AdaptiGraph: Material-Adaptive Graph-Based Neural Dynamics for Robotic Manipulation

Kaifeng Zhang*, Baoyu Li*, Kris Hauser, Yunzhu Li
University of Illinois Urbana-Champaign

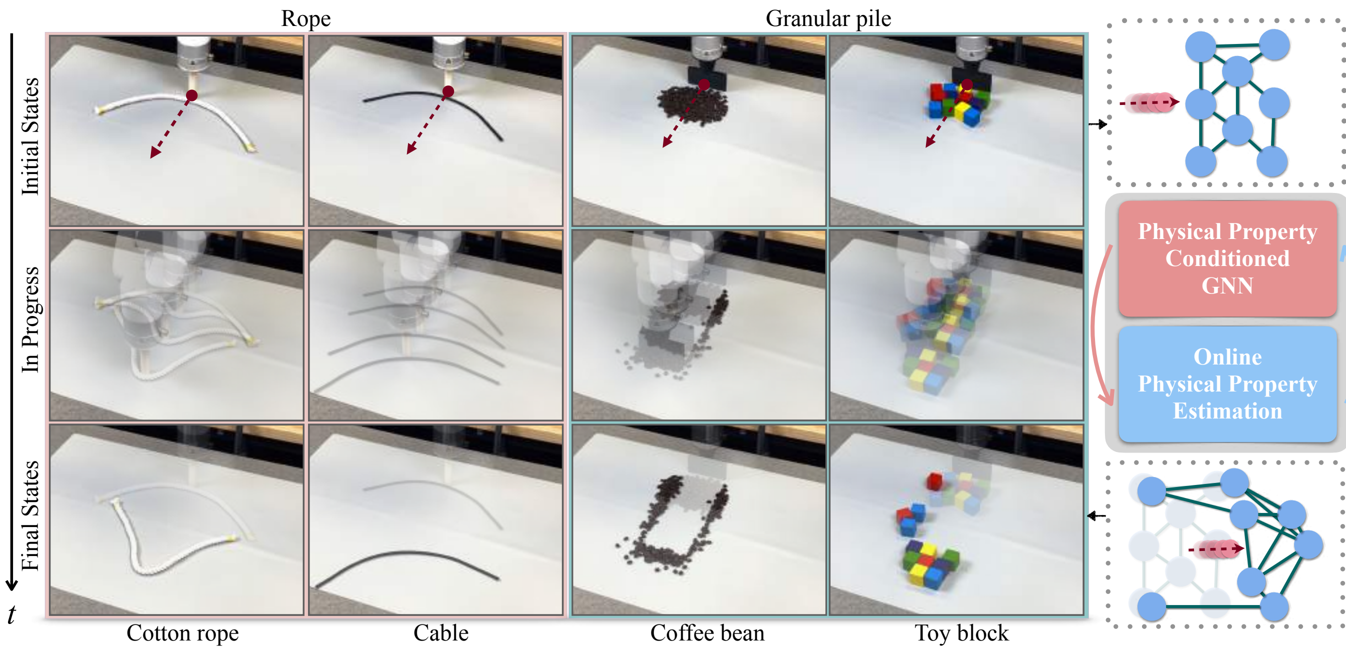

Fig. 1: **Motivation.** Objects made from different materials can exhibit distinct behaviors under interaction. Even within the same object category, varying physical parameters like stiffness can lead to different behaviors. Examples shown here include handling cotton rope and cable, as well as arranging granular piles such as coffee beans and toy blocks. Although the initial configuration and action are the same, different physical parameters result in distinct final states, necessitating the need for online adaptation for effective manipulation. To this end, we introduce **AdaptiGraph**, a unified graph-based neural dynamics framework for real-time modeling and control of various materials with unknown physical properties. **AdaptiGraph** integrates a **physical property-conditioned dynamics model** with **online physical property estimation**. Our framework enables robots to adaptively manipulate diverse objects with varying physical properties and dynamics.

***Abstract*—This paper introduces AdaptiGraph, a learning-based dynamics modeling approach that enables robots to predict, adapt to, and control a wide array of challenging deformable materials with unknown physical properties. AdaptiGraph leverages the highly flexible graph-based neural dynamics (GBND) framework, which represents material bits as particles and employs a graph neural network (GNN) to predict particle motion. Its key innovation is a unified physical property-conditioned GBND model capable of predicting the motions of diverse materials with varying physical properties without retraining. Upon encountering new materials during online deployment, AdaptiGraph utilizes a physical property optimization process for a few-shot adaptation of the model, enhancing its fit to the observed interaction data. The adapted models can precisely simulate the dynamics and predict the motion of various deformable materials, such as ropes, granular media, rigid boxes, and cloth, while adapting to different physical properties, including stiffness, granular size, and center of pressure. On prediction and manipulation tasks involving a diverse set of real-world deformable objects, our method exhibits superior prediction accuracy and task proficiency over non-material-conditioned and non-adaptive models.**

*Denotes equal contribution.

I. Introduction

Learning predictive models, also known as system identification, is a crucial component of many robotic tasks. Whereas classical methods rely on the explicit parameterization of the system state and struggle with systems that have high degrees of freedom, a significant body of work over the last decade has attempted to learn models directly from visual observations. Prior approaches have learned predictive models based on pixels [1, 4] or latent representations of images [2, 3]. However, such representations often overlook the structure of the environment and do not generalize well across different camera poses, object poses, robots, object sizes, and object shapes. Recently, a series of studies have employed Graph Neural Networks (GNN) to model environments as 3D particles and their pairwise interactions [5–8]. A graph representation has proven effective in capturing relational bias and predicting complex motions of deformable objects, but prior works typically only focus on a single material and would require extensive training to model an object of new material or with unknown physical properties. Hence, it is

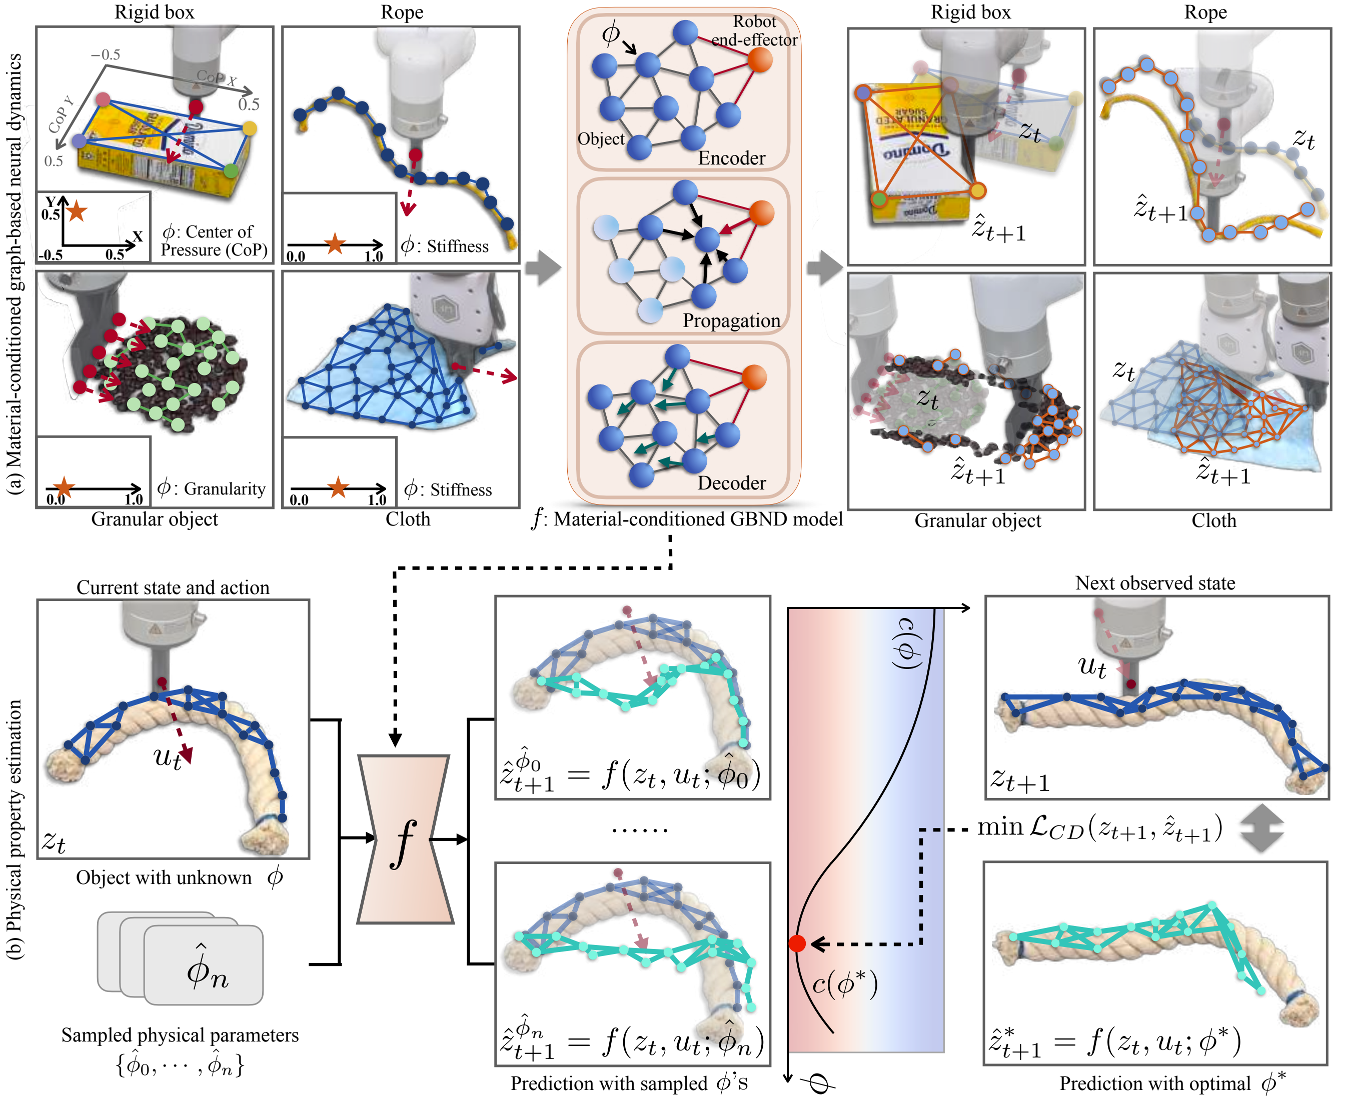

Fig. 2: **Overview of proposed framework: AdaptiGraph. (a)** Our graph-based dynamics model f is conditioned on the discrete material type and continuous physical parameters ϕ. ϕ is encoded as the node features, which will be propagated and updated in the model training process. Our model can accurately predict the future state $\hat{z}_{t+1}$ for a variety of objects with different physical properties. **(b)** Our framework performs physical property estimation for few-shot adaptation. This is achieved through an inverse optimization process to estimate the optimal physical parameters as predicted by the learned dynamics model f. The optimal physical parameter ϕ^* is identified by minimizing the cost function, which is defined as the Chamfer Distance between the predicted graph state and the actual future graph state.

an important challenge to provide such graph-based models to adapt to objects and tasks involving diverse materials and varying physical properties, such as manipulating ropes with different stiffness and granular media with different granularity.

In this work, we present a unified framework for modeling the dynamics of objects with different materials and physical properties. In addition to classifying objects into discrete material types such as rigid objects, ropes, etc., we further consider a range of intra-class physical property variations in each material type. We propose to encode this variation using a continuous variable which we call the physical property variable, and integrate the variable into a Graph-Based Neural Dynamics (GBND) framework (Fig. 1). The physical property variable indicates the important intrinsic properties of each material category, including stiffness for deformable objects and the center of pressure position for rigid objects. By encoding the material type and physical property variables into particles in the graph, the model learns material-specific dynamic functions that predict different physical behaviors for objects with different physical properties. We then employ a test-time adaptation method to reason about the physical properties of novel objects. Specifically, the robot actively interacts with the novel object, observes its response, and estimates its physical properties to optimize the model's fit to the observed reactions. The estimation is performed in a few-shot manner and can be directly applied to planning and trajectory optimization for downstream manipulation tasks.

II. METHODS

A. Problem Formulation

Our aim is to learn a dynamics model, f, that is conditioned on the material type M and continuous physical property variable ϕ, and develop a test-time few-shot adaptation scheme to infer the physical property variable for unseen objects. Specifically, the dynamics model predicts how the environment

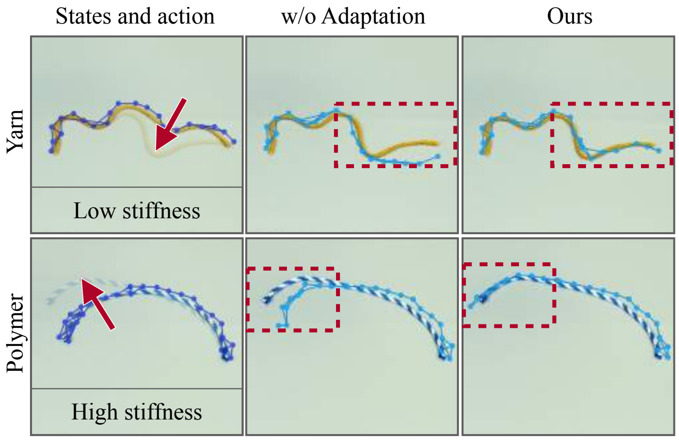

(a) Rope

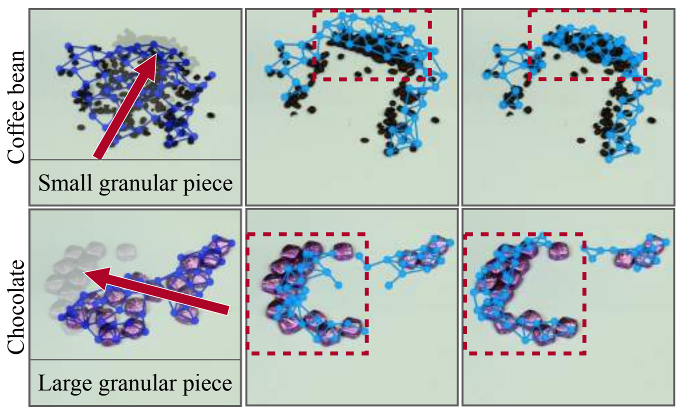

(b) Granular object

Fig. 3: **Qualitative results on dynamics prediction:** We conduct qualitative comparisons to assess the performance of our method against the baseline of a GNN without adaptation. The results, delineated by red dashed boxes, demonstrate that our approach surpasses the baseline in accurately capturing the variations in dynamics that arise due to physical properties.

will change if the robot applies a given action:

$$\hat{z}_{t+1} = f(z_t, u_t; \phi, M), \tag{1}$$

where M indicates the material type (e.g., rigid, granular, rope, cloth), ϕ indicates material-specific physical property variables, and u_t, z_t, z_{t+1} are the robot action, current environment state at time t, and the next state at time $t+1$, respectively. In our approach, we train the dynamics model to minimize the accumulated future prediction loss.

By conditioning on M and ϕ, the model learns to predict material-dependent physical behaviors, based on which we can perform physical property estimation through the following optimization problem:

$$\phi^* = \arg\min_{\phi} \sum_{t=1}^{T} \text{cost}(\hat{z}_{t+1}, z_{t+1}), \tag{2}$$

where T is the iteration number indicating the number of interactions with the unseen object, and $\text{cost}(\cdot, \cdot)$ is the cost function measuring the discrepancy between the predicted future state $\hat{z}_{t+1}$ and the observed state z_{t+1}.

Figure 2 shows the overall framework of AdaptiGraph.

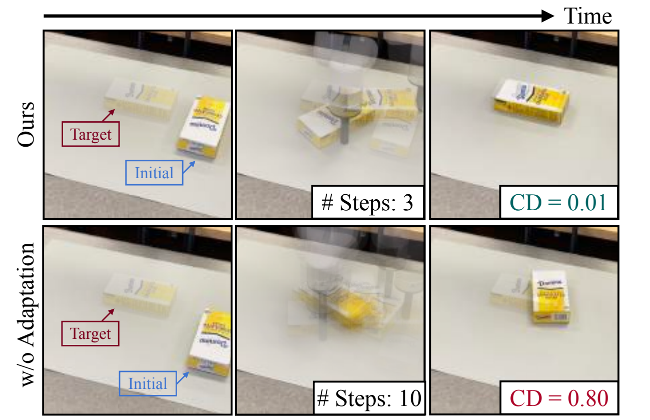

(a) Rigid box pushing

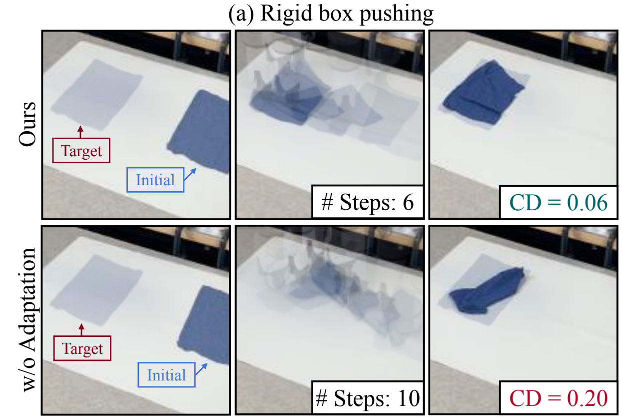

(b) Cloth relocating

Fig. 4: **Qualitative results on closed-loop feedback planning:** We present a qualitative comparison of MPC performance by contrasting our method with the baseline method. Visualizations shown here demonstrate that our method effectively achieves the target configuration within limited steps.

III. Experimental Results

A. Forward Dynamics Prediction

Fig. 3 shows the qualitative comparisons between our material-conditioned GBND model and the baseline method *Ours w/o Adaptation*, which is an ablated version of our material-adaptive model by using only the mean physical property variable $\bar{\phi}$ as input in deployment. The comparisons reveal that, with estimated physical property, the model's prediction matches the interaction outcome more accurately. For instance, in the rope scenario, the baseline model's prediction fails to capture both the below-average stiffness of the yarn object and the above-average stiffness of a polymer rope. In contrast, our method successfully accounts for variations in their motions, exhibiting more precise forecasts of unusual behaviors. Likewise, our model surpasses the baseline in scenarios involving materials with extreme physical properties, such as rigid boxes that differ in center of pressure, granular materials of various sizes, and clothes of differing stiffness.

B. Model-Based Planning

We further show that integrating our material-conditioned GBND model and physical property adaptation into an MPC framework facilitates a series of robotic manipulation tasks. As illustrated in Figure 4, our approach enables more effective and efficient planning compared to the baseline method.

References

[1] Chelsea Finn and Sergey Levine. Deep visual foresight for planning robot motion. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2786–2793. IEEE, 2017. 1

[2] Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*, 2019. 1

[3] Danijar Hafner, Timothy Lillicrap, Mohammad Norouzi, and Jimmy Ba. Mastering atari with discrete world models. *arXiv preprint arXiv:2010.02193*, 2020. 1

[4] Ryan Hoque, Daniel Seita, Ashwin Balakrishna, Aditya Ganapathi, Ajay Kumar Tanwani, Nawid Jamali, Katsu Yamane, Soshi Iba, and Ken Goldberg. Visuospatial foresight for multi-step, multi-task fabric manipulation. *arXiv preprint arXiv:2003.09044*, 2020. 1

[5] Yunzhu Li, Jiajun Wu, Jun-Yan Zhu, Joshua B Tenenbaum, Antonio Torralba, and Russ Tedrake. Propagation networks for model-based control under partial observation. In *ICRA*, 2019. 1

[6] Haochen Shi, Huazhe Xu, Zhiao Huang, Yunzhu Li, and Jiajun Wu. Robocraft: Learning to see, simulate, and shape elasto-plastic objects with graph networks. *arXiv preprint arXiv:2205.02909*, 2022.

[7] Haochen Shi, Huazhe Xu, Samuel Clarke, Yunzhu Li, and Jiajun Wu. Robocook: Long-horizon elasto-plastic object manipulation with diverse tools. *arXiv preprint arXiv:2306.14447*, 2023.

[8] Yixuan Wang, Yunzhu Li, Katherine Driggs-Campbell, Li Fei-Fei, and Jiajun Wu. Dynamic-Resolution Model Learning for Object Pile Manipulation. In *Proceedings of Robotics: Science and Systems*, Daegu, Republic of Korea, July 2023. doi: 10.15607/RSS.2023.XIX.047. 1