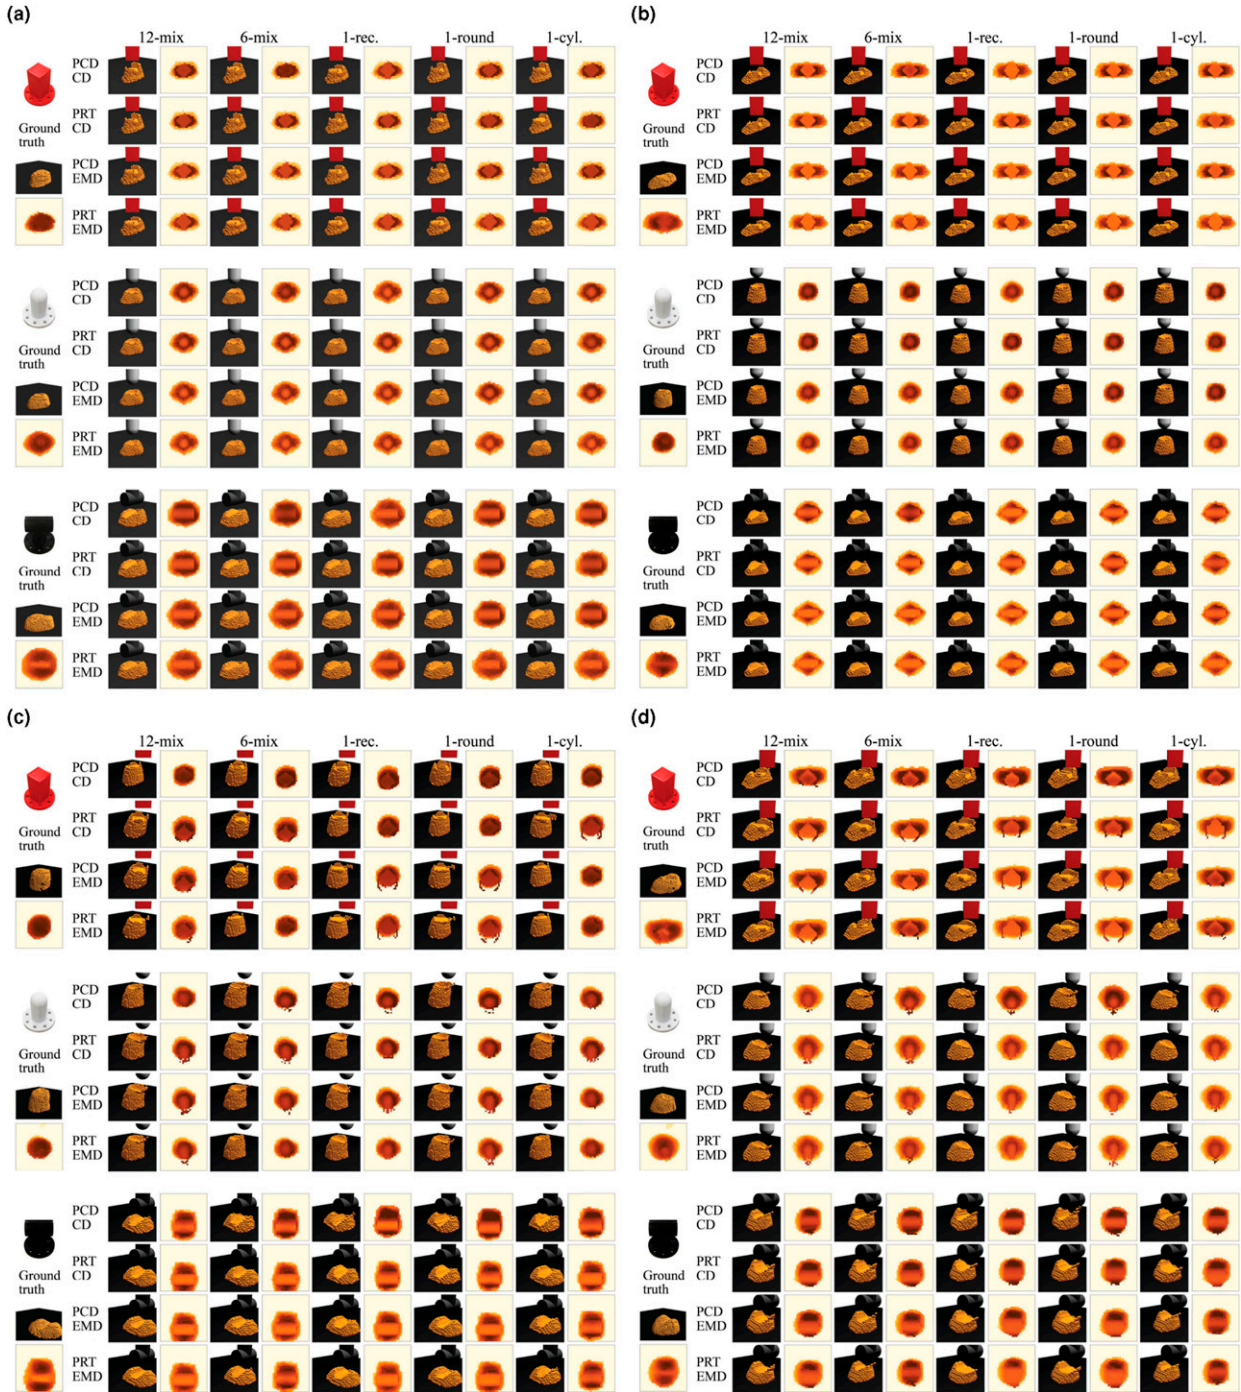


Figure 6. The particles and heightmaps of the objects after applying the second motions at contact level 1 (a and b) and 2 (c and d), simulated with the best set of physics parameters. Darker colours of the heightmaps indicate greater heights of the object. For each combination of the loss function and dataset at each level, three trajectories on two object configurations are simulated (3 effectors \times 2 datapoints). There are three groups in each subfigure, with the ground truth point clouds and heightmaps placed in the leftmost column. In each group, a row shows the results corresponding to a loss function and a column results corresponding to a dataset. (a) Level-1 contact, object configuration 1; (b) Level-1 contact, object configuration 2; (c) Level-2 contact, object configuration 1; and (d) Level-2 contact, object configuration 2.

Secondly, by examining the details of the heightmaps more carefully, it shows that optimising the CD losses tends to produce heightmaps with darker colours (which represent greater heights of the object) while optimising the EMD losses tends to produce heightmaps where

the objects look slightly bigger and wider. This is especially true at level-1 contact complexity, where the heightmaps in the first two rows for each object configuration in Subfigures 6(a) and 6(b) are generally darker than those in the last two rows (see Figure 7 for

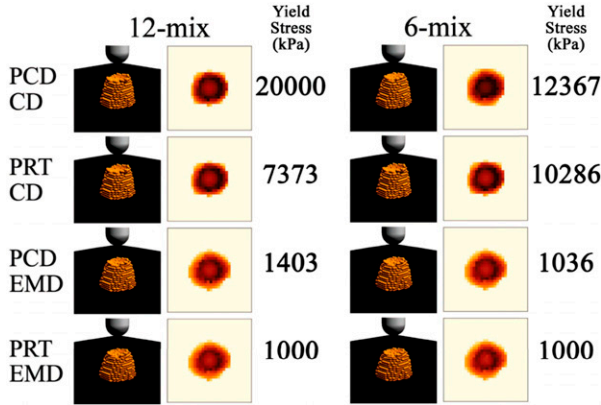


Figure 7. Selected examples from Figure 6(b) to show that the CD losses produce darker-coloured heightmaps while the EMD losses produce lighter-coloured and wider heightmaps, which can be intuitively explained by the optimised yield stress values and their potential difference in spatial focus.

enlarged examples). At level 2, the PRT CD loss sometimes changes its focus on the x and y directions, producing lighter-colour heightmaps (see the second rows of the heightmaps associated with the rectangle and round end-effectors in Subfigures 6(c) and 6(d)). This observation indicates that optimising the CD losses tends to produce particles that match the height of the ground truth shapes, that is, in the z direction, while optimising the EMD losses tends to match the x and y directions.

This may be caused by the fact that these losses distribute differently along the yield stress axis (the CD and EMD losses at level 1, the PRT CD loss at level 2). According to physics intuitions, a smaller yield stress will cause the object to yield more easily and respond more drastically to the poking forces, hence the more spreading in the x and y directions and more compressing in the z direction, and vice versa. One can see that a greater yield stress value in Table 3 corresponds to a darker heightmap in Figure 6, and vice versa.

Thirdly, though the particles and heightmaps are very similar, they are produced by quite different parameter combinations. From the results at contact level 2 in Subfigures 6(c) and 6(d), one can observe a correlation between the friction coefficients and the material density—the three key parameters that determine how much the object will be moved in the shifting motion direction. The parameters in Table 3 and their corresponding visualisations exhibit the following physically plausible relationships (see Figure 8 for a few examples):

1. With similar weights (ρ), the object is moved at a longer distance when the manipulator frictions (μ_m) is greater (compare the parameters and visualisations between the results from the PCD EMD loss with the 12-mix and 1-rec. datasets)

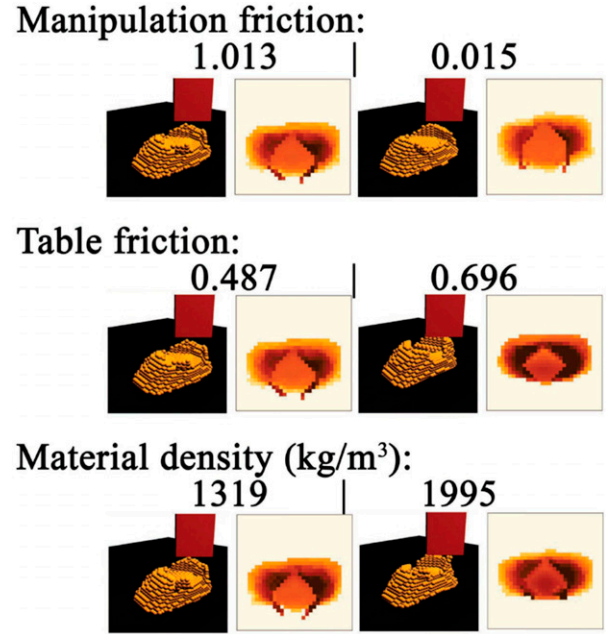


Figure 8. Selected examples from Figure 6(d) to show the physically realistic associations between object movements and three physics parameters, demonstrating the physical plausibility of different physics parameter values.

2. With similar manipulator frictions (μ_m), the object is moved further when the table friction (μ_t) is low (compare the parameters and visualisations between the PCD CD loss with the 1-round dataset and the PRT EMD loss with the 12-mix dataset)
3. With similar friction values, the object is moved further when it is lighter (compare the results from the PCD EMD loss with the 12-mix and the 1-cyl. datasets)

These relationships between parameters and visualisations align with our understanding of real-world physics, demonstrating that DPSI can produce physically plausible and interpretable parameter values. Moreover, the variations in parameter values suggest that, aside from the mismatch of the models and the materials, there is hardly any fixed set of parameter values that can simulate a material in every condition. These parameters will always have some differences in different conditions (e.g., heat, light, and humidity) and they change over time (e.g., drying). Therefore, the fact that DPSI can estimate the parameters with only one interaction makes it highly promising for fast parameter re-identification that can be performed online to recalibrate simulation accuracy efficiently, eliminating the need to keep track of previously identified parameter values.

Fourthly, the visualisations in Figure 6 also reveal limitations and potential improvement directions of the physics model. By comparing with the ground truths, the simulated contact areas of the objects always deform more sharply with insufficient elastic returning. The real objects,

on the other hand, tend to respond more elastically after being plastically deformed, hence the higher and smoother surfaces shown in the ground truths. This is more obvious for contact areas with sharp edges, such as the four edges of the rectangle end-effector or the sharp sides of the cylinder end-effector. In addition, at level-2 cases, with the effectors shifting horizontally, some simulated particles are unrealistically displaced and stay floating. In the real world, the displaced parts would fall because of gravity. What's more, the shifting motion in the real world causes the whole object to tilt in the moving direction, while in the simulation the contact impact tends to remain in a much smaller region around the contact area. These inconsistencies can be caused not only by inaccurate modelling, but also by various computation approximations, such as time integration method, simulation step size, and contact handling.

4.3.3. Summary. In short, this subsection presents a detailed examination of the in-distribution performances of DPSI with motions at two levels of contact complexity. By analysing the validation loss curves, the parameter values and their manipulation visualisations, we can answer the proposed questions:

- The loss functions can be minimised, demonstrating the feasibility of DPSI even in the presence of noisy and incomplete point cloud observations.
- The loss functions produce similar visualisations, but do not always agree with each other on the found parameter values. The CD and EMD losses seem to focus on quite different spatial aspects of the point sets. They distribute differently along the yield stress (σ_y) axis and produce different yield stress values. However, it seems that, as the optimisation motion becomes more complex, the loss functions agree more on the parameter solutions.
- There are many local minima and possible parameter solutions, but they produce visually and physically similar manipulation results with physically intuitive interpretability.
- The number of datapoints has a minor influence on the optimisation performances both quantitatively and qualitatively, indicating that DPSI is not data-hungry even in the presence of real-world perception challenges.
- Discussions on the visualisations and parameter values show that DPSI can produce physics parameter values with physically realistic and interpretations.

4.4. Generalisation

This subsection will look at the out-of-distribution performance of DPSI by visualising the manipulation processes of three unseen and much longer motions that induce more complex contact dynamics with the best parameters (Table 3) found at the level-2 contact complexity identification task.

To recall, there are three unseen motions. The triple-poking motion moves the round end-effector to poke the object at three different locations along the y axis. The flattening motion moves the cylinder end-effector to press and flatten the object towards the positive and negative x directions. The poking-180-rotating motion moves the rectangle end-effector to press the object down and then rotate 180° about the z axis. The triple-poking motion is less relevant to the friction coefficients, while the last two motions involve complex frictional contacts. Table 4 summarises the heightmap distances of the simulated results of the unseen motions (each with two object configurations). Figure 9 visualises the trajectories of the real manipulation trajectory and its simulations with the top two best DPSI-found parameters and the hand-picked parameters. The animation videos of them are included in Supplementary materials. Figure 10 presents the resultant particles and their heightmaps of each motion and object configuration simulated by the DPSI-found parameters. The following observations can be drawn.

Firstly, in Figure 9, (as well as the animation video), it can be seen that, for all motion and object configuration pairs, the real and simulated trajectories are highly similar. Figure 10 also shows that, although discrepancies exist, the simulated particles and their heightmaps of the manipulation results are also highly similar across the loss functions and datasets. This visual similarity proves that DPSI can recover physics parameters from simple interactions that allow the accurate simulation of unseen, longer, and more complex manipulation motions, exhibiting robustness against perception and data challenges. Moreover, we note that it is the single-datapoint cases that achieve 11 out of 18 top-three lowest heightmap losses across the three unseen motions as shown in Table 4, where 2 achieved by the hand-picked parameters and 5 by the 12 mix and 6 mix datasets.

Secondly, the found physics parameters retain their physical meanings and interpretability for the unseen and more complex motions. Evidence supporting this can be found in Figure 10 produced by different parameters (see Figure 11 for selected examples).

- The first example is that when the yield stress (σ_y) is smaller, the resultant heightmaps of the unseen motions tend to be lighter, meaning the objects are more compressed down.
- Another example is that, when moving the cylinder end-effector along the x axis after pressing down the object, a higher manipulator friction coefficient (η_m) causes larger whole-body displacement of the object, while a smaller value causes the effector to slip away and compress down the particles along its moving direction. This is shown more clearly in the bottom-right case (cylinder end-effector) of Figure 10, where the resultant object bodies that are more concentrated to the heightmap centres correspond to larger manipulator friction coefficients, while the objects whose upper part of the heightmaps are flattened correspond to smaller one.

Table 4. The lowest pixel-wise heightmap distances (in mm) achieved by the three unseen motions (each with two object configurations), simulated by the best parameters found at the level 2 contact complexity identification task.

Motion	Object	Num. datapoint	12	6	1 (rec.)	1 (round)	1 (cyl.)	Avg. over losses	
Poking-180 rotating	Config. 1	PCD CD	2384.72	2302.12	2394.64	3827.77	2288.37	2639.52	
		PRT CD	3318.00	2845.03	2316.25	2309.38	4499.31	3057.59	
		PCD EMD	2700.57	<u>2306.85</u>	2971.97	3346.45	2255.47	2716.26	
		PRT EMD	2290.37	<u>2299.37</u>	3605.89	2361.34	2357.05	2582.81	
		Avg. over dataset	<u>2673.42</u>	2438.34	2822.19	<u>2961.24</u>	2850.05	—	
	Hand-picked parameters: *2132.04								
	Config. 2	PCD CD	5367.10	5345.52	5473.18	4822.09	5770.72	5355.72	
		PRT CD	4218.07	4227.91	5802.93	5766.13	*3994.08	4801.82	
		PCD EMD	4187.56	<u>4469.51</u>	4300.81	4026.08	5467.30	4490.25	
		PRT EMD	4175.07	<u>4768.15</u>	4323.36	<u>4226.68</u>	4978.39	4494.33	
		Avg. over dataset	<u>4486.95</u>	4702.77	4975.07	<u>4710.25</u>	5052.62	—	
	Hand-picked parameters: 5283.87								
	Triple-poking	Config. 1	PCD CD	2956.06	2919.90	3028.85	3089.34	2949.11	2988.65
			PRT CD	2960.77	2971.96	2962.99	2963.06	*2798.67	2931.49
PCD EMD			2895.86	<u>2875.01</u>	2839.52	2863.56	3003.25	2895.44	
PRT EMD			<u>3100.87</u>	<u>3035.49</u>	2852.95	<u>3034.08</u>	3059.61	3016.60	
Avg. over dataset			<u>2978.39</u>	2950.59	2921.08	2987.51	2952.66	—	
Hand-picked parameters: 2983.62									
Config. 2		PCD CD	*4175.07	4768.15	4323.36	4226.68	4978.39	4630.09	
		PRT CD	4395.60	4399.77	4731.59	4685.70	4245.47	4491.62	
		PCD EMD	4430.33	<u>4495.40</u>	4417.59	4426.74	4379.40	4429.89	
		PRT EMD	<u>4529.65</u>	<u>4304.59</u>	4300.22	<u>4272.90</u>	4536.17	4388.70	
		Avg. over dataset	<u>4508.76</u>	4467.01	4499.59	4505.60	4444.43	—	
Hand-picked parameters: 4517.24									
Flattening		Config. 1	PCD CD	2443.54	2536.84	2153.18	2244.94	2403.77	2356.45
			PRT CD	2020.59	1988.33	2496.75	2547.35	1860.49	2182.70
	PCD EMD		2149.35	*1835.63	2233.75	2196.98	2545.52	2192.25	
	PRT EMD		<u>2275.80</u>	<u>2439.66</u>	2407.77	<u>2287.72</u>	2539.93	2390.18	
	Avg. over dataset		<u>2222.32</u>	2200.12	2322.86	<u>2319.25</u>	2337.43	-	
	Hand-picked parameters: 2051.06								
	Config. 2	PCD CD	5325.90	5378.31	4752.87	4756.35	5743.95	5191.48	
		PRT CD	5622.73	5387.05	5282.32	5119.63	5817.57	5445.86	
		PCD EMD	5058.42	4697.61	5798.36	5891.09	5190.65	5327.23	
		PRT EMD	4996.47	<u>5460.32</u>	5686.20	<u>5152.61</u>	5493.27	5357.77	
		Avg. over dataset	<u>5250.88</u>	5230.82	5379.94	<u>5229.92</u>	5561.36	—	
	Hand-picked parameters: *4483.75								

Bolded texts denote the top three lowest losses. Underlined texts denote the losses achieved by the cases which have achieved the top-three lowest in-distribution validation performances.

- The last example is that, for the poking-180-rotating motion results shown at the top of Figure 10, when the objects are more rotated by the end-effector, at least two of the following three conditions can be verified true from Table 3 at the same time: a relatively small material density, a relatively small table friction coefficient, and a relatively large manipulator friction coefficient. This is physically intuitive as any two of the three conditions would make the object easier to move.

Thirdly, the found solutions are not globally optimal. Despite the visual similarity, Table 4 shows that the best generalisation results are not always achieved by the cases with the best in-distribution validation performances, as the underlined and bolded texts rarely coincide. In addition, though the hand-picked parameters do achieve the best results in two out of six object simulations, they are not the best fit for every scenario. This again suggests the importance of re-fitting the parameters on the fly for every new

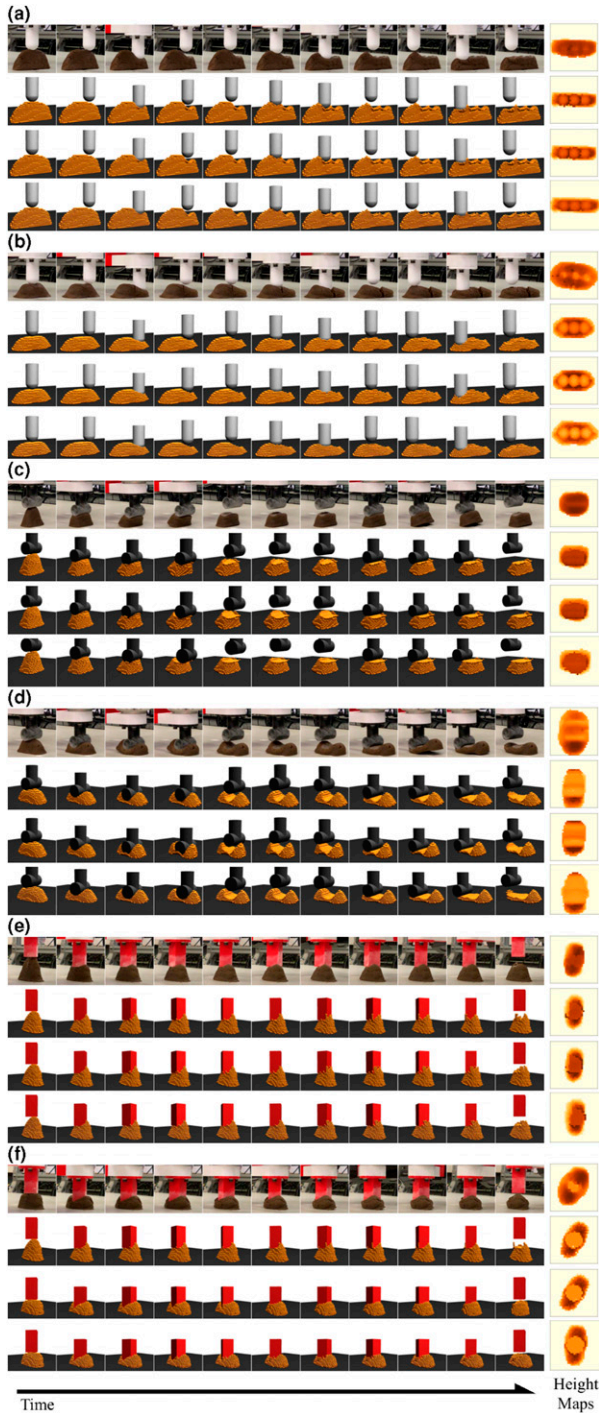


Figure 9. The long-horizon manipulation motions. Every subfigure visualises from top to bottom the real manipulation process and its simulations with the top two best DPSI-recovered parameters and hand-picked parameters. (a) Triple-poking, object configuration 1; (b) triple-poking, object configuration 2; (c) flattening, object configuration 1; (d) flattening, object configuration 2; (e) poking-180-rotating, object configuration 1; and (f) poking-180-rotating, object configuration 2.

scenario as there are hardly two objects that are identical especially for elastoplastic materials. Qualitatively, from the heightmaps and trajectories in Figures 9 and 10, one can also recognise certain visual discrepancies between the real objects and their simulation counterparts, indicating a certain degree of mismatch between the found parameters and the actual solutions. This could be improved by further diversifying the data (and the motions that produce them) used for optimisation. Other directions for improvement may include better observational noise removal or better physics model selection/approximation.

Finally, the modelling limitations observed in the in-distribution visualisation still exist in the unseen motion visualisations (Figures 9 and 10). The insufficient elastic returning at contact areas, especially with sharp manipulator edges, can be seen from all visualisations where the primary contact areas of the object rarely return to heights close to the real object (compare also the enlarged heightmaps in Figure 11). Floating displaced particles can be observed from the poking-180-rotating motion visualisations, where many particles are carried away during rotations. The insufficient spreading of the influence of contacts can be recognised from the first object configuration of the flattening motion visualisations where the object is not tilted enough in the second half of the trajectory. See Figure 12 for enlarged examples.

In summary, the results and discussions above strongly support the argument that DPSI can recover physics parameters that produce the accurate simulation of unseen, longer, and more complex deformable object manipulation dynamics by optimising with very few real-world noisy and incomplete point-cloud data collected by simple and short interacting motions.

4.5. Cloud slime and soil materials

This subsection reports the DPSI performances with cloud slime and soil materials. Note that we use the same elastoplastic physics models in simulation. To do so, we collect a single interaction data point for each material using the cylinder end-effector and the poking-shifting-2 motion, optimise the parameters with the PRT EMD loss from a randomised initialisation for 100 iterations with three random seeds, and validate the results on another data collected with the flattening motion and the cylinder end-effector. The real manipulation trajectory of the flattening motion and its simulations based on the found parameters are shown in Figure 13. The recovered parameters are presented in Table 5.

Figure 13 shows that DPSI manages to recover parameters that simulate cloud slime and soil as closely as

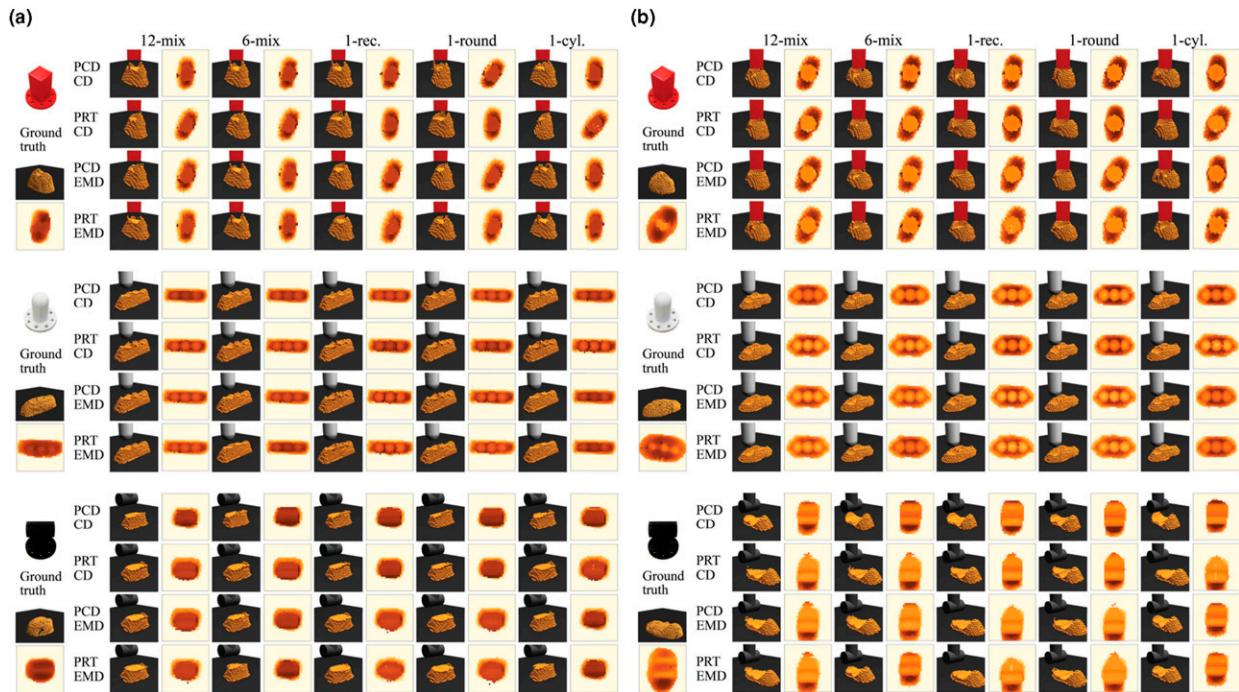


Figure 10. The particles and heightmaps of the objects after applying the long-horizon unseen motions, simulated with the best set of physics parameters found at the level 2 contact complexity experiments. Darker colours of the heightmaps indicate greater heights of the object. For each motion, two object configurations are manipulated with three end-effectors. The results for each trajectory are grouped with the ground truth placed on the left. In each group, a row shows the results corresponding to a loss function and a column results corresponding to a dataset. (a) Object configure 1 and (b) object configure 2.

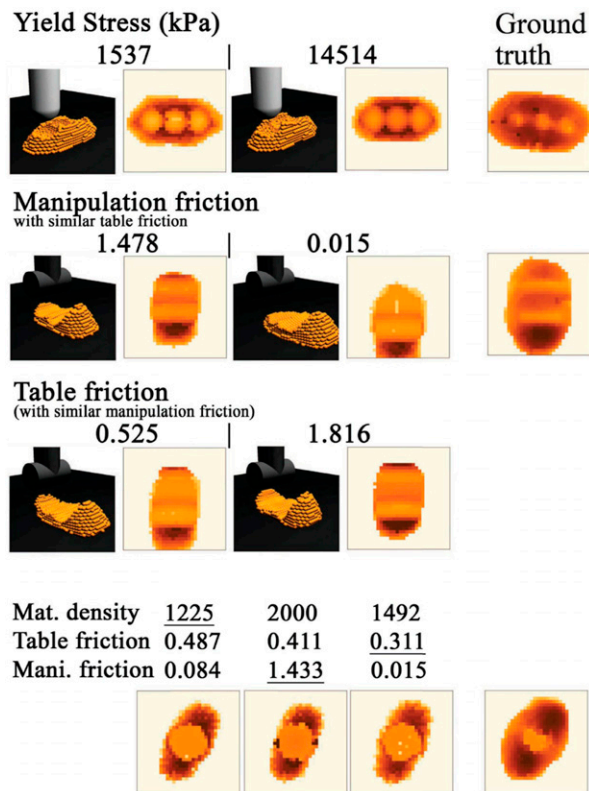


Figure 11. Selected examples from Figure 10 to show the physically realistic associations between the manipulation results and the found parameter values. Comparisons to the ground truths reveal the modelling limitations (insufficient elastic returning).

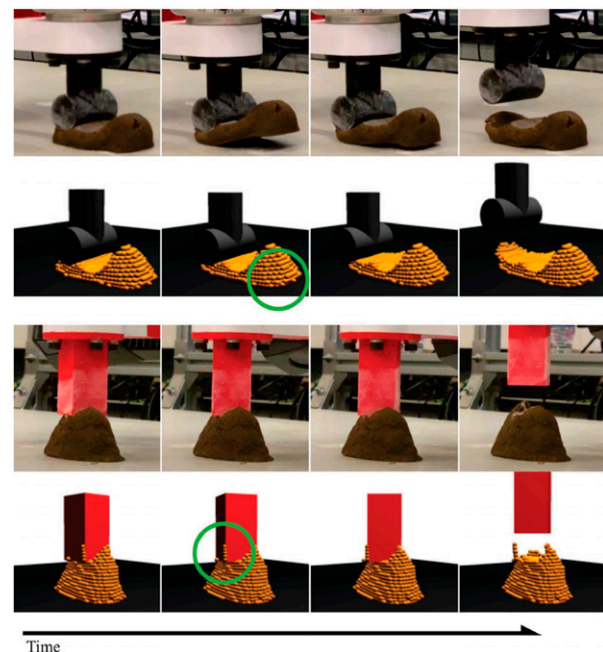


Figure 12. Selected frames from two trajectories in Figure 9 to show modelling limitations (floating particles and insufficient spreading of contacts). The areas around the green circles show that the second part of the flattening motion does not tilt the object up, and there are floating particles in the poking-180-rotating motion.

possible despite the mismatch of the physics models and the target materials. Especially in the case of soil, [Table 5](#) shows that DPSI finds large table friction coefficients and small end-effector friction coefficients to simulate the effect of the soil being flattened, although these are likely not close to the actual values. This suggests that DPSI will seek to find parameters that make the deformations of the simulated object as realistic as possible, regardless of whether the physics models are a good match for the real materials. On the other hand, this also suggests that the manual selection of the physics models is important for DPSI to produce the best deformation simulation.

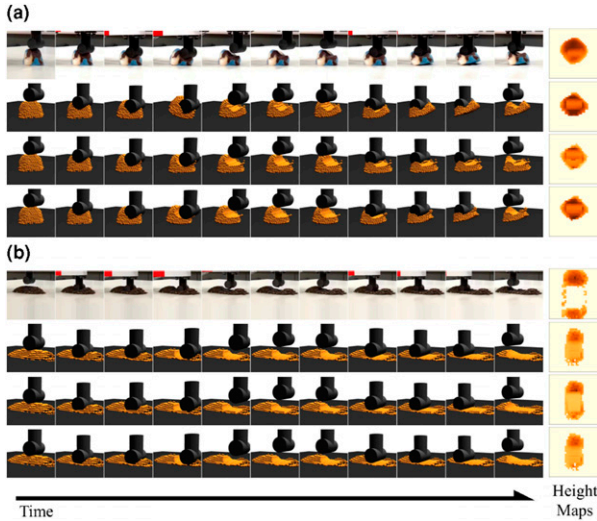


Figure 13. Comparison of the real manipulation trajectories and the resultant height maps for the cloud slime and soil materials and their simulation counterparts based on the recovered physics parameters. In each subfigure, the second to last rows correspond to the results with three random seeds. (a) Cloud slime and (b) soil.

4.6. Cause and treatment to local minima

As shown by the loss landscapes and the optimisation results, although DPSI can find good solutions, it could get stuck in local minima. Therefore, we performed extra experiments to optimise the PRT EMD loss with the three 1-datapoint datasets at contact level 1 with eight more random seeds to examine the random seed sensitivity of DPSI. The results are shown by the black dotted lines in [Figure 14](#).

First of all, these dotted lines show consistent results with the previous conclusions: they converge to stable (though different) loss and parameter values, still reveal the negative relationship between the CD and EMD losses, and can achieve comparable performances with results from larger datasets. On the other hand, it can be seen that the main cause of converging to different solutions is the different initial parameter values, which were determined by the pseudo-random process controlled by different random seed values. This is not surprising as gradient-descent algorithms with complex physics problems are known to be sensitive to the choice of the initial solution ([Antonova et al., 2023](#); [Hu et al., 2020](#)).

To further test the assumption, we perform the same experiments with the four loss functions using the same set of seeds but manually initialise the solutions using the hand-picked parameters in [Table 6](#) selected according to the material/geotechnical engineering literature ([Rehman et al., 2018](#); [Waheed and Asmael, 2023](#); [StructX, 2014–2024](#)). These hand-picked parameter values are less accurate in terms of validation loss as shown in [Table 3](#). However, the (highly overlapped) solid lines in [Figure 14](#) show that with this initial solution, DPSI converges with near-zero variances across eight random seeds, except for optimising the PCD CD loss with the 1-rec. dataset. This confirms that the differentiation process of the DPSI system is largely deterministic and the found solutions rely largely on the initial solutions, the optimisation algorithm and the loss functions. Moreover, with a good guess of the initial values, DPSI could converge very quickly (within 20 iterations shown by the red/green solid lines).

Table 5. Physics parameters recovered by the DPSI method for cloud slime and soil materials with three random seeds.

Random seed		1	2	3
Cloud slime	Young's modulus (kPa)	30,0000	10,6971	11,7831
	Poisson's ratio	0.4	0.33	0.03
	Yield stress (kPa)	13,836	3907	10,469
	Material density (kg/m ³)	1000	1000	1491
	Table friction	1.18	0.13	0.54
	Manipulator friction	2.0	0.01	0.58
	Random seed		1	2
Soil	Young's modulus (kPa)	10,000	88,963	10,696
	Poisson's ratio	0.32	0.33	0.08
	Yield stress (kPa)	9333	1000	10,037
	Material density (kg/m ³)	2000	1653	1823
	Table friction	2.0	1.08	1.83
	Manipulator friction	0.01	0.01	0.01

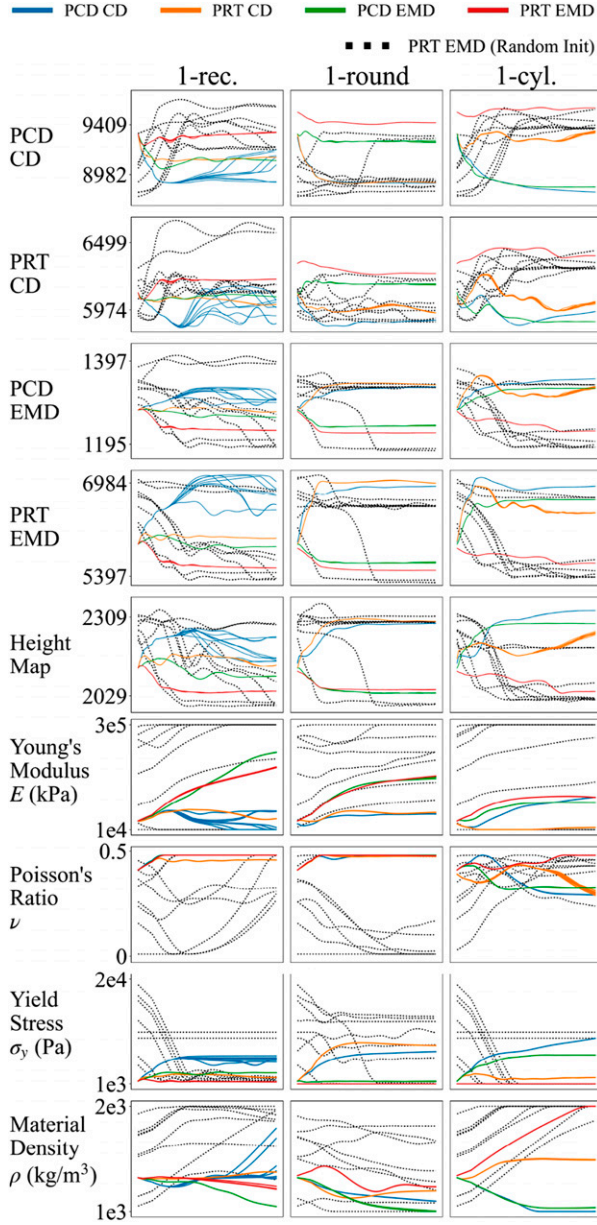


Figure 14. Validation losses (top five rows) and parameter values (last four rows) over 100 gradient updates at level-1 contact complexity from optimising the loss functions with an extra set of 8 random seeds. Solid lines (highly overlapped) are results with the same initial solution, while dotted lines show results of optimising the PRT EMD loss with randomised initial solutions. Each column presents the results of optimising with a different dataset. Each row shows the changes of an evaluation metric or a parameter, denoted on the left. In each figure, different colours indicate the results of optimising a different loss, as labelled by the legend on the top.

4.7. Remark on CD and EMD losses

As mentioned previously, the CD and EMD losses seem to focus on different spatial aspects of the object’s geometry, leading to deformations that match either the horizontal or the vertical directions. With the results of optimising the four losses with the same initial solutions shown in

Table 6. Hand-picked physics parameter values based on the loss landscapes, published literature and open-source website in geotechnical engineering (Rehman et al., 2018; Waheed and Asmael, 2023; StructX, 2014–2024). The yield stress value is slightly higher than the reported values because the reported values were found to cause the material collapse in simulation.

E (kPa)	ν	ρ (kg/m ³)	σ_y (kPa)	η_t	η_m
30000	0.4	1330	1000	0.7	0.7

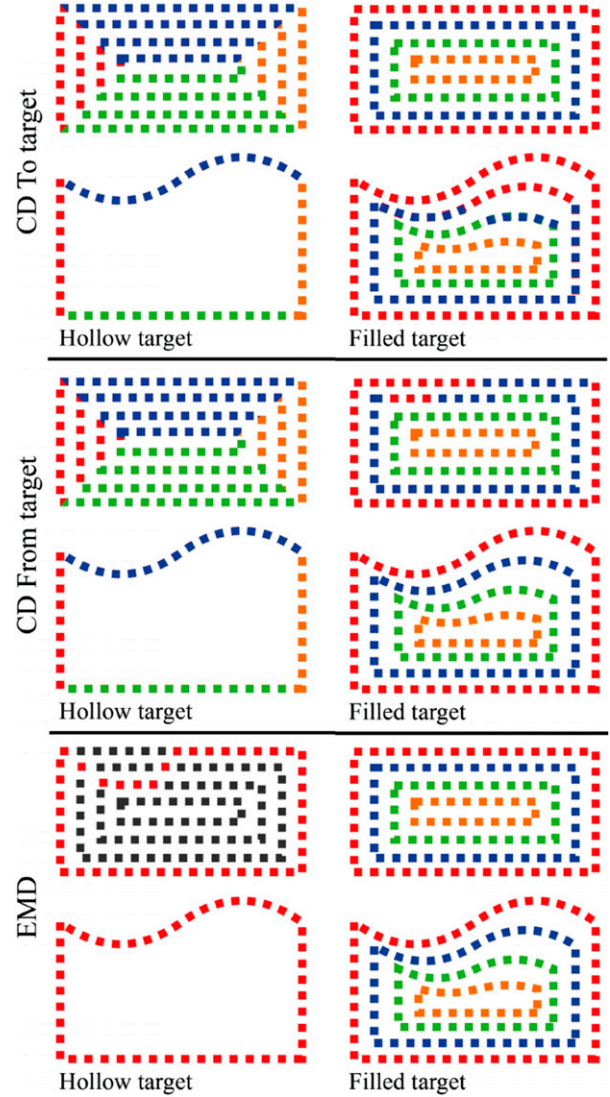


Figure 15. Correspondence examples for the CD and EMD computations with hollow and filled point sets as targets. Colours indicate corresponding points. Note the CD loss is directional.

Figure 14, we can analyse the differences between CD and EMD losses with hollow or filled point sets.

4.7.1. Hollow or filled point sets. For the CD losses, it can be seen by comparing the blue and orange solid lines in Figure 14. The blue solids in the first two rows indicate that

optimising the CD loss with a hollow point cloud target can achieve smaller losses in both CD loss metrics. However, the third to fifth rows in Figure 14 show that optimising CD with a filled particle system target (PRT CD loss) outperforms PCD CD in terms of both EMD loss metrics and the height map distance. On the other hand, optimising the EMD loss with a filled particle system target (red lines) exhibits the same pattern, consistently showing inferior performances in terms of the CD loss metrics while outperforming the PCD EMD losses (green lines) in terms of both EMD loss metrics and the height map distance. Therefore, the results indicate that optimising with a filled particle system target generated from a reconstructed watertight mesh is better than optimising with a hollow point cloud target.

The cause behind it is the inaccurate point-to-point correspondences generated with a hollow point set. As shown by Figure 15, a hollow point set either averages the spatial information of an area of particles over a thin layer of target points in computing the CD loss or ignores the information of a large body of particles within the particle body in computing the EMD loss. On the contrary, with a non-hollow target, the point correspondences connect spatially closer points and do not disregard information within the object body.

4.7.2. *CD or EMD.* Figure 14 shows that either with a hollow or non-hollow target, the EMD losses (green and red lines) consistently outperform the CD losses (blue and orange lines) in terms of the EMD and height map distance metrics while being inferior in terms of the CD metrics. Again, we believe that this is another consequence of the different ways CD and EMD compute the point correspondences.

Figure 15 shows that, with a hollow target, CD distributes the spatial information of the points within over the surface points, which is problematic. Take the shape in Figure 15 as an example, the particles above the target surface will be forced down, while the particles below the target surface will be pushed up, which is not supposed to happen. On the other hand, the EMD loss does not suffer from this effect because it corresponds the particles to the closest points from the target surface only. With a non-hollow target point set, the CD loss produces different correspondences in the two computation directions, which results in a mixture of correct and wrong spatial relations, while the EMD loss is more likely to find the best match spatially due to the one-to-one correspondence generated by a linear assignment algorithm.

In short, this analysis leads to the conclusion that the PRT EMD loss, which computes point correspondences between the filled simulation particle system and a non-hollow target

Table 7. Computational costs (time and GPU memory consumption) of the TaiChi-based differentiable simulator used in the DPSI framework with an Nvidia GeForce RX 4090 GPU. The costs of simulations with three motions and particle densities are reported. cp.: compile-time costs. The first time a program is run takes extra time to compile TaiChi kernels. 100 FB-iterations (min): the runtime of 100 forward and backward computations in minutes.

Motion	Particle density	4×10^7	6×10^7	8×10^7
Poking-1	Number of particles	1946	2808	3926
	CP. Forward time (s)	9.56	9.83	9.23
	Forward time (s)	1.45	1.61	2.17
	Forward GPU (GB)	~1.6	~1.8	~1.9
	CP. Backward time (s)	24.69	24.71	24.42
	Backward time (s)	2.27	2.39	2.56
	Backward GPU (GB)	~1.7	~1.8	~1.9
	100 FB-Iterations (min)	6.71	7.18	8.37
Poking-shifting-1	Number of particles	3880	5836	7738
	CP. Forward time (s)	11.09	14.63	23.97
	Forward time (s)	2.87	7.06	16.78
	Forward GPU (GB)	~1.9	~2.2	~2.5
	CP. Backward time (s)	26.15	27.67	27.34
	Backward time (s)	4.31	5.42	5.79
	Backward GPU (GB)	~1.9	~2.2	~2.5
	100 FB-iterations (min)	12.47	21.30	38.10
Flattening	Number of particles	3059	4688	6271
	CP. Forward time (s)	14.02	15.19	15.94
	Forward time (s)	6.08	6.54	7.88
	Forward GPU (GB)	~1.8	~1.9	~2.1
	CP. Backward time (s)	31.95	34.83	34.93
	Backward time (s)	10.14	12.29	13.03
	Backward GPU (GB)	~1.8	~1.9	~2.1
	100 FB-iterations (min)	27.53	31.90	35.35

point cloud, is the most ideal choice for capturing the correct spatial differences between point sets.

4.8. Remark on computation cost

In this subsection, we report the running time and GPU memory needed for forward and backward computations for the DPSI framework with different particle densities and motions. The statistics are summarised in Table 7, which were obtained using an Nvidia GeForce RTX 4090 GPU.

It can be seen that the computation time and memory increase as the number of simulation particles and the length of motion increase. Because of the pre-compilation of kernels when they are first called in a TaiChi programme, the first forward and backward passes always take much longer and the later repetitive calls are much faster, as denoted by ‘CP. Forward time (s)’.

With particle numbers around 2000 to 4000, a DPSI optimisation run that consists of 100 iterations of forward-backward computations using 1 datapoint can be done in about 10 min. As shown by Figures 4 and 5, most runs converge within 50 iterations, which means that a plausible solution can be found by DPSI in about 5 min. With a good initial guess of the parameters, as shown by the dotted lines in Figure 14, this runtime can be further reduced to around 2 min. In addition, the memory consumption for these simulations is within 3 GB, which permits implementations on portable and small GPUs. With no more than 2 min for the 3D reconstruction pipeline, a parameter identification process could take less than 5 min to complete. These statistics confirm that DPSI meets the criteria in computational costs for practical deployment. It also suggests again that it is practical to perform parameter identification with DPSI at any time without the need to rely on previously found parameters or slow manual measuring techniques.

5. Conclusion

This work addresses the important problem of closing the gap between simulated and real-world manipulations of elastoplastic objects. In particular, we propose a differentiable physics-based system identification framework (DPSI) that can identify physics parameters through gradient descent algorithms. Substantial experiments demonstrate that the proposed framework can identify parameters that reproduce quantitatively and qualitatively realistic elastoplastic object manipulation dynamics in the presence of real-world perceptual and data collection challenges. These challenges include (1) simple and short motions, (2) incomplete trajectories, (3) object occlusions, (4) point cloud noises and (5) small data. The proposed framework is the first example of system identification with differentiable physics-based, particle-based simulation for robotic volumetric elastoplastic object manipulation. It serves as the foundation for faster and more accurate real-world deployments of deformable object manipulation.

In addition, with the use of physics-based dynamics models, the identified parameters are physically meaningful. The experiments reveal that different parameter values found through optimisation can be interpreted in a way that is aligned with our understanding and intuitions about real-world physics. Therefore, the DPSI framework not only gives users confidence in the simulation controlled by the physics models and these parameters but also provides users with intuitive angles to identify the limitations of the reconstructed manipulation dynamics.

5.1. Limitations and future research

Several limitations can be observed from the experiments conducted in this work. The most obvious would be the under-representativeness of the physics models that describe the underlying manipulation dynamics. As discussed in the Results section, there is a lack of elastic returning effect at the areas of contact with sharp edges of objects, the insufficient spread of force impact that causes the deformation to stay near the local contact area, and the artefacts of floating displaced particles. Though these inaccuracies may be negligible for coarse manipulations that appear more often in our daily lives, they would lead to unacceptable solutions for high-precision manipulation tasks such as surgery and soft object assembly. In addition, various approximations in the MLS-MPM and programming implementations also contribute to the simulation inaccuracy. As such, future research is needed to develop more accurate physics models, numerical approximations, and coding implementations to simulate object deformations.

Secondly, an important assumption of this work is that the deformation behaviour of the target material in need of parameter identification can be simulated by the selected elastoplasticity model. Thus, it should be noted that when the target material’s behaviour is largely underrepresented by the selected physics model, DPSI would not be able to produce very realistic simulations. In other words, the manual selection of the elastoplasticity model is of high importance for the proposed DPSI framework. For example, our results with soil reveal that DPSI found parameters to make simulations appear as close to the real materials as possible, but the use of the von Mises plasticity model will make it difficult to recover parameter values that accurately reflect realistic soil or foam behaviours, which are better captured by the Drucker-Prager model.

Thirdly, an important lesson learnt from the results is that increased complexity of motion and contact mode leads to fewer local minima and better system identification accuracy. The number of local minima indicates that the collected data is insufficient to fully induce the correct values of the concerned parameters. Thus, in the future, we envision a better framework that incorporates learning-based approaches to allow the automatic selection of diverse tools and interaction motions to achieve more efficient and accurate parameter identification.

Another limitation comes from the means of capturing real-world observations. In this work, we employ the multi-view fusion point clouds as the observation space, which suffers from noise and the inaccurate estimate of the camera extrinsic matrix. In the future, new methods should be developed to reduce point cloud noises and improve the precision of camera calibration. Lastly, the simulated scene in this work is limited to an end-effector, a table surface, and a target elastoplastic object. Efforts are needed to extend the framework to support more realistic and complex contact dynamics. More experiments are also needed to examine the feasibility of large-scale simulation in terms of efficiency and accuracy.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was supported by the Engineering and Physical Sciences Research Council (grant No. EP/X018962/1). The authors would like to thank Samuel Moeller and Dr Prasad Rayamane for their help with 3D printing.

ORCID iD

Xintong Yang  <https://orcid.org/0000-0002-7612-614X>

Supplemental Material

Supplemental material for this article is available online.

Notes

1. Visit our project page at: <https://ianyanchina.github.io/SI4RP-data/>.
2. Modelling Clay (<https://amzn.eu/d/fK3HIGa>). Cloud Slime Kit (<https://amzn.eu/d/9BYM1kx>).

References

- Antonova R, Yang J, Jatavallabhula KM, et al. (2023) Rethinking optimization with differentiable simulation from a global perspective. In: *Conference on robot learning*, Atlanta, GA: PMLR. pp. 6276–9286, Nov 2023.
- Arnavaz K, Nielsen MK, Kry PG, et al. (2023) Differentiable depth for real2sim calibration of soft body simulations. *Computer Graphics Forum* 42: 277–289. Wiley Online Library.
- Arriola-Rios VE, Guler P, Ficuciello F, et al. (2020) Modeling of deformable objects for robotic manipulation: a tutorial and review. *Frontiers in Robotics and AI* 7: 82.
- Baker K (2005) Singular value decomposition tutorial. *The Ohio State University* 24: 22.
- Bernardini F, Mittleman J, Rushmeier H, et al. (1999) The ball-pivoting algorithm for surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics* 5(4): 349–359.
- Caporali A, Kicki P, Galassi K, et al. (2024) Deformable linear objects manipulation with online model parameters estimation. *IEEE Robotics and Automation Letters* 9: 2598–2605.
- Chen S, Werling K, Wu A, et al. (2022) Real-time model predictive control and system identification using differentiable simulation. *IEEE Robotics and Automation Letters* 8(1): 312–319.
- Cherubini A, Ortenzi V, Cosgun A, et al. (2020) Model-free vision-based shaping of deformable plastic materials. *The International Journal of Robotics Research* 39(14): 1739–1759.
- Collins J, Chand S, Vanderkop A, et al. (2021) A review of physics simulators for robotic applications. *IEEE Access* 9: 51416–51431.
- Featherstone R (2014) *Rigid Body Dynamics Algorithms*. Berlin: Springer.
- Gao M, Tampubolon AP, Jiang C, et al. (2017) An adaptive generalized interpolation material point method for simulating elastoplastic materials. *ACM Transactions on Graphics* 36(6): 1–12.
- Görner M, Haschke R, Ritter H, et al. (2019) Moveit! task constructor for task-level motion planning. In: 2019 International Conference on Robotics and Automation (ICRA), Montreal, QC, 20–24 May 2019, pp. 190–196. IEEE.
- Güler P, Pieropan A, Ishikawa M, et al. (2017) Estimating deformability of objects using meshless shape matching. In: 2017 IEEE/RSJ international conference on Intelligent Robots and Systems (IROS), Vancouver, BC, 24–28 September 2017, pp. 5941–5948. IEEE.
- Heiden E, Millard D, Coumans E, et al. (2021) Neursim: augmenting differentiable simulators with neural networks. In: 2021 IEEE International Conference on Robotics and Automation (ICRA), Xi'an, 30 May 2021–05 June 2021, pp. 9474–9481. IEEE.
- Heiden E, Denniston CE, Millard D, et al. (2022) Probabilistic inference of simulation parameters via parallel differentiable simulation. In: 2022 International Conference on Robotics and Automation (ICRA), Philadelphia, PA, 23–27 May 2022, pp. 3638–3645. IEEE.
- Hu Y, Anderson L and Li TM, et al. (2020) DiffTaichi: differentiable programming for physical simulation. In: The Eighth International Conference on Learning Representations (ICLR), Virtual conference, 26 Apr.-1 May.
- Hu Y, Fang Y, Ge Z, et al. (2018) A moving least squares material point method with displacement discontinuity and two-way rigid body coupling. *ACM Transactions on Graphics* 37(4): 1–14.
- Hu Y, Li TM, Anderson L, et al. (2019) Taichi: a language for high-performance computation on spatially sparse data structures. *ACM Transactions on Graphics* 38(6): 201.
- Huang Z, Hu Y, Du T, et al. (2021) Plasticinelab: a soft-body manipulation benchmark with differentiable physics. In: *The Ninth International Conference on Learning Representations (ICLR)*, Virtual conference, 3-7 May 2021.
- Jaques M, Asenov M, Burke M, et al. (2022) Vision-based system identification and 3d keypoint discovery using dynamics constraints. In: proceedings of The 4th annual learning for dynamics and control conference (eds Firoozi R, Mehr N, Yel E, et al.), Vol. 168, pp. 316–329. Westminster: PMLR.

- Jiang C, Schroeder C, Selle A, et al. (2015) The affine particle-in-cell method. *ACM Transactions on Graphics* 34(4): 1–10.
- Jiang C, Schroeder C, Teran J, et al. (2016) The material point method for simulating continuum materials. In: *ACM SIGGRAPH 2016 Courses*. New York, NY: ACM, 1–52.
- Jones RM (2009) Deformation theory of plasticity. *Bull Ridge Corporation*.
- Kaneko T (2024) *Improving physics-augmented continuum neural radiance field-based geometry-agnostic system identification with Lagrangian particle optimization*. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, Seattle WA, USA, pp. 175470-215480 June, 2024.
- Kingma DP and Ba J (2015) Adam: a method for stochastic optimization. In: *International conference on learning representations*, Vancouver, BC, Canada, 7-9 May 2015.
- Kroemer O, Niekum S and Konidaris G (2021) A review of robot learning for manipulation: challenges, representations, and algorithms. *Journal of machine learning research* 22(30): 1–82.
- Kuka (2024) Kuka lbr iiwa. <https://www.kuka.com/en-de/products/robot-systems/industrial-robots/lbr-iiwa> (accessed 15 January 2024).
- Latombe JC (2012) *Robot motion planning*. Berlin: Springer Science & Business Media, Vol. 124.
- Li X, Qiao YL, Chen PY, et al. (2023) PAC-neRF: physics augmented continuum neural radiance fields for geometry-agnostic system identification. In: *The eleventh international conference on learning representations, Kigali Rwanda*, 1-5 May 2023. <https://openreview.net/forum?id=tVkrbkz42vc>
- Lin X, Huang Z, Li Y, et al. (2022) Diffskill: skill abstraction from differentiable physics for deformable object manipulations with tools. In: ICLR, Virtual conference, 25-29 Apr. 2022.
- Lloyd B, Székely G and Harders M (2007) Identification of spring parameters for deformable object simulation. *IEEE Transactions on Visualization and Computer Graphics* 13(5): 1081–1094.
- McConachie D, Dobson A, Ruan M, et al. (2020) Manipulating deformable objects by interleaving prediction, planning, and control. *The International Journal of Robotics Research* 39(8): 957–982.
- Murthy JK, Macklin M, Golemo F, et al. (2020) gradsim: differentiable simulation for system identification and visuomotor control. In: The Eighth International conference on learning representations, Virtual conference, 26 Apr. - 1 May 2020.
- Navarro-Alarcon D, Yip HM, Wang Z, et al. (2016) Automatic 3-d manipulation of soft objects by robotic arms with an adaptive deformation model. *IEEE Transactions on Robotics* 32(2): 429–441.
- Park JJ, Florence P, Straub J, et al. (2019) Deepsdf: learning continuous signed distance functions for shape representation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Long Beach, CA, 15–20 June 2019, pp. 165–174.
- Quigley M, Conley K, Gerkey B, et al. (2009) Ros: an open-source robot operating system. In: *ICRA Workshop on Open Source Software*. Japan: Kobe, Vol. 3, p. 5.
- Rehman Z, Khalid U, Farooq K, et al. (2018) On yield stress of compacted clays. *International Journal of Geo-Engineering* 9: 1–16.
- Schoenholz S and Cubuk ED (2020) Jax md: a framework for differentiable physics. *Advances in Neural Information Processing Systems* 33: 11428–11441.
- Shen B, Jiang Z, Choy C, et al. (2024) Action-conditional implicit visual dynamics for deformable object manipulation. *The International Journal of Robotics Research* 43(4): 437–455.
- Shetab-Bushehri M, Aranda M, Mezouar Y, et al. (2023) Lattice-based shape tracking and servoing of elastic objects. *IEEE Transactions on Robotics*.
- Shi H, Xu H, Clarke S, et al. (2023) Robocook: long-horizon elasto-plastic object manipulation with diverse tools. In: CoRL, Atlanta, GA, 6-9 Nov. 2023.
- Shi H, Xu H, Huang Z, et al. (2024) Robocraft: learning to see, simulate, and shape elasto-plastic objects in 3d with graph networks. *The International Journal of Robotics Research* 43(4): 533–549.
- Stomakhin A, Howes R, Schroeder CA, et al. (2012) Energetically consistent invertible elasticity. In: *Symposium on computer animation, Lausanne Switzerland*, 29-31 July 2012, Vol. 1.
- Stomakhin A, Schroeder C, Chai L, et al. (2013) A material point method for snow simulation. *ACM Transactions on Graphics* 32(4): 1–10.
- StructX (2014–2024) Structx. <https://structx.com/home.html>
- Sucan IA, Moll M and Kavraki LE (2012) The open motion planning library. *IEEE Robotics & Automation Magazine* 19(4): 72–82.
- Sun C, Sun M and Chen HT (2022) Direct voxel grid optimization: super-fast convergence for radiance fields reconstruction. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, New Orleans, LA, 18–24 June 2022, pp. 5459–5469.
- Sundaresan P, Antonova R and Bohgl J (2022) Diffcloud: real-to-sim from point clouds with differentiable simulation and rendering of deformable objects. In: 2022 IEEE/RSJ international conference on Intelligent Robots and Systems (IROS), Kyoto, Japan, 23–27 October 2022, pp. 10828–10835. IEEE.
- Waheed MQ and Asmael NM (2023) Evaluation of elasticity modulus of clayey soil from undrained shear strength. In: E3S Web of Conferences. Les Ulis: EDP Sciences, Vol. 427, p. 01028.
- Xian Z, Zhu B, Xu Z, et al. (2023) Fluidlab: a differentiable environment for benchmarking complex fluid manipulation. In: *The eighth international conference on learning representations (ICLR), Kigali Rwanda*, 1-5 May 2023.
- Yang Y, Stork JA and Stoyanov T (2022) Learning differentiable dynamics models for shape control of deformable linear objects. *Robotics and Autonomous Systems* 158: 104258.
- Yang B, Lu B, Chen W, et al. (2023) Model-free 3-d shape control of deformable objects using novel features based on modal analysis. *IEEE Transactions on Robotics* 39(4): 3134–3153.
- Yin H, Varava A and Kragic D (2021) Modeling, learning, perception, and control methods for deformable object manipulation. *Science Robotics* 6(54): eabd8803.
- Zivid (2024) Zivid one plus 3d cameras. <https://www.zivid.com/zivid-one-plus> (accessed 15 January 2024).

Appendix

A. Optimisation results with height map loss

As mentioned in the main text, optimising the parameters using the height map loss results in unsatisfactory performances. Here, we compare the optimisation and

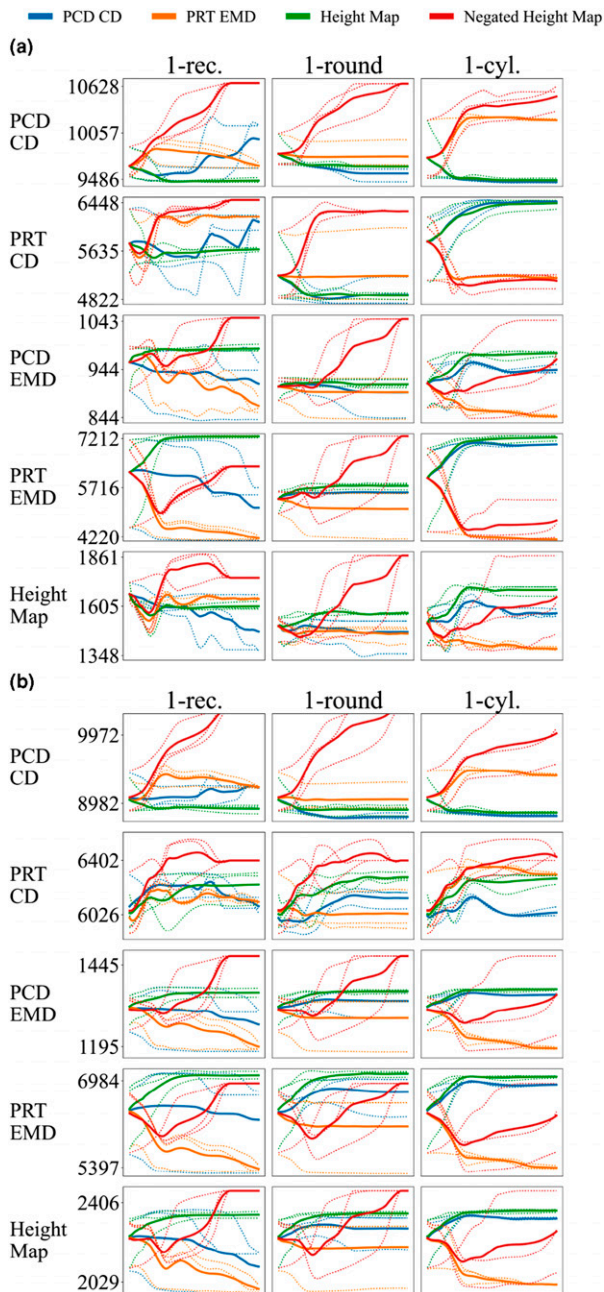


Figure 16. Optimisation and validation losses over 100 gradient updates at level-1 contact complexity from optimising the PRT CD, PRT EMD, Height Map and Negated Height Map loss functions. Each row shows the changes of an evaluation metric or a parameter, denoted on the left. In each figure, different colours indicate the results of minimising a different loss function, as labelled by the legend on the top. Each colour has three dotted lines corresponding to the results of three random seeds and a solid line corresponding to their means. (a) Optimisation losses. (b) Validation losses.

validation losses of minimising the height map loss with those of minimising the PCD CD and PRT EMD losses in Figure 16.

The green curves show that minimising the height map loss results in an increase of the height map loss with both optimisation and validation datasets, which is counterintuitive. To verify that the counter effect is not due to mistaken loss computations, we further optimise the negated height map loss (by effectively using the negated gradients from the height map loss), which resulted in worse performances (red curves). In addition, the height map loss is found to perform worse than the PCD CD loss, which was found to be the worst in all point-based loss functions (see the blue curves in Figures 4 and 5). Therefore, we discarded the results from the main article to reduce complications during the result analysis.

The possible cause of this phenomenon could be the loss of spatial information and geometric precision during height map computations. While height maps appear to be visually more intuitive to humans, it is computed by collecting the highest z values of the points within an area defined by the height map resolution. As such, the gradients from the height map loss are only relevant to the z values of the highest points in each cell, having no connections with the points that are lower in each cell nor with the horizontal aspect of the selected points.

B. Skill planning with the MPM simulator

To demonstrate the feasibility of planning with the MPM simulator, we create a simple greedy exhaustive search algorithm that selects 10 actions to manipulate a piece of clay to match a target height map using the cylinder end-effector. The process is summarised in Algorithm 1.

In short, the algorithm starts with an initial object configuration in the simulator and performs exhaustive searches to find eight consecutive actions that result in the lowest height map distance from the target height map. An action effectively applies one of the two poking-shifting motions at one of the nine locations surrounding the object. As shown by Figure 17, the neighbouring locations are 0.03 m away from each other in the x and y directions. The height from which the skill motion starts is set to the highest z value of the object before each action. For each action, the algorithm queries the simulator for the height map loss values of the manipulation results of applying the two poking-shifting motions at the nine locations (hence 18 forward simulations). An empty action that does not move the end-effector is also considered in each search. The action with the lowest loss is selected and the algorithm proceeds to the next action search with a new object state from the last action. The algorithm returns a plan of eight actions.

Figures 18 and 19 visualise two plans returned by the algorithm that successfully manipulate the clays to a shape close to their respective target height maps. These

examples showcase the feasibility of planning with the MPM simulator. Each plan takes roughly 7.5 min to finish on a 4090 GPU, where each iteration (search over 19 actions) takes roughly 0.9 min. This is not yet real-time but promising considering the high dimensionality of deformable object states.

Algorithm 1: Greedy exhaustive search

```

Initialise object particle system
Initialise MPM simulator: ENV
Load target height map
Load DPSI-found parameters: ENV.set_parameters( $\theta^{sim}$ )
Effector_init_pos  $\leftarrow$  ENV.get_object_centre()
New_state  $\leftarrow$  ENV.get_state()
Plan = []
For n = 0 to Number_of_actions:
  ▷ Evaluate all skills and start locations
  Min_loss  $\leftarrow$   $\infty$ 
  For i = 1 to 2:
    For j = -1 to 1:
      For k = -1 to 1:
        Load skill trajectory: skilli
        Effector_init_pos' = Effector_init_pos.copy()
        Effector_init_pos'.x + = 0.3 × j
        Effector_init_pos'.y + = 0.3 × k
        Effector_init_pos'.z = ENV.get_object_height()
        ENV.set_state(New_state)
        ENV.set_agent_pos(Effector_init_pos')
        ENV.simulate(skilli)
        Loss = ENV.get_loss()
        If Loss < Min_loss:
          Min_loss  $\leftarrow$  Loss
          Best_actionn  $\leftarrow$  (i, j, k)
          Final_state  $\leftarrow$  ENV.get_state()
  ▷ Evaluate empty action
  ENV.set_state(New_state)
  ENV.set_agent_pos(Effector_init_pos')
  ENV.simulate(empty_action)
  Loss = ENV.get_loss()
  If Loss < Min_loss:
    Min_loss  $\leftarrow$  Loss
    Best_actionn  $\leftarrow$  (empty_action)
    Final_state  $\leftarrow$  ENV.get_state()
  Plan.append(Best_actionn)
  New_state  $\leftarrow$  Final_state
Return Plan

```

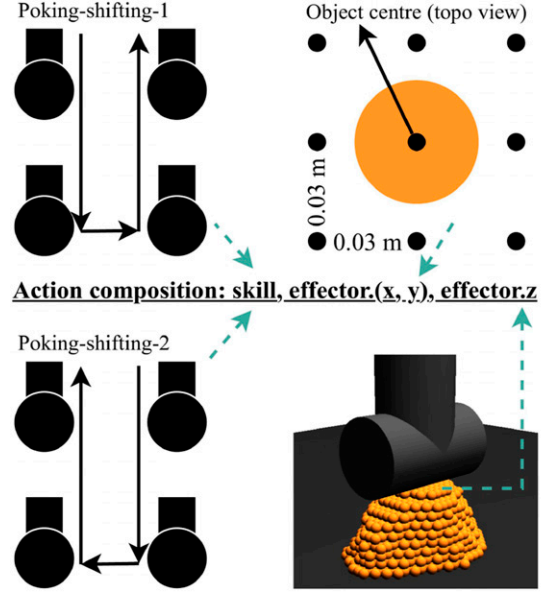


Figure 17. Action composition for the skill planning algorithm with the cylinder end-effector. An action applies one of the two poking-shifting motions at one of the nine neighbouring locations surrounding the object from the highest z value of the object.

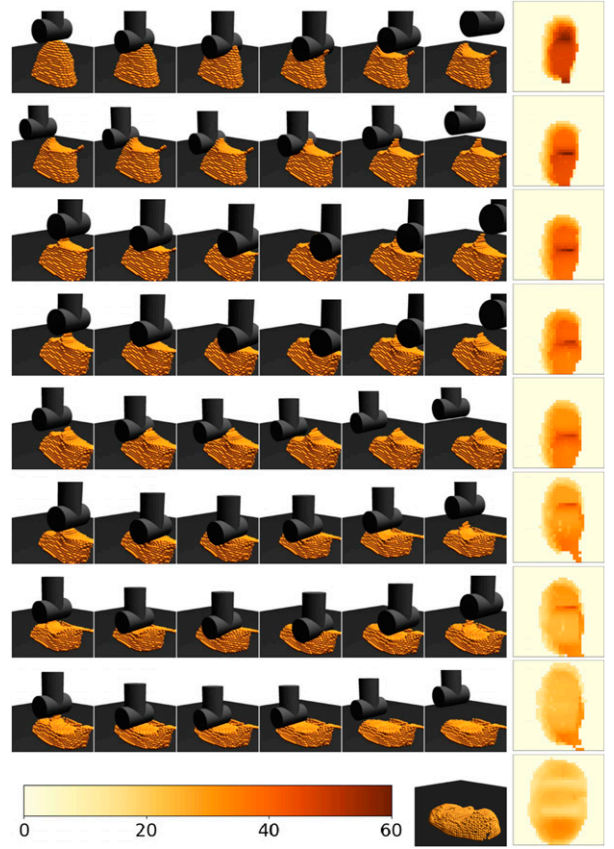


Figure 18. Visualising a plan returned by the exhaustive search algorithm that takes eight actions to flatten a piece of clay to match a target height map (bottom right). All height maps are plotted under the same colour/value distribution shown by the colour bar.

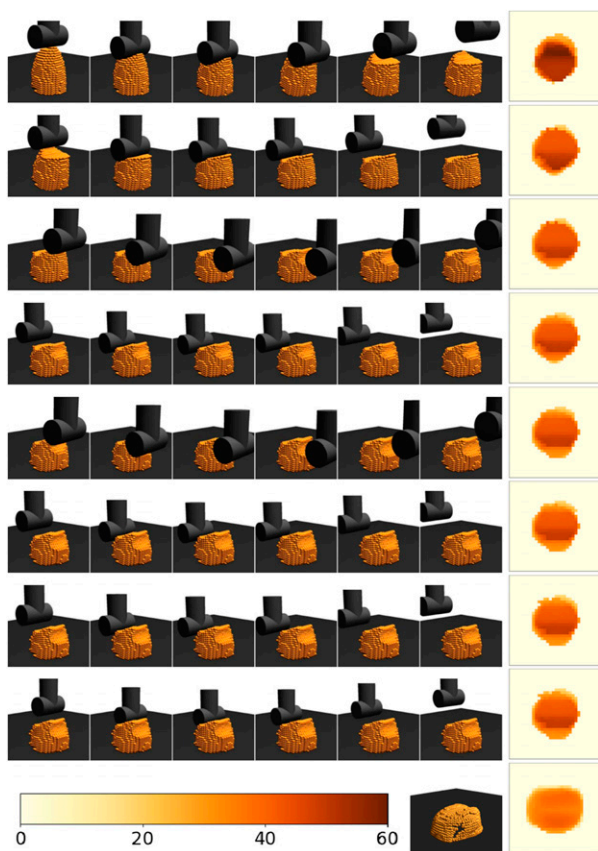


Figure 19. Visualising a plan returned by the exhaustive search algorithm that takes eight actions to flatten a piece of clay to match a target height map (bottom right). All height maps are plotted under the same colour/value distribution shown by the colour bar.