LATENT TASK-SPECIFIC GRAPH NETWORK SIMULA TORS

Anonymous authors

Paper under double-blind review

ABSTRACT

Simulating object deformations is a critical challenge in many scientific domains, with applications ranging from robotics to materials science. Learned Graph Network Simulators (GNSs) are an efficient alternative to traditional mesh-based physics simulators. Their speed and inherent differentiability make them particularly well-suited for inverse design problems such as process optimization. However, these applications typically offer limited available data, making GNSs difficult to use in real-world scenarios. We frame mesh-based simulation as a meta-learning problem and apply conditional Neural Processes to adapt to new simulation scenarios with little data. In addition, we address the problem of error accumulation common in previous step-based methods by combining this approach with movement primitives, allowing efficient predictions of full trajectories. We validate the effectiveness of our approach, called Movement-primitive Meta-MeshGraphNet (M3GN), through a variety of experiments, outperforming state-of-the-art step-based baseline GNSs and step-based meta-learning methods.

024 025

026 027

003 004

010 011

012

013

014

015

016

017

018

019

021

1 INTRODUCTION

028 The simulation of complex physical systems is of paramount importance in a wide variety of engineering disciplines, including structural mechanics (Yazid et al., 2009; Zienkiewicz & Taylor, 029 2005; Stanova et al., 2015), fluid dynamics (Chung, 1978; Zienkiewicz et al., 2013; Connor & Brebbia, 2013), and electromagnetism (Jin, 2015; Polycarpou, 2022; Reddy, 1994). In particular, the simulation 031 of object deformations under external forces is crucial for, e.g., robotic applications (Scheikl et al., 032 2022; Wang & Zhu, 2023; Linkerhägner et al., 2023). Mesh-based simulations are appealing for 033 such problems due to the computational efficiency and accuracy of the underlying finite element 034 method (Brenner & Scott, 2008; Reddy, 2019). However, the diversity of the problems to be 035 modeled usually necessitates the development of task-specific simulators to accurately capture the relevant physical quantities (Reddy & Gartling, 2010). Such specialized simulators can be slow and 037 cumbersome to use, especially for large-scale simulations (Paszynski, 2016; Hughes et al., 2005). 038

Thus, data-driven models trained on reference simulations have gained attention as an appealing alternative (Guo et al., 2016; Da Wang et al., 2021; Li et al., 2022). Among them, general-purpose Graph Network Simulators (GNSs) have recently become increasingly popular (Battaglia et al., 2018; Pfaff et al., 2021; Allen et al., 2022b; 2023; Linkerhägner et al., 2023). GNSs encode the simulated system as a graph of interacting entities whose dynamics are predicted using Graph Neural Networks (GNNs) (Bronstein et al., 2021). These models are one to two orders of magnitude faster than classical simulators (Pfaff et al., 2021) while being fully differentiable, which makes them highly effective for, e.g., inverse design problems (Allen et al., 2022b; Xu et al., 2021).

GNSs are typically trained through simple next-step supervision (Battaglia et al., 2018; Pfaff et al., 2021; Allen et al., 2023). During inference, entire trajectories are simulated by iteratively predicting per-node dynamics from an initial system state in an autoregressive manner. This approach is prone to error accumulation over time, especially as the input distribution diverges from the training set (Brandstetter et al., 2022; Han et al., 2022). To mitigate this issue, existing approaches add noise to the input during training and predict the dynamics based on the original inputs, thus following an implicit de-noising objective (Pfaff et al., 2021; Brandstetter et al., 2022). While this significantly improves simulation stability over longer rollouts, the auto-regressive inference of next-step GNSs still propagates and accumulates errors made in earlier prediction steps. These methods also struggle



Figure 1: Final M3GN (Movement-Primitive Meta-MeshGraphNet) simulation steps for different
evaluation tasks. From left to right: a planar sheet bending under two orthogonal forces, a falling
collider deforming a 2D plate, a surgical tool dragging tissue, and a falling teddy bear. All visualizations present the **predicted mesh** alongside a reference **wireframe** of the ground-truth simulation.
M3GN takes the deformable object positions from a few initial time steps and applies meta-learning
techniques to infer physical properties, such as Poisson's ratio or Young's modulus. With this latent
task description, it predicts the remaining simulation steps using per-node movement primitives.

- with partially known initial system states (Linkerhägner et al., 2023), which are common in, e.g.,
 robotic planning (Antonova et al., 2022). Extensions that address this uncertainty (Linkerhägner
 et al., 2023) require a consistent stream of auxiliary information, such as point cloud observations,
 throughout the simulation. Furthermore, GNSs typically require large amounts of training data, which
 prevents their application in real-world scenarios where data is scarce, highlighting the need for
 efficient adaptation techniques to novel tasks.
- To address these limitations, we reformulate learned mesh-based simulation as a trajectory-level 079 meta-learning problem. Here, the initial mesh states of a trajectory serve as a context set on which to condition future predictions. We employ Conditional Neural Processs (CNPs) (Garnelo 081 et al., 2018a) to aggregate these context sets and the dynamics inferred from them into a latent descriptor, which is then used to predict the rest of the trajectory. This method enables efficient 083 training and rapid adaptation to trajectory-specific simulation parameters, such as unknown object 084 material properties, during inference. In addition, we mitigate the problem of error accumulation 085 by directly predicting full simulation trajectories of the mesh nodes instead of iteratively predicting their next-step dynamics. To this end, we represent the simulation using node-level Probabilistic 087 Dynamic Movement Primitives (ProDMPs) (Schaal, 2006; Paraschos et al., 2013; Li et al., 2023), 088 which allows for an efficient encoding of higher-order dynamics at arbitrary temporal resolution. 089 The resulting method, called Movement-Primitive Meta-MeshGraphNet (M3GN), allows efficient generation of context-dependent simulation trajectories that accurately infer and integrate unknown 090 system properties. Figure 1 shows examples for different tasks, while Figure 2 provides an overview 091 of our approach. 092
- To validate the effectiveness of M3GN, we introduce three novel tasks based on challenging deformable object simulations with varying object materials. These include a planar plate bending under stress, and a variety of falling objects that collide with the ground. In addition to these new tasks, we evaluate M3GN on an existing suite of experiments (Linkerhägner et al., 2023).Our results show that our method provides superior simulation accuracy compared to several variants of MeshGraphNet (MGN) Pfaff et al. (2021), the state-of-the-art GNS.¹ Further, the ProDMPs trajectory representation of M3GN reduces the number of required model calls, improving inference runtime by up to 32 times compared to MGN.
- In summary, we (i) propose M3GN, a novel GNS that combines meta-learning and movement
 primitives to predict node-level simulation trajectories based on initial system state contexts;
 (ii) introduce two challenging deformation prediction tasks involving varying object materials;
 (iii) evaluate and compare our method to state-of-the-art GNSs, demonstrating superior prediction
 performance.
- 106 107

¹Code is provided in the supplement. Here, the reviewers can also find a video showing the main results including HD renders of M3GN predictions and comparisons to existing baselines.

Context Data

Shared GNN

Context Aggregation

Encod

Shared GNN

Encoder

108 109 110

Time

111 112 113

114 115 116

117 118

119 120

121

123

124

Figure 2: M3GN architecture schematic. Given a context set of initial system states, we calculate node-level latent features for every pair of states using a shared GNN encoder. These feature sets are then aggregated, yielding a node-level latent task description z_v . We concatenate this description with the last system state to predict ProDMP weights that are used to compute per-node trajectories.

Latent Task

Description z_1

Prediction

Simulato

GNN

ProDMP

Weights

 w_v

ProDMP

Traiectory

Generator

125 126 127

128

2 RELATED WORK

Shared GNN

Encoder

129 Graph Network Simulators. Deep neural networks for physical simulations can provide signifi-130 cant speedups over traditional simulators while being fully differentiable (Pfaff et al., 2021; Allen 131 et al., 2022a), making them a natural choice for applications like model-based Reinforcement Learning (Mora et al., 2021) and Inverse Design problems (Baqué et al., 2018; Durasov et al., 2021; 132 Allen et al., 2022a). A popular class of learned neural simulators are Graph Network Simulators 133 (GNSs) (Battaglia et al., 2016; Sanchez-Gonzalez et al., 2020). GNSs utilize Message Passing 134 Networks (MPNs), a special type of GNN (Scarselli et al., 2009; Bronstein et al., 2021) that repre-135 sentationally encompasses the function class of many classical solvers (Brandstetter et al., 2022). 136 GNS handle physical data by modeling arbitrary entities and their relations as a graph. Applications 137 of GNSs include particle-based simulations (Li et al., 2019; Sanchez-Gonzalez et al., 2020; Whitney 138 et al., 2023), atomic force prediction (Hu et al., 2021), and fluid dynamic problems (Brandstetter et al., 139 2022). These models have additionally been applied to the mesh-based prediction of deformable 140 objects (Pfaff et al., 2021; Weng et al., 2021; Han et al., 2022; Fortunato et al., 2022; Linkerhägner 141 et al., 2023). Recent extensions handle rigid objects (Allen et al., 2022b; 2023; Lopez-Guevara et al., 142 2024) and integrate learned adaptive mesh refinement strategies (Plewa et al., 2005; Freymuth et al., 2023) into the simulator (Wu et al., 2023). Existing work that considers unknown material properties 143 in simulations of deformable objects combines the GNS prediction with point cloud information to 144 improve long-term predictions Linkerhägner et al. (2023). This method requires a constant stream of 145 point clouds to ground the simulation in, but can not aggregate this information into a description of 146 the material properties. Additionally, the DEL method (Wang et al., 2024) integrates physical priors 147 from the Discrete Element Analysis (DEA) framework with learnable graph kernels, addressing 148 the challenges of simulating 3D particle dynamics from 2D images. In the context of larger-scale 149 simulations, foundation models are gaining traction in neural simulation tasks, as exemplified by 150 Aurora (Bodnar et al., 2024), a large-scale model trained on extensive climate data. While Aurora 151 demonstrates impressive performance on atmospheric predictions, including global air pollution 152 and weather forecasts, it requires significantly more data for fine-tuning compared to our approach, which focuses on efficient adaptation with fewer data points. Notably, all previously mentioned 153 GNSs predict system dynamics iteratively from a given state, whereas we directly estimate entire 154 trajectories, improving rollout stability and reducing function calls. Related to our approach is the 155 Equivariant Graph Neural Operator (EGNO) (Xu et al., 2024), which also predicts full trajectories 156 using SE(3) equivariance to model 3D dynamics and capture spatial and temporal correlations. In 157 contrast to the related work, which rely on supervised learning or fine-tuning a foundation model, we 158 employ meta-learning for efficient adaptation to new trajectories. 159

Meta-Learning. Meta-learning (Schmidhuber, 1992; Thrun & Pratt, 1998; Vilalta & Drissi, 2005;
 Hospedales et al., 2022) extracts inductive biases from a training set of related tasks in order to increase data efficiency on unseen tasks drawn from the same task distribution. In contrast to other

162 multi-task learning methods, such as transfer learning (Krizhevsky et al., 2012; Golovin et al., 2017; 163 Zhuang et al., 2020), which merely fine-tune or combine standard single-task models, meta-learning 164 makes the multi-task setting explicit in the model architecture (Bengio et al., 1991; Ravi & Larochelle, 165 2017; Andrychowicz et al., 2016; Volpp et al., 2019; Santoro et al., 2016; Snell et al., 2017). This 166 explicit architecture allows the resulting meta-models to learn how to learn new tasks from a small number of example contexts. A popular variant is Model-Agnostic Meta-Learning (MAML) (Finn 167 et al., 2017; Grant et al., 2018; Finn et al., 2018; Kim et al., 2018), which employs standard single-task 168 models and formulates a multi-task optimization procedure. Neural Processes (NPs) (Garnelo et al., 2018a;b; Kim et al., 2019; Gordon et al., 2019; Louizos et al., 2019; Volpp et al., 2021; 2023) instead 170 build on a multi-task model architecture (Heskes, 2000; Bakker & Heskes, 2003) but employ standard 171 gradient based optimization algorithms (Kingma & Ba, 2015; Kingma & Welling, 2014; Rezende 172 et al., 2014; Zaheer et al., 2017). Here, we use Conditional Neural Processes (CNPs) (Garnelo 173 et al., 2018a), which aggregate learned features over a variable-sized context set to yield a latent task 174 description that our downstream GNS is conditioned on. Compared to regular NPs, CNPs assume 175 a deterministic task description, eliminating the need for a distribution over latent variables. This 176 assumption simplifies and accelerates the training process, as our objective is to predict a single 177 precise simulation trajectory from the context set.

178 179

180 181

187 188 189

190

3 MOVEMENT-PRIMITIVE META-MESHGRAPHNETS

In this section, we present the theoretical foundation of the M3GN method, detailing the algorithmic design choices that guided its development. **Graph Network Simulators.** Consider a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{X}_{\mathcal{V}}, \mathbf{X}_{\mathcal{E}})$ with nodes \mathcal{V} , edges \mathcal{E} , and associated vector-valued node and edge features **X**_{\mathcal{V}} and **X**_{\mathcal{E}}. An MPN (Sanchez-Gonzalez et al., 2020; Pfaff et al., 2021) consists of *M* message passing steps, which iteratively update the node and edge features based on the graph topology. Each such step is given as

$$egin{aligned} \mathbf{h}_{e}^{m+1} &= f_{\mathcal{E}}^{m}(\mathbf{h}_{v}^{m},\mathbf{h}_{e}^{m}), \ \mathbf{h}_{v}^{m+1} &= f_{\mathcal{V}}^{m}(\mathbf{h}_{v}^{m},igoplus_{e\in\mathcal{E}_{v}}\mathbf{h}_{e}^{m+1}), \end{aligned}$$

where \mathbf{h}_{v}^{m} and \mathbf{h}_{e}^{m} denote embeddings of the system state per node and edge at message passing iteration m, respectively. $\mathcal{E}_{v} \subset \mathcal{E}$ are the edges connected to v. Further, \bigoplus denotes a permutationinvariant aggregation operation such as the sum, the max, or the mean. The functions $f_{\mathcal{V}}^{m}$ and $f_{\mathcal{E}}^{m}$ are learned Multilayer Perceptrons (MLPs). The network's final output are the node-wise learned representations $\mathbf{h}_{v} := \mathbf{h}_{v}^{M}$ that encode local information of the initial node and edge features.

196 Conventional GNSs encode the state of the simulated system as a graph, feed it through the MPN, 197 and interpret the per-node outputs as velocities or accelerations (Pfaff et al., 2021). These dynamics are used to forward the simulation in time using, e.g., a forward-Euler integrator (Sanchez-Gonzalez 199 et al., 2020). The graph encodes relative distances and velocities between entities instead of absolute 200 ones, as the resulting equivariance to translation improves generalization (Sanchez-Gonzalez et al., 201 2020). GNSs usually minimize a next-step Mean Squared Error (MSE) per node during training, adding carefully tuned implicit denoising strategies (Pfaff et al., 2021; Brandstetter et al., 2022) to 202 stabilize long-term predictions. During inference, they compute trajectories by iteratively predicting 203 and integrating their output in an autoregressive fashion. If some simulated objects, like the col-204 lider, are known, only the remaining nodes are predicted. Our method instead uses a ProDMP to 205 predict a compact representation of a whole trajectory per system node, reducing the effect of error 206 accumulation, similar to temporal bundling (Brandstetter et al., 2022). 207

Probabilistic Dynamic Movement Primitives. Movement Primitives (MPs) (Schaal, 2006; 208 Paraschos et al., 2013) allow for compact and smooth trajectory representations y via a set of 209 basis functions parameterized by a set of weights w. This temporal smoothness is highly beneficial 210 for, e.g., robotic applications (Li et al., 2024; Otto et al., 2022). Recent methods integrate MPs with 211 neural networks to enhance their expressive capabilities (Seker et al., 2019; Bahl et al., 2020; Li 212 et al., 2023). Dynamic Movement Primitives (DMPs) (Schaal, 2006) use a spring-damper dynamical 213 system governed by parameters α and β . To manipulate the trajectory, an external forcing term f is 214 added, before the system converges to a predefined goal g: 215

$$\tau^2 \ddot{y} = \alpha \left(\beta(g - y) - \tau \dot{y}\right) + f(x), \quad f(x) = x \varphi^{\mathsf{T}} w.$$
(1)

225

226

227

228

229 230

231

232

233

234 235

253



Figure 3: Left: M3GN and MGN task setup. Both methods predict mesh positions based on the initial mesh at the anchor time step. M3GN utilizes previous mesh positions and the last step of the collider trajectory for its latent task description, whereas MGN disregards past information and integrates the ground truth collider trajectory into its step-based model. **Right:** Exemplary final simulation steps on the Planar Bending task of M3GN given different context set sizes. A larger context size results in a more accurate **prediction** of the **ground truth** simulation.

Here, τ influences execution speed, while f depends on the basis functions φ in force-space, the weights w and the exponential decaying phase x. Solving this equation typically is computationally intensive, particularly when the gradient dy/dw is required (Bahl et al., 2020). ProDMPs (Li et al., 2023) instead solve Equation equation 7 with pre-computed basis functions Φ in position-space as

$$y(t) = c_1 y_1(t) + c_2 y_2(t) + \boldsymbol{\Phi}(t)^{\top} \boldsymbol{w}.$$

The term $c_1y_1(t) + c_2y_2(t)$ only depends on the initial conditions $[y(t_0), \dot{y}(t_0)]$. ProDMPs thus generate smooth trajectories at an arbitrary temporal resolution from low-dimensional weights w. They crucially allow for efficient gradient computation, and can respect different initial conditions such as positions or velocities. We provide an extensive mathematical background of ProDMPs in Appendix A.

241 Meta-Learning and Graph Network Simulators. To enable generalization across tasks with 242 varying properties, we frame GNS as a meta-learning problem. In this setup, each task corresponds 243 to a simulation of a deformable object with unknown material properties. The goal is to learn 244 a simulator that can adapt quickly to a specific scenario using a limited amount of context data. 245 Following the notation of Volpp et al. (2021), the meta-dataset $\mathcal{D} = \mathcal{D}_{1:L}$ consists of simulation 246 trajectories $\mathcal{D}_l = \{\mathcal{G}_{l,1} \dots \mathcal{G}_{l,T}\}$ of length T, where T is the trajectory length. Each simulation step 247 $\mathcal{G}_{l,t} = (m_{l,t}, u_{l,t})$ represents a graph capturing both the deformable object mesh $m_{l,t}$ (describing 248 its position and topology) and an optional rigid collider $u_{l,t}$. Physical proximity is used to define graph edges that model interactions between the deformable object and the collider. At test time, 249 the first $T^c \ll T$ simulation frames, $\mathcal{G}_{l,1:T^c}$, are observed as a context set to predict the remaining 250 trajectory. Following prior work (Pfaff et al., 2021), we assume access to the full collider trajectory 251 during prediction, resulting in the complete context set: 252

$$\mathcal{D}_l^c = \{\mathcal{G}_{l,1}, \dots, \mathcal{G}_{l,T^c}\} \cup \{\boldsymbol{u}_{l,T^c+1}, \dots, \boldsymbol{u}_{l,T}\}.$$
(2)

To provide a clear reference point for discussion, we define the *anchor time step* as the final time step T^c of the context set. The corresponding *anchor graph*, \mathcal{G}_{l,T^c} , represents the system's state at this point and serves as the starting state for trajectory prediction by the GNSs. Notably, the anchor graph \mathcal{G}_{l,T^c} alone does not capture the complete system state, as the material properties of the deformable object remain unknown. These properties must be inferred from the prior simulation steps, $\mathcal{G}_{l,1:T^c}$, to enable accurate trajectory predictions. Figure 3 illustrates this setup.

260 Model Architecture. Our model architecture, M3GN, is designed to learn from context data and 261 predict future simulation steps by leveraging a combination of graph network simulation and meta-262 learning techniques. The architecture consists of two parts: the computation of the latent task 263 description from the context data and the actual graph network simulation of future simulation steps. 264 We base our context processing on the Conditional Neural Process (CNP) (Garnelo et al., 2018a), 265 as it efficiently encodes a latent description over tasks given a set of context observations. Omitting 266 the task index l to avoid clutter, CNPs expect a context set $\{(x_1, y_1), \ldots, (x_{N^c}, y_{N^c})\}$ consisting of inputs x_i and corresponding targets y_i . We translate our context set \mathcal{D}^c from Equation 2 to this format 267 by using each graph as an input, and setting its labels as the node-wise velocities. This approach 268 allows the model to focus on dynamics rather than absolute positions, which are more task-specific. 269 Assuming a forward-Euler integration scheme with a time step of 1, we numerically approximate the

velocities as the difference between consecutive simulation steps. The input graph x_i represents the simulation state at time step *i*, including mesh and collider, while y_i encodes the change in positions between consecutive time steps.

273 274

$$\boldsymbol{x}_i = \mathcal{G}_i, \qquad \boldsymbol{y}_i = \operatorname{pos}(\mathcal{G}_{i+1}) - \operatorname{pos}(\mathcal{G}_i)$$

To account for the known collider trajectory, we add its relative position as an additional node feature. Specifically, we include the position of the collider at time step T, $pos(u_T)$, relative to its current position, $pos(u_i)$. Preliminary testing indicated that incorporating the complete future collider trajectory $pos(u_{T^c})$, ... $pos(u_T)$ did not improve the results on our task suite. Given a context set \mathcal{D}^c with anchor time step T^c , this results in $T^c - 1$ tuples (x_i, y_i) . A shared GNN encoder h_{θ} computes node-level latent features

293

294

308 309

$$\boldsymbol{z}_{i,v} = h_{\boldsymbol{\theta}}(\boldsymbol{x}_i, \boldsymbol{y}_i) \in \mathbb{R}^{T^c \times |\mathcal{V}| \times d_z}$$
(3)

for each context time step with feature dimension d_z . We then aggregate over the context set to obtain $z_v = \bigoplus z_{i,v} \in \mathbb{R}^{|\mathcal{V}| \times d_z}$, using $\bigoplus = \max$ as the aggregation operator. Intuitively, z_v is a representation of the task inferred from the context data \mathcal{D}^c , and encodes material properties, future collider movements, and high-level deformations of the simulation. While we use node-level latent features for M3GN, one could additionally aggregate over the nodes to obtain a graph-global task descriptor $z = \bigotimes z_v \in \mathbb{R}^{d_z}$. We explore this choice and different aggregation functions \bigoplus , \bigotimes in Section 4.

290 Once the task descriptor has been computed, it serves as the input to the predictive stage, enabling 291 simulation of future trajectories. We concatenate the latent description z_v with the node features of 292 the anchor graph \mathcal{G}_{T^c} and subsequently use a GNN g_{θ} to predict per-node ProDMP weights

1

$$\boldsymbol{v}_{v} = g_{\boldsymbol{\theta}}(\mathcal{G}_{T^{c}}, \boldsymbol{z}_{v}) \in \mathbb{R}^{|\mathcal{V}| \times d_{w}}.$$
(4)

295 For certain tasks, incorporating the current node velocities as an additional node feature in the 296 simulator GNN can be advantageous. The ProDMP trajectory generator $f(w_v) \in \mathbb{R}^{T \times |\mathcal{V}| \times d_{world}}$ 297 transforms the predicted outputs of the simulator GNN into per-node object trajectories over the 298 entire simulation horizon. This approach can be seen as a form of temporal bundling (Brandstetter 299 et al., 2022), requiring a single function call. In comparison, existing GNS train mostly on next-step dynamics and require one call per step during their auto-regressive inference scheme (Pfaff et al., 300 2021; Allen et al., 2023). The trajectory-level view further allows us to omit noise injection during 301 training, which MGN requires to generalize from learned next-step predictions to multi-step rollouts 302 during inference. We provide a visualization of our model architecture in Figure 2 and refer to 303 Appendix B for further details. 304

305 Meta Training. The goal of meta-learning is to automatically encode inductive biases towards the 306 task distribution extracted from the meta-dataset \mathcal{D} into the task-global parameter θ . To this end, we 307 minimize the negative conditional log probability (Garnelo et al., 2018a)

$$\mathcal{L}(\boldsymbol{\theta}) = -\mathbb{E}_{l\sim 1:L} \Big[\mathbb{E}_{T^c \sim T_{\min}:T_{\max}} \Big[\log p_{\boldsymbol{\theta}}(\operatorname{pos}(\boldsymbol{m}_{l,1:T}) \mid \mathcal{D}_l^c) \Big] \Big].$$
(5)

Each training batch consists of a task \mathcal{D}_l for which we sample a context size T^c uniformly between T_{\min} and T_{\max} to ensure that the model learns to handle different context set sizes. We then compute the latent task descriptor z_v and subsequently the predicted node trajectories $f(g_{\theta}(\mathcal{G}_{l,T^c}, z_v))$ as described in Equation 3 and Equation 4. The likelihood p_{θ} is defined to be the Gaussian

$$p_{\boldsymbol{\theta}}(\operatorname{pos}(\boldsymbol{m}_{l,1:T}) \mid \mathcal{D}_{l}^{c}) = \mathcal{N}(\operatorname{pos}(\boldsymbol{m}_{l,1:T}) \mid f(g_{\boldsymbol{\theta}}(\mathcal{G}_{l,T^{c}}, \boldsymbol{z}_{v})), \boldsymbol{\sigma}_{o}).$$
(6)

Since the training simulations are not affected by noise, we are not modeling the output variance and set it to $\sigma_0 = 1$. Together with taking the mean over the nodes and time steps to stabilize training, optimizing the Gaussian log likelihood from Equation 6 is equivalent to minimizing the MSE

$$\log p_{\boldsymbol{\theta}}(\operatorname{pos}(\boldsymbol{m}_{l,1:T}) \mid \mathcal{D}_{l}^{c}) \simeq \frac{1}{T \mid \mathcal{V} \mid d_{\operatorname{world}}} \sum_{t,v,i} \left(\operatorname{pos}(\boldsymbol{m}_{l,t})_{v,i} - f(g_{\boldsymbol{\theta}}(\mathcal{G}_{l,T^{c}}, \boldsymbol{z}_{v}))_{t,v,i} \right)^{2}$$

321 322

320

315

The whole architecture is trained end-to-end using the loss $\mathcal{L}(\theta)$ from Equation 5. After the metatraining, we fix θ , which now encodes inductive biases towards the meta-data \mathcal{D} .



Figure 4: Comparison of the final simulation step between **M3GN** and **MGN** on all datasets. From (**Left**) to (**Right**): (**Top**) *Planar Bending* and *Deformable Plate* with an anchor time step of 2. (**Bottom**) *Tissue Manipulation* with a context size of 6, and *Falling Teddy Bear* and *Mixed Objects Falling* with an anchor time step of 20. M3GN provides much better alignment to the **ground truth** simulation on all tasks, except for *Tissue Manipulation*, where MGN also solves the task well.

4 EXPERIMENTS

336

337

338

339 340 341

342 343

344

345

346

347

348

Setup. Our experimental setup largely follows previous work (Linkerhägner et al., 2023). The graph representation views the mesh vertices as nodes and adds edges according to the mesh topology within an object, and based on Euclidean distance between different objects. We employ one-hot encoding to differentiate between deformable objects and colliders, omitting explicit world edges (Pfaff et al., 2021). The edges additionally contain the relative distances between their nodes.

For both the context MPN as well as the simulator MPN, we use 15 message passing steps. Each 349 message passing step uses separate 1-layer MLPs with a latent dimension of 128 and LeakyReLU 350 activations for its node and edge updates. We evaluate the Full Rollout MSE, which is calculated 351 as the average of all simulation MSEs following the anchor time step, and the 10-step MSE, which 352 is the average performance for the next 10 simulation steps following the anchor time step. Both 353 metrics are averaged over all test set trajectories. We report the interquartile mean and bootstrapped 354 confidence intervals (Agarwal et al., 2021) over 8 random seeds for each experiment. We evaluate the 355 metrics for various context sizes ranging from 2 to 30 steps, always setting the anchor time step to 356 the last context mesh position. Appendix C provides additional details on our experimental setup.

- 357 **Datasets.** We validate our method on five different simulation datasets based on three different 358 mesh-based physics simulators. All task spaces are normalized to $[-1, 1]^3$. These include a 2D 359 Deformable Plate (DP) task and a 3D Tissue Manipulation (TM) task (Linkerhägner et al., 2023) In 360 both datasets, the material property, Poisson's ratio (Lim, 2015) was randomized. Deformable Plate 361 simulates different trapezoids that are deformed by a circular collider with constant velocity and 362 varying size and starting position. Each trajectory consists of a mesh with 81 nodes that is deformed 363 over 39 time steps. Tissue Manipulation considers a surgical robotics scenario where a piece of tissue is deformed by a gripper. The gripper is attached to a fixed object position and moves in a random 364 direction with constant velocity. The mesh comprises 361 nodes and the simulation has 100 steps. 365
- 366 We further propose three additional 3D datasets. Planar Bending (PB) simulates the bending of a 2D 367 plane when two constant forces perpendicular to the plane are applied at different positions. As the 368 Young's modulus is varied between sheets, this dataset constitutes a simplified stamp forming process, as common in material engineering (Zimmerling et al., 2022). The simulations are generated with 369 Abaqus (Smith, 2009), comprising 50 time steps and a plate with 225 nodes. We test two different 370 data splits: The in-distribution (ID) split uses material properties in the test set that the model has 371 seen during training, while the out-of-distribution (OOD) split uses Young's modulus values outside 372 the training domain for the test dataset. 373
- The other two tasks place a randomly rotated deformable object at a specific height and let it fall to and collide with the ground. *Falling Teddy Bear* (FTB) considers the titular teddy bear as its only object, whereas six different objects are considered for *Mixed Objects Falling* (MOF). Each trajectory assigns a random Poisson's ratio and random Young's modulus to the falling object, thus influencing its deformation upon contact with the floor. Each trajectory consists of 200 time steps. The object



Figure 5: Log-scale MSE over full rollouts for different methods under all tasks, including an 396 additional figure for Planar Bending task using an out-of-distribution (OOD) test dataset of material 397 properties. Overall, M3GN steadily improves its performance when provided with additional context 398 information and a later anchor time step. Our method generally improves over MGN, likely due to the MP trajectory representation, and outperforms MGN (Oracle) when provided with sufficient 400 context information. Equivariant Graph Neural Operator (EGNO) generally performs unstable for 401 later anchor time steps and can only compete in the *Planar Bending (ID)* task. 402

403 meshes have up to 350 nodes and are shown in Figure 9 in the appendix. For simplicity, we only 404 consider the triangular surface meshes for the experimental setup. Further information about the 405 graph encoding are given in Appendix B.1, while we provide detailed information of the dataset sizes 406 and preprocessing steps in Appendix D. 407

Baselines and Ablations. We compare our method to MGN(Pfaff et al., 2021), evaluating its 408 performance both with and without additional material property information provided as a node 409 feature. When this *oracle* information is available, the simulation becomes deterministic with respect 410 to the initial system state. Importantly, we never supply this node feature to M3GN. 411

MGN generates the next mesh state by iteratively predicting the velocities for the current simulation 412 step. It is trained to minimize the 1-step MSE over node velocities, incorporating Gaussian input 413 noise during training (Brandstetter et al., 2022). This noise serves to mitigate error accumulation and 414 stabilize auto-regressive rollouts during inference. To ensure a fair comparison, we adopt the same 415 hyperparameters as our method and fine-tune the input noise level for each task, maximizing MGN's 416 performance. 417

We also explore the effect of incorporating historical information, specifically previous velocities, as 418 node features for both MGN and M3GN. For MGN, using both the current and previous velocities 419 improves performance significantly on many tasks. For M3GN, including only the current velocity 420 yields similar benefits. The results of a hyperparameter optimization on the validation split are 421 presented in Figure 10 in the Appendix. Additionally, Table 1 summarizes the specific history 422 configurations used for each method and task, along with other relevant hyperparameters. 423

As an additional baseline, we compare our approach to the Equivariant Graph Neural Operator 424 (EGNO) method, which employs equivariant message-passing layers and predicts the remaining 425 simulation steps in a single pass, closely aligning with our setup. However, training EGNO proved 426 unstable with 15 message-passing steps, and the best results were achieved using only 5 steps. We 427 hypothesize that this instability may stem from the longer prediction horizon of up to 200 steps in our 428 experiments, as the baseline was originally evaluated on tasks with a much shorter prediction horizon 429 of only 8 steps. Further details on the implementation of these baselines can be found in Appendix C. 430

Additionally, we ablate different design choices of M3GN on Planar Bending and Deformable Plate. 431 To investigate the effect of the meta-learning approach, we train an MGN (MP) variant that uses



Figure 6: Log-scale MSE over full rollouts for the *Planar Bending* (Left) and *Deformable Plate* (**Right**) tasks for different meta-learning and MP variants. Using a ProDMP representation for MGN improves performance. CNPs and NPs with a next-step prediction do not improve over standard MGN. A NP instead of an CNP architecture for M3GN slightly reduces performance.

447 ProDMPs predictions, but omits a context aggregation and thus has no latent task description z_n . 448 Similarly, we compare to M3GN (Step-based), which performs a next-step prediction of the dynamics 449 instead of predicting ProDMP parameters, but otherwise follows the CNP training scheme to learn a 450 latent task description. Finally, we compare the deterministic CNP approach to both MP and step-451 based probabilistic Neural Process (NP) approaches. Here, we get diagonal Gaussian distributions 452 as the outputs of the context MPN, which we aggregate using Bayesian context aggregation (Volpp 453 et al., 2021). We further investigate if node aggregation of the latent task description is beneficial by applying a maximum aggregation of the node features before the context aggregation. While standard 454 CNPs require a permutation-invariant context aggregation, our context has a temporal structure. 455 We thus experiment with a small transformer model with 4 transformer blocks, 4 attention heads, 456 temporal encoding and a latent dimension of 32 as an aggregator. The transformer takes the sequence 457 of outputs of the context MPN and predicts the aggregated node-level task description z_v . 458

Results. Figure 4 visualizes exemplary final simulation steps for M3GN and MGN for all tasks.
M3GN aggregates context information that it uses to condition node-level ProDMP representations of the simulated trajectory. This approach leads to accurate simulations, providing much better alignment to the ground truth simulation than the step-based MGN on all tasks. Appendix E.3 shows visualizations of full simulation rollouts for all tasks and methods.²

464 Figure 5 shows the full rollout MSE for M3GN, MGN and MGN (Oracle) on all tasks. In general, the 465 additional material information improves performance of MGN (Oracle) on both datasets compared 466 to MGN. A later anchor time step improves performance for all methods, presumably because the remaining simulation is shorter. On Deformable Plate, M3GN surpasses MGN (Oracle) for a context 467 size of 10, likely because the additional context improves the latent task description. For the *Planar* 468 Bending task, M3GN significantly outperforms the step-based baselines across context sizes, likely 469 because the temporal smoothness of the ProDMP trajectory is a strong inductive bias for the gradual 470 bending of the simulated plane. Furthermore, M3GN generalizes well to the OOD task, while the 471 MGN methods fail to extrapolate to unseen material properties. 472

Next, in a more difficult task, such as *Tissue Manipulation*, increasing the context size greatly 473 improves performance for M3GN, whereas the step-based methods only slightly benefit from a later 474 anchor time step. On the last two tasks, i.e. Falling Teddy Bear and Mixed Objects Falling, the step-475 based MGN methods fail to provide accurate long-term simulations. Here, the predicted trajectories 476 usually qualitatively deviate from the ground truth, either causing an object drift or a misalignment 477 of, e.g., teddy limbs, to the point that providing additional material information only yields marginal 478 improvements. For M3GN, the ProDMP's temporally consistent movements combined with context 479 aggregation to provide consistent simulations alleviates these issues, significantly improving over 480 the baselines. The bottom of Figure 4 provides examples. Interestingly, providing more context 481 information does not improve performance for either of these two tasks. A likely reason for this 482 behavior is that most early context steps consist of a falling object, resulting in similar graph 483 representations, which may cause the GNN context encoders to overfit.

484 485

442

443

444

²Videos of these simulations are in the supplementary material.

486 Figure 6 provides evaluations for various ablations, further supporting these results. Equipping MGN 487 with a ProDMP representation improves performance, especially for *Planar Bending*. Combining 488 NPs and ProDMPs uniquely leads to the strong performance of M3GN, suggesting that the latent task 489 description extracted from the context set is particularly well-suited for trajectory-level representations 490 of the simulations. Performance decreases slightly when using an NP instead of a CNP, indicating that the latent distribution in NPs is not beneficial for our GNS setup. Additional ablations in Figure 11 of 491 the Appendix show similar performance for different aggregation schemes. Our node-level maximum 492 context aggregation is the simplest and works slightly better than the alternatives for *Planar Bending*. 493

To provide a more detailed understanding of our model's performance, we include additional plots in Appendix E showing the Mean Squared Error (MSE) over time for the trajectory, rather than just the final averaged MSE. These plots demonstrate the temporal progression of errors, offering insights into the model's behavior throughout the simulation.

We additionally report the 10-step MSE in Appendix E.2, find-498 ing that the relative improvement of M3GN compared to the 499 baselines matches or exceeds that on the full rollout MSE for 500 all tasks. Furthermore, we present a visualization of the latent 501 space for M3GN in the Appendix, Figure 20, which reveals that 502 simulations with similar material properties are clustered together. This structure in the latent space highlights the model's 504 ability to effectively differentiate between different material 505 behaviors while preserving the relationships between similar 506 properties. Finally, we compare the inference speed between 507 M3GN and MGN in Figure 7. The ProDMPs trajectory rep-508 resentation of M3GN decreases the amount of required model calls, resulting in an inference-time speedup of up to 32 times 509 compared to MGN, and up to 400 times compared to the ground 510 truth simulators (Gth Simulators). We encode the context set 511 for M3GN in parallel, resulting in a relatively minor cost for the 512 context computation and aggregation for all tasks. In addition, 513 our M3GN requires only one GNN forward pass via ProDMP 514



Figure 7: Runtime comparison on four tasks between the learned methods and the different ground truth simulators. Note the log scale on the y-axis.

to compute the full trajectory, enabling much faster inference. While MGN does not perform any context processing, its iterative rollout is inherently sequential and requires multiple forward passes.

517

5 CONCLUSION

518 519

520 We introduce Movement-Primitive Meta-MeshGraphNet (M3GN), a novel Graph Network Simulator 521 that combines movement primitives and trajectory-level meta-learning for efficient and accurate 522 long-term predictions in physical simulations. Our method dynamically adapts to provided context 523 information during inference, allowing for an accurate prediction of deformations under unknown 524 object properties. Additionally, it effectively addresses the issue of error accumulation while reducing 525 the number of required simulator function calls. To validate the effectiveness of M3GN, we propose three novel deformation prediction tasks with uncertain material properties. Results on these tasks and 526 existing datasets show that our method consistently outperforms a strong Graph Network Simulators 527 baseline, even when providing the baseline with oracle information about the material property. 528

Limitations and Future Work. We currently consider each trajectory as a task, and require initial states of this trajectory as a context set during inference. As generating such data is often impractical, we plan to group simulations with the same system properties into the same task. This adaptation will enable data-efficient generalization to unseen properties regardless of the underlying simulated objects. We also plan to integrate online re-planning of trajectories, predicting trajectory segments with every model forward pass. This process may increase coordination between simulated nodes across segments, while maintaining the benefits of a compact multi-step trajectory representation.

Broader Impact Statement Our proposed Graph Network Simulator can positively impact various
 fields relying on computational modeling and simulation by significantly reducing computational
 cost compared to traditional simulators while providing accurate simulations. However, efficient and
 accurate simulation of physical systems also comes with potential negative impacts, such as, e.g., the
 development of advanced weapon models.

540 REFERENCES

570

571

- Agarwal, R., Schwarzer, M., Castro, P. S., Courville, A. C., and Bellemare, M. Deep reinforcement learning at the edge of the statistical precipice. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), Advances in Neural Information Processing Systems, volume 34, pp. 29304–29320. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/ file/f514cec81cb148559cf475e7426eed5e-Paper.pdf.
- Allen, K., Lopez-Guevara, T., Stachenfeld, K. L., Sanchez Gonzalez, A., Battaglia, P., Hamrick, J. B., and Pfaff, T. Inverse design for fluid-structure interactions using graph network simulators. In Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., and Oh, A. (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 13759–13774. Curran Associates, Inc., 2022a. URL https://proceedings.neurips.cc/paper_files/paper/2022/ file/59593615e358d52295578e0d8e94ec4a-Paper-Conference.pdf.
- Allen, K. R., Guevara, T. L., Rubanova, Y., Stachenfeld, K., Sanchez-Gonzalez, A., Battaglia, P., and
 Pfaff, T. Graph network simulators can learn discontinuous, rigid contact dynamics. *Conference* on Robot Learning (CoRL)., 2022b.
- Allen, K. R., Rubanova, Y., Lopez-Guevara, T., Whitney, W. F., Sanchez-Gonzalez, A., Battaglia, P., and Pfaff, T. Learning rigid dynamics with face interaction graph networks. In *The Eleventh International Conference on Learning Representations*, 2023.
- Andrychowicz, M., Denil, M., Colmenarejo, S. G., Hoffman, M. W., Pfau, D., Schaul, T., and de Freitas, N. Learning to Learn by Gradient Descent by Gradient Descent. *Advances in Neural Information Processing Systems*, 2016.
- Antonova, R., Yang, J., Sundaresan, P., Fox, D., Ramos, F., and Bohg, J. A bayesian treatment of real-to-sim for deformable object manipulation. *IEEE Robotics and Automation Letters*, 7(3): 5819–5826, 2022.
- Bahl, S., Mukadam, M., Gupta, A., and Pathak, D. Neural dynamic policies for end-to-end sensori motor learning. *Advances in Neural Information Processing Systems*, 33:5058–5069, 2020.
 - Bakker, B. and Heskes, T. Task clustering and gating for Bayesian multitask learning. *Journal of Machine Learning Research*, 2003.
- Baqué, P., Remelli, E., Fleuret, F., and Fua, P. Geodesic convolutional shape optimization. In Dy, J. G. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018, volume 80 of Proceedings of Machine Learning Research*, pp. 481–490. PMLR, 2018. URL http://proceedings.mlr.press/v80/baque18a.html.
- Battaglia, P., Pascanu, R., Lai, M., Jimenez Rezende, D., and kavukcuoglu, k. Interaction networks
 for learning about objects, relations and physics. In Lee, D., Sugiyama, M., Luxburg, U., Guyon,
 I., and Garnett, R. (eds.), Advances in Neural Information Processing Systems, volume 29. Curran Associates, Inc., 2016. URL https://proceedings.neurips.cc/paper/2016/
 file/3147da8ab4a0437c15ef51a5cc7f2dc4-Paper.pdf.
- Battaglia, P. W., Hamrick, J. B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V. F., Malinowski,
 M., Tacchetti, A., Raposo, D., Santoro, A., Faulkner, R., Gülçehre, Ç., Song, H. F., Ballard,
 A. J., Gilmer, J., Dahl, G. E., Vaswani, A., Allen, K. R., Nash, C., Langston, V., Dyer, C.,
 Heess, N., Wierstra, D., Kohli, P., Botvinick, M., Vinyals, O., Li, Y., and Pascanu, R. Relational
 inductive biases, deep learning, and graph networks. *CoRR*, abs/1806.01261, 2018. URL http:
 //arxiv.org/abs/1806.01261.
- Bengio, Y., Bengio, S., and Cloutier, J. Learning a synaptic learning rule. *International Joint Conference on Neural Networks*, 1991.
- Bodnar, C., Bruinsma, W. P., Lucic, A., Stanley, M., Brandstetter, J., Garvan, P., Riechert, M., Weyn,
 J. A., Dong, H., Vaughan, A., Gupta, J. K., Thambiratnam, K., Archibald, A., Heider, E., Welling,
 M., Turner, R. E., and Perdikaris, P. Aurora: A foundation model of the atmosphere. *CoRR*,

594 595	abs/2405.13063, 2024. doi: 10.48550/ARXIV.2405.13063. URL https://doi.org/10.48550/arXiv.2405.13063.
596 597 598	Brandstetter, J., Worrall, D. E., and Welling, M. Message passing neural pde solvers. In <i>International Conference on Learning Representations</i> , 2022.
599 600 601	Brenner, S. C. and Scott, L. R. <i>The mathematical theory of finite element methods</i> , volume 3. Springer, 2008.
602 603 604	Bronstein, M. M., Bruna, J., Cohen, T., and Velickovic, P. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. <i>CoRR</i> , abs/2104.13478, 2021. URL https://arxiv.org/abs/2104.13478.
605 606 607	Chung, T. Finite element analysis in fluid dynamics. NASA STI/Recon Technical Report A, 78:44102, 1978.
608 609	Connor, J. J. and Brebbia, C. A. Finite element techniques for fluid flow. Newnes, 2013.
610 611	Da Wang, Y., Blunt, M. J., Armstrong, R. T., and Mostaghimi, P. Deep learning in pore scale imaging and modeling. <i>Earth-Science Reviews</i> , 215:103555, 2021.
612 613 614	Durasov, N., Lukoyanov, A., Donier, J., and Fua, P. Debosh: Deep bayesian shape optimization. <i>arXiv preprint arXiv:2109.13337</i> , 2021.
615 616 617 618 619 620 621	Faure, F., Duriez, C., Delingette, H., Allard, J., Gilles, B., Marchesseau, S., Talbot, H., Courtecuisse, H., Bousquet, G., Peterlik, I., and Cotin, S. SOFA: A Multi-Model Framework for Interactive Physical Simulation. In Payan, Y. (ed.), Soft Tissue Biomechanical Modeling for Computer Assisted Surgery, volume 11 of Studies in Mechanobiology, Tissue Engineering and Biomaterials, pp. 283– 321. Springer, June 2012. doi: 10.1007/8415_2012_125. URL https://hal.inria.fr/ hal-00681539.
622 623	Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. <i>International Conference on Machine Learning</i> , 2017.
624 625 626	Finn, C., Xu, K., and Levine, S. Probabilistic Model-Agnostic Meta-Learning. Advances in Neural Information Processing Systems, 2018.
627 628 629	Fortunato, M., Pfaff, T., Wirnsberger, P., Pritzel, A., and Battaglia, P. Multiscale meshgraphnets. In <i>ICML 2022 2nd AI for Science Workshop</i> , 2022. URL https://openreview.net/forum?id=G3TRIsmMhhf.
630 631 632	Freymuth, N., Dahlinger, P., Würth, T., Kärger, L., and Neumann, G. Swarm reinforcement learning for adaptive mesh refinement. <i>Neural Information Processing Systems</i> , 37, 2023.
633 634 635 636	Garnelo, M., Rosenbaum, D., Maddison, C., Ramalho, T., Saxton, D., Shanahan, M., Teh, Y. W., Rezende, D., and Eslami, S. M. A. Conditional neural processes. <i>International Conference on Machine Learning</i> , 2018a.
637 638 639	Garnelo, M., Schwarz, J., Rosenbaum, D., Viola, F., Rezende, D. J., Eslami, S. M. A., and Teh, Y. W. Neural processes. <i>ICML Workshop on Theoretical Foundations and Applications of Deep Generative Models</i> , 2018b.
640 641 642 643	Golovin, D., Solnik, B., Moitra, S., Kochanski, G., Karro, J., and Sculley, D. Google vizier: a service for black-box optimization. <i>International Conference on Knowledge Discovery and Data Mining</i> , 2017.
644 645	Gordon, J., Bronskill, J., Bauer, M., Nowozin, S., and Turner, R. E. Meta-Learning Probabilistic Inference for Prediction. <i>International Conference on Learning Representations</i> , 2019.
646 647	Grant, E., Finn, C., Levine, S., Darrell, T., and Griffiths, T. L. Recasting Gradient-Based Meta- Learning as Hierarchical Bayes. <i>International Conference on Learning Representations</i> , 2018.

660

666

687

691

- 648 Guo, X., Li, W., and Iorio, F. Convolutional neural networks for steady flow approximation. In 649 Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and 650 Data Mining, KDD '16, pp. 481–490, New York, NY, USA, 2016. Association for Computing 651 Machinery. ISBN 9781450342322. doi: 10.1145/2939672.2939738. URL https://doi.org/ 652 10.1145/2939672.2939738.
- Han, X., Gao, H., Pfaff, T., Wang, J., and Liu, L. Predicting physics in mesh-reduced space with 654 temporal attention. In The Tenth International Conference on Learning Representations, ICLR 655 2022, Virtual Event, April 25-29, 2022. OpenReview.net, 2022. URL https://openreview. 656 net/forum?id=XctLdNfCmP. 657
- 658 Heskes, T. Empirical Bayes for learning to learn. International Conference on Machine Learning, 659 2000.
- Hospedales, T., Antoniou, A., Micaelli, P., and Storkey, A. Meta-learning in neural networks: a 661 survey. IEEE Transactions on Pattern Analysis and Machine Intelligence, 2022. 662
- 663 Hu, W., Shuaibi, M., Das, A., Goyal, S., Sriram, A., Leskovec, J., Parikh, D., and Zitnick, C. L. 664 Forcenet: A graph neural network for large-scale quantum calculations. In ICLR 2021 SimDL 665 Workshop, volume 20, 2021.
- Hughes, T. J., Cottrell, J. A., and Bazilevs, Y. Isogeometric analysis: Cad, finite elements, nurbs, 667 exact geometry and mesh refinement. Computer methods in applied mechanics and engineering, 668 194(39-41):4135-4195, 2005. 669
- 670 Jin, J.-M. The finite element method in electromagnetics. John Wiley & Sons, 2015. 671
- Kim, H., Mnih, A., Schwarz, J., Garnelo, M., Eslami, A., Rosenbaum, D., Vinyals, O., and Teh, Y. W. 672 Attentive neural processes. International Conference on Learning Representations, 2019. 673
- 674 Kim, T., Yoon, J., Dia, O., Kim, S., Bengio, Y., and Ahn, S. Bayesian Model-Agnostic Meta-Learning. 675 Advances in Neural Information Processing Systems, 2018. 676
- 677 Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In Bengio, Y. and LeCun, 678 Y. (eds.), 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings, 2015. URL http://arxiv.org/abs/ 679 1412.6980. 680
- 681 Kingma, D. P. and Welling, M. Auto-encoding variational bayes. International Conference on 682 Learning Representations, 2014. 683
- 684 Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional 685 neural networks. In Pereira, F., Burges, C. J. C., Bottou, L., and Weinberger, K. Q. (eds.), Advances 686 in Neural Information Processing Systems, volume 25. Curran Associates, Inc., 2012.
- Li, G., Jin, Z., Volpp, M., Otto, F., Lioutikov, R., and Neumann, G. Prodmp: A unified perspective 688 on dynamic and probabilistic movement primitives. *IEEE Robotics and Automation Letters*, 8(4): 689 2325-2332, 2023. 690
- Li, G., Zhou, H., Roth, D., Thilges, S., Otto, F., Lioutikov, R., and Neumann, G. Open the black box: 692 Step-based policy updates for temporally-correlated episodic reinforcement learning, 2024.
- Li, J., Du, X., and Martins, J. R. Machine learning in aerodynamic shape optimization. Progress in 694 Aerospace Sciences, 134:100849, 2022. 695
- 696 Li, Y., Wu, J., Tedrake, R., Tenenbaum, J. B., and Torralba, A. Learning particle dynamics for 697 manipulating rigid bodies, deformable objects, and fluids. In International Conference on Learning Representations, 2019. URL https://openreview.net/forum?id=rJgbSn09Ym. 699
- Lim, T.-C. Auxetic Materials and Structures. Springer Singapore, 01 2015. ISBN 978-700 981-287-274-6. doi: 10.1007/978-981-287-275-3. URL https://doi.org/10.1007/ 978-981-287-275-3.

702 703 704 705	Linkerhägner, J., Freymuth, N., Scheikl, P. M., Mathis-Ullrich, F., and Neumann, G. Grounding graph network simulators using physical sensor observations. In <i>The Eleventh International Conference on Learning Representations</i> , 2023.
706 707	Lopez-Guevara, T., Rubanova, Y., Whitney, W. F., Pfaff, T., Stachenfeld, K., and Allen, K. R. Scaling face interaction graph networks to real world scenes. <i>arXiv preprint arXiv:2401.11985</i> , 2024.
708 709 710	Louizos, C., Shi, X., Schutte, K., and Welling, M. The functional neural process. <i>Advances in Neural</i> <i>Information Processing Systems</i> , 2019.
711 712 713 714 715	Mora, M. A. Z., Peychev, M., Ha, S., Vechev, M., and Coros, S. Pods: Policy optimization via differentiable simulation. In Meila, M. and Zhang, T. (eds.), <i>Proceedings of the 38th International Conference on Machine Learning</i> , volume 139 of <i>Proceedings of Machine Learning Research</i> , pp. 7805–7817. PMLR, 18–24 Jul 2021. URL http://proceedings.mlr.press/v139/mora21a.html.
716	NVIDIA. Isaac sim, 2022a. URL https://developer.nvidia.com/isaac-sim.
717 718	NVIDIA. Nvidia physx sdk, 2022b. URL https://developer.nvidia.com/physx-sdk.
719 720 721 722 723	Otto, F., Celik, O., Zhou, H., Ziesche, H., Vien, N. A., and Neumann, G. Deep black-box rein- forcement learning with movement primitives. In Liu, K., Kulic, D., and Ichnowski, J. (eds.), <i>Conference on Robot Learning, CoRL 2022, 14-18 December 2022, Auckland, New Zealand</i> , volume 205 of <i>Proceedings of Machine Learning Research</i> , pp. 1244–1265. PMLR, 2022. URL https://proceedings.mlr.press/v205/otto23a.html.
724 725 726	Paraschos, A., Daniel, C., Peters, J. R., and Neumann, G. Probabilistic movement primitives. <i>Advances in neural information processing systems</i> , 26, 2013.
727	Paszynski, M. Fast solvers for mesh-based computations. CRC Press, 2016.
728 729 730 731	Pfaff, T., Fortunato, M., Sanchez-Gonzalez, A., and Battaglia, P. W. Learning mesh-based simulation with graph networks. In <i>International Conference on Learning Representations</i> , 2021. URL https://arxiv.org/abs/2010.03409.
732 733	Plewa, T., Linde, T., Weirs, V. G., et al. <i>Adaptive mesh refinement-theory and applications</i> . Springer, 2005.
734 735 736	Polycarpou, A. C. Introduction to the finite element method in electromagnetics. Springer Nature, 2022.
737 738	Ravi, S. and Larochelle, H. Optimization as a model for few-shot learning. <i>International Conference on Learning Representations</i> , 2017.
739 740 741	Reddy, C. <i>Finite element method for eigenvalue problems in electromagnetics</i> , volume 3485. NASA, Langley Research Center, 1994.
742	Reddy, J. N. Introduction to the finite element method. McGraw-Hill Education, 2019.
743 744 745	Reddy, J. N. and Gartling, D. K. <i>The finite element method in heat transfer and fluid dynamics</i> . CRC press, 2010.
746 747	Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. <i>International Conference on Machine Learning</i> , 2014.
748 749 750 751 752 753	Sanchez-Gonzalez, A., Heess, N., Springenberg, J. T., Merel, J., Riedmiller, M., Hadsell, R., and Battaglia, P. Graph networks as learnable physics engines for inference and control. In Dy, J. and Krause, A. (eds.), <i>Proceedings of the 35th International Conference on Machine Learning</i> , volume 80 of <i>Proceedings of Machine Learning Research</i> , pp. 4470–4479. PMLR, 10–15 Jul 2018. URL https://proceedings.mlr.press/v80/sanchez-gonzalez18a.html.
754 755	Sanchez-Gonzalez, A., Godwin, J., Pfaff, T., Ying, R., Leskovec, J., and Battaglia, P. Learning to simulate complex physics with graph networks. In <i>Proceedings of the 37th International Conference on Machine Learning</i> , pp. 8459–8468. PMLR, 2020.

756 757 758	Santoro, A., Bartunov, S., Botvinick, M. M., Wierstra, D., and Lillicrap, T. P. Meta-learning with memory-augmented neural networks. <i>International Conference on Machine Learning</i> , 2016.
759 760 761	Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M., and Monfardini, G. The graph neural network model. <i>IEEE Transactions on Neural Networks</i> , 20(1):61–80, 2009. doi: 10.1109/TNN.2008. 2005605.
762 763 764	Schaal, S. Dynamic movement primitives-a framework for motor control in humans and humanoid robotics. In <i>Adaptive motion of animals and machines</i> , pp. 261–280. Springer, 2006.
765 766 767	Scheikl, P. M., Tagliabue, E., Gyenes, B., Wagner, M., Dall'Alba, D., Fiorini, P., and Mathis-Ullrich, F. Sim-to-real transfer for visual reinforcement learning of deformable object manipulation for robot-assisted surgery. <i>IEEE Robotics and Automation Letters</i> , 8(2):560–567, 2022.
768 769 770	Schmidhuber, J. Learning to control fast-weight memories: an alternative to dynamic recurrent networks. <i>Neural Computation</i> , 1992.
771 772	Seker, M. Y., Imre, M., Piater, J. H., and Ugur, E. Conditional neural movement primitives. In <i>Robotics: Science and Systems</i> , volume 10, 2019.
773 774 775	Smith, M. ABAQUS/Standard User's Manual, Version 6.9. Dassault Systèmes Simulia Corp, United States, 2009.
776 777	Snell, J., Swersky, K., and Zemel, R. S. Prototypical networks for few-shot learning. Advances in Neural Information Processing Systems, 2017.
778 779 780 781	Stanova, E., Fedorko, G., Kmet, S., Molnar, V., and Fabian, M. Finite element analysis of spiral strands with different shapes subjected to axial loads. <i>Advances in engineering software</i> , 83:45–58, 2015.
782	Thrun, S. and Pratt, L. Learning to Learn. Kluwer Academic Publishers, 1998.
783 784 785 786	<pre>van der Maaten, L. and Hinton, G. Visualizing data using t-sne. Journal of Machine Learning Research, 9(86):2579-2605, 2008. URL http://jmlr.org/papers/v9/ vandermaaten08a.html.</pre>
787 788	Vilalta, R. and Drissi, Y. A Perspective View and Survey of Meta-Learning. <i>Artificial Intelligence Review</i> , 2005.
789 790 791 792	Volpp, M., Fröhlich, L. P., Fischer, K., Doerr, A., Falkner, S., Hutter, F., and Daniel, C. Meta-learning acquisition functions for transfer learning in Bayesian optimization. <i>International Conference on Learning Representations</i> , 2019.
793 794	Volpp, M., Flürenbrock, F., Großberger, L., Daniel, C., and Neumann, G. Bayesian context aggregation for neural processes. <i>International Conference on Learning Representations</i> , 2021.
795 796 797 798 799	Volpp, M., Dahlinger, P., Becker, P., Daniel, C., and Neumann, G. Accurate bayesian meta-learning by accurate task posterior inference. In <i>The Eleventh International Conference on Learning</i> <i>Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023.</i> OpenReview.net, 2023. URL https://openreview.net/pdf?id=sb-IkS8DQw2.
800 801 802 803	Wang, J., SUN, J., Zhang, Z., He, J., Zhang, Q., Sun, M., and Xu, R. DEL: Discrete element learner for learning 3d particle dynamics with neural rendering. In <i>The Thirty-eighth Annual</i> <i>Conference on Neural Information Processing Systems</i> , 2024. URL https://openreview. net/forum?id=2nvkD0sPOk.
804 805 806	Wang, L. and Zhu, J. Deformable object manipulation in caregiving scenarios: A review. <i>Machines</i> , 11(11):1013, 2023.
807 808 809	Weng, Z., Paus, F., Varava, A., Yin, H., Asfour, T., and Kragic, D. Graph-based task-specific prediction models for interactions between deformable and rigid objects. In 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 5741–5748, 2021. doi: 10.1109/IROS51168.2021.9636660. URL https://arxiv.org/abs/2103.02932.

810 811 812	Whitney, W. F., Lopez-Guevara, T., Pfaff, T., Rubanova, Y., Kipf, T., Stachenfeld, K., and Allen, K. R. Learning 3d particle-based simulators from rgb-d videos. <i>arXiv preprint arXiv:2312.05359</i> , 2023.
813 814 815	Wu, T., Maruyama, T., Zhao, Q., Wetzstein, G., and Leskovec, J. Learning controllable adaptive simulation for multi-resolution physics. In <i>The Eleventh International Conference on Learning Representations</i> , 2023.
816 817 818	Xu, J., Chen, T., Zlokapa, L., Foshey, M., Matusik, W., Sueda, S., and Agrawal, P. An end-to-end differentiable framework for contact-aware robot design. In <i>Robotics: Science & Systems</i> , 2021.
819 820 821 822	Xu, M., Han, J., Lou, A., Kossaifi, J., Ramanathan, A., Azizzadenesheli, K., Leskovec, J., Ermon, S., and Anandkumar, A. Equivariant graph neural operator for modeling 3d dynamics. In <i>Forty-first</i> <i>International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024.</i> OpenReview.net , 2024. URL https://openreview.net/forum?id=dccRCYmL5x.
823 824	Yazid, A., Abdelkader, N., and Abdelmadjid, H. A state-of-the-art review of the x-fem for computa- tional fracture mechanics. <i>Applied Mathematical Modelling</i> , 33(12):4269–4282, 2009.
825 826 827	Zaheer, M., Kottur, S., Ravanbakhsh, S., Póczos, B., Salakhutdinov, R., and Smola, A. J. Deep sets. <i>Advances in Neural Information Processing Systems</i> , 2017.
828 829	Zhuang, F., Qi, Z., Duan, K., Xi, D., Zhu, Y., Zhu, H., Xiong, H., and He, Q. A comprehensive survey on transfer learning. <i>Proceedings of the IEEE</i> , 2020.
830 831 832	Zienkiewicz, O. C. and Taylor, R. L. <i>The finite element method for solid and structural mechanics</i> . Elsevier, 2005.
833 834	Zienkiewicz, O. C., Taylor, R. L., and Nithiarasu, P. <i>The finite element method for fluid dynamics</i> . Butterworth-Heinemann, 2013.
835 836 837 838 839	Zimmerling, C., Poppe, C., Stein, O., and Kärger, L. Optimisation of manufacturing process parameters for variable component geometries using reinforcement learning. <i>Materials and Design</i> , 214:Art.–Nr.: 110423, 2022. ISSN 0264-1275, 0141-5530, 0261-3069, 1873-4197, 1878-2876. doi: 10.1016/j.matdes.2022.110423.
840	
841	
842	
843	
844	
845	
846	
847	
848	
849	
850	
050	
952	
85/	
855	
856	
857	
858	
859	
860	
861	
862	
863	

A MATHEMATICAL FORMULATIONS OF MOVEMENT PRIMITIVES

We provide an overview of the probabilistic dynamic movement primitives (ProDMP) formulations utilized in this paper, starting with the foundational methods: Dynamic Movement Primitives (DMPs) and Probabilistic Movement Primitives (ProMPs).

A.1 DMPs

Schaal (2006) introduced Dynamic Movement Primitives (DMPs), which integrate a forcing term
into a dynamical system to generate smooth trajectories from given initial conditions³, such as a
robot's position and velocity at a particular time. A DMP trajectory is governed by a second-order
linear ordinary differential equation (ODE) as follows:

864

865 866

867

868

869 870

871

$$\tau^{2}\ddot{y} = \alpha(\beta(g-y) - \tau\dot{y}) + f(x), \quad f(x) = x\frac{\sum\varphi_{i}(x)w_{i}}{\sum\varphi_{i}(x)} = x\varphi_{x}^{\mathsf{T}}\boldsymbol{w}, \tag{7}$$

878

888

889 890

896

897

898

904 905

906 907 908

909

910

879 where y = y(t), $\dot{y} = dy/dt$, and $\ddot{y} = d^2y/dt^2$ denote the position, velocity, and acceleration of the 880 system at a specific time t, respectively. Constants α and β are spring-damper parameters, g is the 881 goal attractor, and τ is a time constant modulating the speed of trajectory execution.

The functions $\varphi_i(x)$ represent the basis functions for the forcing term, as shown in Fig. 8a, while the phase variable $x = x(t) \in [0, 1]$ captures the execution progress. The trajectory's shape is determined by the weight parameters $w_i \in w$ for i = 1, ..., N and the goal term g. The trajectory $[y_t]_{t=0:T}$ is typically computed by numerically integrating the dynamical system from the start to the endpoint. However, this numerical process is computationally expensive (Bahl et al., 2020; Li et al., 2023), as its cost scales with the trajectory length and the resolution of the numerical integration.

A.2 PROMPS

Paraschos et al. (2013) introduced the Probabilistic Movement Primitives (ProMPs) framework for modeling trajectory distributions, effectively capturing both temporal and inter-dimensional correlations. Unlike DMPs, which rely on a forcing term, ProMPs directly model the desired trajectory and its distribution using a linear basis function representation. Given a weight vector wor a weight vector distribution $p(w) \sim \mathcal{N}(w|\mu_w, \Sigma_w)$, the corresponding trajectory or trajectory distribution is computed as follows:

Compute Trajectory:
$$[y_t]_{t=0:T} = \mathbf{\Phi}^{\mathsf{T}} \boldsymbol{w},$$
 (8)

Compute Distribution:
$$p([y_t]_{t=0:T}; \boldsymbol{\mu}_{\boldsymbol{y}}, \boldsymbol{\Sigma}_{\boldsymbol{y}}) = \mathcal{N}(\boldsymbol{\Phi}^{\mathsf{T}} \boldsymbol{\mu}_{\boldsymbol{w}}, \boldsymbol{\Phi}^{\mathsf{T}} \boldsymbol{\Sigma}_{\boldsymbol{w}} \boldsymbol{\Phi}).$$
 (9)

Here, the matrix Φ contains the basis functions for each time step $t \in [0, T]$, shown in Fig. 8a. The trajectory shape is determined by the weight parameters $w_i \in w$ through matrix-vector multiplication. Despite their simplicity and computational efficiency, ProMPs lack an intrinsic dynamic system, limiting their ability to specify a given initial condition for a trajectory or predict smooth transitions between two ProMP trajectories with differing parameter vectors.

A.3 PRODMPS

Solving the ODE underlying DMPs Li et al. (2023) observed that the governing equation of DMPs, as described in Eq. (7), admits an analytical solution. We re-express the original ODE from Eq. (7) and its homogeneous counterpart in standard ODE forms as follows:

Non-homo. ODE:
$$\ddot{y} + \frac{\alpha}{\tau}\dot{y} + \frac{\alpha\beta}{\tau^2}y = \frac{f(x)}{\tau^2} + \frac{\alpha\beta}{\tau^2}g \equiv F(x,g),$$
 (10)

Homo. ODE:
$$\ddot{y} + \frac{\alpha}{\tau}\dot{y} + \frac{\alpha\beta}{\tau^2}y = 0.$$
 (11)

³In mathematics, an initial condition refers to the value of a function or its derivatives at a starting point, which can be specified at any time, not necessarily at t = 0.

930

931

932

933

934 935

938

944

945

954

955

966

967 968



Figure 8: Illustration of basis functions used in MP methods. (a) Normalized radial basis functions used in DMPs in Eq.(7) and ProMPs in Eq.(8), respectively. (b) Positional basis functions of ProDMPs' weights w and (c) ProDMPs' goal g in Eq.(17). In ProDMPs, g is concatenated with the weights vector w and treated as one dimension of the resulting vector w_g . Both weights and goal basis functions are computed from solving the DMPs' underlying ODE, following the procedure from Eq.(12) to Eq.(16)

936 The solution to this ODE is essentially the position trajectory, and its time derivative yields the 937 velocity trajectory. They are formulated through several time-dependent function as:

$$y = [y_2 p_2 - y_1 p_1 \quad y_2 q_2 - y_1 q_1] \begin{bmatrix} w \\ g \end{bmatrix} + c_1 y_1 + c_2 y_2$$
(12)

$$\dot{y} = [\dot{y}_2 \mathbf{p}_2 - \dot{y}_1 \mathbf{p}_1 \quad \dot{y}_2 q_2 - \dot{y}_1 q_1] \begin{bmatrix} \mathbf{w} \\ g \end{bmatrix} + c_1 \dot{y}_1 + c_2 \dot{y}_2.$$
(13)

Here, the learnable parameters $[w, g]^T$ which control the shape of the trajectory, are separable from the remaining time-dependent functions $y_1, y_2, p_1, p_2, q_1, q_2$. These functions are computed by solving the ODE in Eq. (10), (11):

$$y_1(t) = \exp\left(-\frac{\alpha}{2\tau}t\right),$$
 $y_2(t) = t\exp\left(-\frac{\alpha}{2\tau}t\right),$ (14)

$$\boldsymbol{p}_{1}(t) = \frac{1}{\tau^{2}} \int_{0}^{t} t' \exp\left(\frac{\alpha}{2\tau}t'\right) x(t') \boldsymbol{\varphi}_{x}^{\mathsf{T}} \mathrm{d}t', \qquad \boldsymbol{p}_{2}(t) = \frac{1}{\tau^{2}} \int_{0}^{t} \exp\left(\frac{\alpha}{2\tau}t'\right) x(t') \boldsymbol{\varphi}_{x}^{\mathsf{T}} \mathrm{d}t', \tag{15}$$

$$q_1(t) = \left(\frac{\alpha}{2\tau}t - 1\right)\exp\left(\frac{\alpha}{2\tau}t\right) + 1, \qquad q_2(t) = \frac{\alpha}{2\tau}\left[\exp\left(\frac{\alpha}{2\tau}t\right) - 1\right]. \tag{16}$$

Here, the function y_1, y_2 are the complementary solutions to the homogeneous ODE presented in Eq.(11), with \dot{y}_1, \dot{y}_2 their time derivatives respectively.

956 It's worth noting that p_1 and p_2 cannot be derived analytically due to the complexity of the forcing 957 basis terms φ_x . Consequently, these terms must be computed numerically. However, isolating the 958 learnable parameters, namely w and g, enables the reuse of other time-dependent functions across all 959 generated trajectories.

960 ProDMPs identify these reusable terms as the position and velocity basis functions, denoted by $\Phi(t)$ 961 and $\dot{\Phi}(t)$, respectively. Fig. 8b and Fig. 8c illustrate the resulting position basis functions for the 962 weights w and the goal g, respectively. These functions are pre-computed offline and treated as 963 constants during online learning. When w and g are combined into a concatenated vector, represented 964 as w_g , the position and velocity trajectories can be expressed in a manner similar to that used by 965 ProMPs:

Position:
$$y(t) = \Phi(t)^{\mathsf{T}} w_g + c_1 y_1(t) + c_2 y_2(t),$$
 (17)

Velocity:
$$\dot{y}(t) = \mathbf{\Phi}(t)^{\mathsf{T}} \boldsymbol{w}_g + c_1 \dot{y}_1(t) + c_2 \dot{y}_2(t).$$
 (18)

969 In the main paper, for simplicity and notation convenience, we use w instead of w_g to describe the 970 parameters and goal of ProDMPs. 971

972 **Trajectory's Initial Condition** The coefficients c_1 and c_2 are solutions to the initial value problem 973 defined by Eqs.(17)(18). Assuming the trajectory starts at time t_b with position y_b and velocity \dot{y}_b , 974 we denote the values of the complementary functions and their derivatives at the condition time t_b 975 as y_{1b} , y_{2b} , \dot{y}_{1b} and \dot{y}_{2b} . Similarly, the values of the position and velocity basis functions at t_b are 976 denoted as Φ_b and Φ_b respectively. Using these notations, c_1 and c_2 are computed as:

$$\begin{bmatrix} c_1 \\ c_2 \end{bmatrix} = \begin{bmatrix} \frac{\dot{y}_{2_b} y_b - y_{2_b} \dot{y}_b}{y_{1_b} \dot{y}_{2_b} - y_{2_b} \dot{y}_{1_b}} + \frac{y_{2_b} \dot{\Phi}_b^{\mathsf{T}} - \dot{y}_{2_b} \Phi_b^{\mathsf{T}}}{y_{1_b} \dot{y}_{2_b} - y_{2_b} \dot{y}_{1_b}} \boldsymbol{w}_g \\ \frac{y_{1_b} \dot{y}_b - \dot{y}_{1_b} y_b}{y_{1_b} \dot{y}_{2_b} - y_{2_b} \dot{y}_{1_b}} + \frac{\dot{y}_{1_b} \Phi_b^{\mathsf{T}} - y_{1_b} \dot{\Phi}_b^{\mathsf{T}}}{y_{1_b} \dot{y}_{2_b} - y_{2_b} \dot{y}_{1_b}} \boldsymbol{w}_g \end{bmatrix}.$$
(19)

Set Goal Convergence Relative to Initial Condition The goal attractor g in the ProDMPs framework represents an asymptotic convergence point for the dynamical system as $t \to \infty$, typically defined as an absolute coordinate. However, the goal term can also be modeled relative to the initial position y_b . In this approach, the relative goal g_{rel} is predicted, and its absolute counterpart is computed as $g_{abs} = g_{rel} + y_b$. This approach is particularly useful for predicting the goal in the coordinate system relative to a node's starting position. Since we aim to achieve a translationequivariant approach (where absolute node positions are encoded as relative edge features between nodes), predicting relative goal positions aligns well with this design principle.

991 В ARCHITECTURE AND METHOD DETAILS 992

This section offers detailed insights into our methodology and the architectural decisions guiding our 994 approach.

996 **B.1** GRAPH ENCODINGS

In processing the initial graph \mathcal{G}_{*,T^c} , we create edges between the mesh and the collider based on 998 a radius graph. Specifically, we connect mesh and collider nodes for *Deformable Plate* and *Tissue* 999 Manipulation if their euclidean distance is smaller than 0.3. In the *Tissue Manipulation* task, the 1000 collider is given as a single node which is connected to the tip of the tissue. It marks the grasping 1001 point of a gripper. In the *Planar Bending* task, we add an additional node feature to the nodes which 1002 get directly influenced by the external force. Therefore, no collider is used in this task. For the 1003 Falling Teddy Bear and the Mixed Objects Falling task, we implicitly model the ground as a collider 1004 by adding the current z position of every node to its node features (Sanchez-Gonzalez et al., 2018). 1005 This quantity gets updated for the step-based methods. 1006

B.2 PRODMP DETAILS

Initialization of ProDMPs necessitates node velocities for the anchor time step T^c . We employ a 1009 linear approximation, leveraging data from the previous time step $T^c - 1$. 1010

1011 Similar to the relative encoding of node positions in the MPN, we employ a technique in ProDMP 1012 to derive relative trajectories. Initially, we integrate a relative goal position as part of the node 1013 weights w_v . Utilizing this approach, trajectories commence from the origin and traverse towards their 1014 respective relative goals. Subsequently, we adjust all positions by the initial position. This strategy fosters model generalization across various nodes. 1015

- 1016 The parameter τ , as described in Equation 7, is learned globally across all tasks using a compact 1017 MLP. The model's final layer employs a scaled sigmoid function for parameter estimation.
- 1018 1019

1008

983

984

985

986

987

988

989 990

993

995

997

С EXPERIMENTAL PROTOCOL

1020 1021

In order to promote reproducibility, we provide details of our experimental methodology. Table 1 1022 presents the hyperparameters used in our experiments. For a comprehensive description of the 1023 creation of all datasets, please refer to Appendix D. 1024

The training took place on an NVIDIA A100 GPU, with each method given the same computation 1025 budget of 48 hours, except for the *Planar Bending* task, where the computation budget was set to 24 hours. Consequently, the number of epochs varied, as the batching differed significantly between the trajectory-based method M3GN and the step-based MGN. We adapted the batchsize of the step-based methods in order to use the GPU memory efficiently. M3GN is always trained on one full trajectory per batch. Here, the whole context is processed in parallel and the remaining trajectory is predicted and compared to the ground truth.

We conducted a multi-staged grid-based hyperparameter search for the learning rate, input noise, and other hyperparameters as the latent task description dimension. In general, we optimized all methods on all tasks separately, however, we noticed that over different tasks and methods some parameters had the same best configuration. We did not use the test data for this, but tuned all hyperparameters on a separate validation split. This split was also used to determine the best epoch checkpoint to mitigate any overfitting effects.

1037 In the end, all methods worked well with a learning rate of 5.0×10^{-4} except in the *Planar Bending* 1038 task. Here, our hyperparameter optimization indicated that the trajectory based methods benefit from 1039 a smaller learning rate of 1.0×10^{-5} . For MGN, we experimented with different input noise scales. 1040 Notably, for the *Deformable Plate* and the *Planar Bending* task, a smaller noise scale improved 1041 performance significantly. In the falling objects tasks, we also explored second-order predictions, 1042 such as node accelerations, instead of velocity predictions. Following the approach in Pfaff et al. 1043 (2021), we adjusted the labels accordingly and conducted preliminary evaluations. However, since direct velocity predictions yielded superior results, we opted for them as our final approach, as 1044 presented in the main paper. 1045

- 1046
- 1047

1048 C.1 EGNO TRAINING 1049

For the Equivariant Graph Neural Operator (EGNO) method, we used the original code from Xu et al.
(2024) for the model implementation. Since EGNO can only predict for a fixed next horizon, we
cut the remaining prediction when using a later anchor time step. This is done during training and
evaluation.

1054 1055 C.2 MGN TRAINING

We mainly follow Pfaff et al. (2021) for the training of the MGN baseline. The only difference is the incorporation of current and historic velocity node features. Pfaff et al. (2021) consider this in their experiments but they show in their experiment suite that it does not improve the results and can lead to overfitting. This is different to our results. For us, on all tasks except *Mixed Objects Falling*, adding the current and historic velocities of nodes improves the results. We follow the Gaussian random walk noise injection for the velocity features from Sanchez-Gonzalez et al. (2020).

- 1062
- 1063 1064

D DATASETS AND PREPROCESSING INFORMATION

In this section, we give detailed information about the datasets we used. We report a general overview of all datasets in Table 2. Here each dataset is abbreviated for brevity as the following:

Table 2 lists in detail the datasets used in the paper. Each dataset is abbreviated for brevity and explained as follows:

- **PB.**: Planar Bending
- **DP.**: Deformable Plate
 - TM.: Tissue Manipulation
 - FTB.: Falling Teddy Bear
 - MOF.: Mixed Objects Falling
- 1075 1076

1074

1070

- 1077 D.1 PLANAR BENDING
- 1079 We select 9 different Young's modulus ranging between 10 and 1000 from a very deformable to an almost stiff sheet. Then, per material, we compute 100 simulations using Abaqus where the positions

1081		
1082	Parameter	Value
1083	Node feature dimension	128
1084	Latent task description dimension	64
1085	Decoder hidden dimension	128
1086	Message passing blocks	15
1087	Message passing blocks (EGNO)	5
1007	GNN Aggregation function	Mean
1000	GNN Activation function	Leaky ReLU
1009	M3GN Context Aggregation method	Max Aggr.
1090	M3GN Latent Node Aggregation method	No Aggr.
1091	Learning rate	5.0×10^{-4}
1092	Learning rate (Plan. Bend. MP-based methods)	$1.0 imes 10^{-5}$
1093	Number of ProDMP basis functions	30
1094	ProDMP $ au$	learned
1095		range: [0.3,3.0]
1096	ProDMP Relative start position	True
1097	MGN input mesh noise	0.01
1098	MGN input mesh noise (Def. Plate)	0.001
1099	MGN input mesh noise (<i>Plan. Bend.</i>)	0.0001
1100	MGN history length (all tasks except <i>Mixed Objects Fall</i>	2
1101	MGN history length (<i>Mixed Objects Fall</i>)	0
1101	M3GN history length (all tasks except <i>Tiss. Man.</i> and <i>Plan. Bend. (OOD)</i>)	1
1102	M3GN history length (<i>Tiss. Man.</i> and <i>Plan. Bend.</i> (<i>OOD</i>))	0
1103	Minimum/Maximum Train Context Size (<i>Plan. Bend.</i>)	2/15
1104	Minimum/Maximum Train Context Size (Deformable Plate)	2/15
1105	Minimum/Maximum Train Context Size (Tissue Manipulation)	2/40
1106	Minimum/Maximum Train Context Size (Falling Teddy Bear)	10 / 50
1107	Minimum/Maximum Train Context Size (<i>Mixed Objects Fall</i>)	10/50
1108	Threshold to create collider-mesh edge	0.3
1109		

Table 1: Table listing the hyperparameters and configurations of the experiments

Table 2: Table listing the datasets and their configurations

112	Name	Train/Val/Test Splits	Number of steps	Number of Nodes	Collider interaction
113	PB.	630/135/135	50	225	External Force
114	DP.	675/135/135	39	81	Rigid Collider
115	TM.	600/120/120	100	361	Grasping point
116	FTB.	700/150/150	200	304	Boundary Condition
117	MOF.	1800/360/360	200	up to 350	Boundary Condition
1118				1	5

1119 1120

1121

1110

1111

1080

of the two acting forces are randomized. The boundary nodes of the sheet are kept in place. From every material configuration, we take 70 simulations for training and 15 simulations for validation 1122 and testing respectively. For the out-of-distribution task, we only trained on Young's modulus ranging 1123 between 60 and 500, while testing on Young's modulus values 10, 30, 750, and 1000. 1124

1125

1126 D.2 DEFORMABLE PLATE 1127

1128 The original task was introduced in Linkerhägner et al. (2023), generated using Simulation Open 1129 Framework Architecture (SOFA) (Faure et al., 2012). It uses 3 different Poisson's ratios and 9 1130 different trapezoidal meshes. We increase the difficulty of this dataset by introducing more complex initial starting conditions. This is done by selecting a random Poisson's ratio, simulating for 11 steps, 1131 and then switching to another Poisson's ratio. Then, the simulation continues for 39 steps. The first 1132 11 steps are then discarded and step 12 is then the initial step for the dataset (and is referred to step 0 1133 throughout the paper).

Figure 9: Six objects used in the Mixed Objects Falling task. From left to right: Teddy Bear, Bunny, Gummy Bear, Gummy Worm 1, Gummy Worm 2, and Traffic Cone.

D.3 TISSUE MANIPULATION

We use the original task introduced in Linkerhägner et al. (2023) without alterations. This dataset was also generated using SOFA (Faure et al., 2012).

D.4 FALLING TEDDY BEAR

Each trajectory of the dataset was created by choosing an angle from $[0^{\circ}, 360^{\circ}]$ for the first time step. To vary the material properties, we randomly select one possible combination of the Young's modulus and Poisson's ratio from the sets generated by np.linspace $(1 \times 10^5, 1 \times 10^6, 1000)$ and np.linspace (0.0, 0.499, 100), respectively. This dataset is generated using NVIDIA Isaac Sim (NVIDIA, 2022a), which utilizes PhysX 5.0 (NVIDIA, 2022b) to simulate tetrahedral meshes based on initial CAD models.

D.5 MIXED OBJECTS FALLING

The simulation uses the setup from Falling Teddy Bear. In addition to the Teddy, we include other objects to encourage diversity. In total, there are six different objects presented in the dataset. We report an image of their high-resolution meshes in Figure 9. From every object, we use 300 simulations for the training split and 60 simulations for the test and validation split respectively.



Figure 11: Log-scale MSE over full rollouts for the *Planar Bending* (Left) and *Deformable Plate* (Right) tasks for different context and node aggregation methods. The node-level maximum context aggregation of M3GN performs best for *Planar Bending*, while all methods work roughly equally well on the *Deformable Plate* task.

E ADDITIONAL RESULTS

1227 1228 1229

1225 1226

E.1 HYPERPARAMETER OPTIMIZATION

We observed that the history inclusion of previous velocities has a big impact on the result of the simulation, depending on the task. To obtain optimal performance, we did an hyperparameter optimization on the validation split comparing history features. The results for M3GN andMGN are given in Figure 10.

1235

1237

1236 E.2 EVALUATIONS.

Aggregation Figure 11 shows results for different context aggregation schemes, comparing global and node-level aggregation, and a transformer-based aggregation over the context set to a simple maximum operator. The node-level maximum context aggregation of M3GN works best for *Planar Bending*. For *Deformable Plate*, there is no significant difference between node-level and global contexts, or between maximum and transformer-based aggregations.



Figure 12: Log-scale MSE over the first 10 steps after the anchor time step for the *Deformable Plate* (Left) task, the *Planar Bending* (Center) task, and the *Planar Bending* task with out-of-distribution material properties for the test trajectories. M3GN steadily improves its performance when provided with additional context information and a later anchor time step. When evaluating this metric, we also outperform MGN Oracle on the *Deformable plate* task.



Figure 13: Log-scale MSE over the first 10 steps after the anchor time step for the *Tissue Manipulation* (Left), *Falling Teddy Bear* (Middle) and *Mixed Objects Falling* (Right) tasks for different methods. M3GN significantly outperforms all baselines on *Tissue Manipulation* when provided with a context size larger than 3. For the other tasks, M3GN significantly improves over both MGN variants, but does not benefit much from additional context information.

1283

1294

1295

Figure 12 and Figure 13 show the performance of M3GN compared to the baselines evaluated on the 10 steps MSE. Here, instead of the whole trajectory, only the first 10 steps after the anchor time step are used to evalute the MSE.

MSE over time To gain better insights into the rollout stability of the model predictions, we report the Mean Squared Error (MSE) over timesteps in Figure 14, 15, 16, 17, 18, and Figure 19. Overall, our model M3GN demonstrates great robustness against error accumulation, benefiting from the inherent trajectory representation provided by the ProDMP method. MGN works well on the *Tissue Manipulation* task, but fails to incorporate the correct context information on other tasks. EGNO performs in general worse except on the *Planar Bending* tasks.

1284 E.3 VISUALIZATIONS.

In Figure 20, we include a latent space visualization of the Planar Bending task, where simulations with 9 different Young's Modulus values are clustered according to their material properties. The
 t-SNE projection of the 64-dimensional latent node vectors demonstrates clear clustering, indicating that the model effectively captures and differentiates material characteristics based on learned task representations.

We further provide additional visualizations for M3GN, MGN and MGN (Oracle) for exemplary
 simulations of all tasks. Each visualization shows the same simulated trajectory for different time
 steps (columns) and different methods (rows).

- Figure 21 shows a simulation of the *Planar Bending* task for a context set size of 5.
- Figure 22 visualizes the *Deformable Plate* task for a context set size of 6.







Figure 20: This figure shows a latent space visualization of the Planar Bending task for trajectories with 9 different Young's Modulus values, using a context size of 10. Each dot represents a 64dimensional latent node vector projected to 2D using the t-SNE algorithm (van der Maaten & Hinton, 2008). Dots of the same color correspond to latent node descriptions for the same task, each simulated with a unique Young's Modulus. The visualization reveals distinct clustering in the latent space, with similar material properties grouped closer together, highlighting the relationship between material characteristics and the learned task representations. To improve clarity, points corresponding to nodes on the plate's edge were excluded, as their constant boundary condition resulted in unvarying latent descriptions.



Figure 21: Simulation over time of an exemplary test trajectory from the **Planar Bending** task by M3GN, MGN, and MGN with oracle information. The **context set size** is set to 5. All visualizations show the colored **predicted mesh**, a **collider or floor**, and a **wireframe** of the ground-truth simulation. M3GN can accurately predict the correct material properties, resulting in a highly accurate simulation.



Figure 22: Simulation over time of an exemplary test trajectory from the Deformable Plate task by M3GN, MGN, and MGN with oracle information. The context set size is set to 6. All visualizations show the colored predicted mesh, a collider or floor, and a wireframe of the ground-truth simulation. M3GN can accurately predict the correct material properties, resulting in a highly accurate simulation.



Figure 23: Simulation over time of an exemplary test trajectory from the **Tissue Manipulation** task by M3GN, MGN, and MGN with oracle information. The **context set size** is set to 6. All visualizations show the colored **predicted mesh**, a **collider or floor**, and a **wireframe** of the ground-truth simulation. All methods can solve the task, however MGN is drifting a tiny bit to the left over time.



Figure 24: Simulation over time of an exemplary test trajectory from the **Falling Teddy Bear** task by M3GN, MGN, and MGN with oracle information. The **context set size** is set to 20. All visualizations show the colored **predicted mesh**, a **collider or floor**, and a **wireframe** of the ground-truth simulation. M3GN significantly outperforms both step-based baselines.



Figure 25: Simulation over time of a bunny from the Mixed Objects Fall task by M3GN, MGN, and MGN with oracle information. The context set size is set to 20. All visualizations show the colored predicted mesh, a collider or floor, and a wireframe of the ground-truth simulation. M3GN significantly outperforms both step-based baselines.



Figure 26: Simulation over time of a pylon from the Mixed Objects Fall task by M3GN, MGN, and MGN with oracle information. The context set size is set to 20. All visualizations show the 1615 colored predicted mesh, a collider or floor, and a wireframe of the ground-truth simulation. M3GN 1616 significantly outperforms both step-based baselines. The simulation generated by MGN is severly 1617 affected by drift. 1618

1614

1587

1588

1589