Elastic Context: Encoding Elasticity for Data-driven Models of Textiles

Alberta Longhini[1], Marco Moletta[1], Alfredo Reichlin[1], Michael C. Welle[1], Alexander Kravberg[1],
Yufei Wang[2], David Held[2], Zackory Erickson[2], and Danica Kragic[1]

***Abstract*— Physical interaction with textiles, such as assistive dressing or household tasks, requires advanced dexterous skills. The complexity of textile behavior during contact interactions is influenced by the material properties of the yarn and by the textile's construction technique, which are often unknown in real-world settings. Moreover, identification of physical properties of textiles through sensing commonly available on robotic platforms remains an open problem. To address this, we introduce Elastic Context (EC), a method to encode the elasticity of textiles using stress-strain curves adapted from textile engineering for robotic applications. We employ EC to learn a data-driven model of textiles that captures the variations of elastic force responses of textiles in contact-rich scenarios. Moreover, we examine the effect of EC dimension on accurate force modeling of real-world non-linear elastic responses during contact interactions.**

I. Introduction and Background

Manipulation of deformable objects such as textiles is common in medical robotics [1], human-robot interaction [2], automation of household tasks [3], assistive dressing [4], [5]. In such robotic tasks, physical properties of textiles such as elasticity, surface friction, and flexibility play a fundamental role on the interaction dynamics involved in the manipulation. An example of how elasticity influences a pulling interaction is shown in Fig. 1. However, textile objects are challenging to manipulate due to complex dynamics and often unknown physical properties such as elasticity, friction, density distribution [6]. Moreover, these properties are difficult to estimate in online robotics manipulation scenarios [7].

We address the problem of encoding the elastic properties of textile objects and learning data-driven models of contact interactions that generalize to variations of elastic properties. In particular, these interactions might come either from robot actions, such as stretching or shearing the object, or they could be due to collisions with external objects, like the body of a person in an assistive dressing scenario. In both cases, we can reason about the interaction dynamics by identifying three different stages. In the first stage (*free* manipulation), no stress is involved during the manipulation as the deformable nature of the textile makes it compliant with the robot's movement or external objects. In the second phase (*stress* manipulation), the textile starts to induce forces on the end-effector and its elastic properties become relevant to characterize the interaction dynamics of the manipulation.

[1]The authors are with the Robotics, Perception and Learning Lab, EECS, at KTH Royal Institute of Technology, Stockholm, Sweden albertal, moletta, alfrei, mwelle, okr, dani@kth.se

[2]The authors are with Carnegie Mellon University, Pittsburgh, USA yufeiw2, dheld, zerickso@andrew.cmu.edu

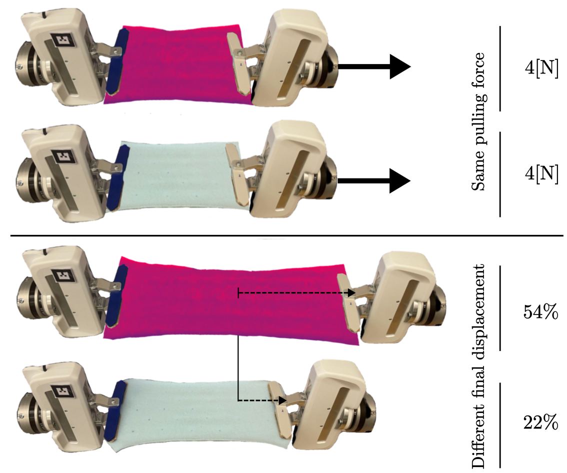

Fig. 1: The role of elasticity in the manipulation of different textile samples, exhibiting different behaviors under the same pulling force.

Finally, after a certain force known as rupture force, the textile breaks. Often in robotic tasks, objects are manipulated within the *free* manipulation phase. Nevertheless, precise modeling of the *stress* manipulation phase can be useful for applications requiring a constraint over the forces exerted on the deformable object or to assure that a specific action is correctly performed [2], [6]. Modeling the elastic behavior of textiles has been addressed in two rather distant communities: textile engineering [8] and computer graphics [9], [10]. The analytical models from these communities are computationally expensive and commonly not applicable in real-time robotic manipulation. Moreover, such models build on parameters measured with high-precision devices under controlled experiments [11], [12]. Today, there are no commonly adopted and annotated datasets for which these parameters are provided, requiring thus alternative strategies to encode elastic properties of textiles.

Our previous work introduced a taxonomy considering the yarn material and the fabric construction technique to understand textile properties [13]. To meet this goal, it relied on physical interactions (pulling and twisting) and force-torque measurements. We build upon this work and encode elastic behaviors of textiles through Elastic Context (EC). To formulate the Elastic Context (EC) we draw inspiration from industrial techniques to characterize textiles, namely the stress-strain curve [14]. These curves can be obtained by recording the forces perceived by a robot when manipulating different textiles.

We experimentally evaluate the EC on synthetic data obtained in the PyBullet simulator [15], demonstrating the ability of our approach to encode the complex elastic response of a wide range of textile. Furthermore, with force measurements collected on two Franka-Emika Panda robots interacting with textiles and an external rigid object, we show that increasing the dimensionality of the EC plays an important role in accurately modeling the complex elastic response of textiles when an interactive scenario is considered.

II. Elastic Context for Data-Driven Models

To capture the diverse range of potential elastic responses of textiles during contact interactions, we utilize a context-informed dynamics model that incorporates our proposed encoding of textile's elastic properties, which we refer to as the *Elastic Context* (EC).

A. Elastic Context

A potential measure of a material's elasticity is the elastic modulus, which evaluates the material's resistance to deformation under stress [16]. Its value can be derived by measuring the slope of the stress-strain curve corresponding to a specific material or textile sample. The stress σ is defined as the deformation force F [N] acting on the cross-sectional area A [m^2] of the sample, while the strain ε corresponds to the percentage of displacement Δl of the sample with respect to its original length l_0. Using these quantities, the elastic modulus is calculated as $e = \sigma/\varepsilon$.

Fig. 2 presents stress-strain curves for two real-world textile samples. The values were obtained from force-feedback readings of the dual-arm robotic setup shown in Fig. 3. The robots were pulling the samples with $l_0 = 0.18$ m and $A = 0.18 \times 10^{-3}$ m^2 until a stress $\sigma_{max} = 30$ KPa was reached. Fig. 2 indicates how encoding elastic properties of textile with the elastic modulus can only describe small displacements and linear behaviors. We observe that a linear approximation can accurately describe the rigid sample (blue line) but loses accuracy for the elastic sample (red line), which increases its rigidity with increasing stress.

To overcome the limitation of the elastic modulus, we define EC as the combination of elastic modules of a given textile evaluated at n_{EC} equidistant points between 0 and σ_{max} on the stress-strain curve, where n_{EC} determines the dimension of the EC. We thus represent the EC as a vector $EC = [e_1, ..., e_{n_{EC}}] \in \mathbb{R}^{n_{EC}}$, where $e_i = \sigma_i/\varepsilon_i$ is the elastic modulus evaluated from the textile stress-strain curve at the corresponding stress σ_i. Since both stress and strain measurements are normalized by the size of the sample, the EC is a consistent definition of elasticity that is comparable between textiles of different sizes as long as σ_{max} is fixed.

B. Context Informed Dynamics Model

We propose to leverage the EC along a data-driven dynamic model of the force response of the textile during a contact interaction. We assume the contact interaction to happen between the textile and an external rigid object, and we aim at predicting the force response of the textile from the

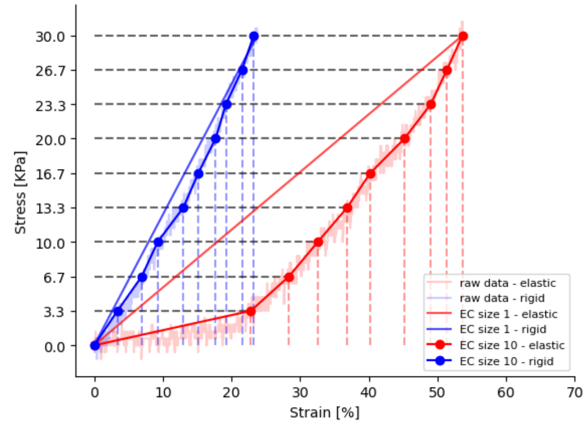

Fig. 2: Stress-strain curves of a rigid (blue) and an elastic (red) textile sample: with $n_{EC} = 1$ we recover the elastic modulus, describing a linear elastic behavior that accurately represents the rigid sample (blue) but not the elastic one (red); with a higher-dimensional EC ($n_{EC} = 10$), we are able to better describe the non-linear behavior of the elastic sample.

moment of first contact with the rigid object. In particular, at time t, a dual arm robotic manipulator grasping the textile, which is in contact with an external object, applies an action $a^t = \Delta x^t_{\text{gripper}} \in \mathbb{R}^3$, resulting in new positions x^{t+1}_{gripper} of the gripper and a textile specific force F^{t+1}_k perceived at the end-effectors. Our goal is to learn a model that leverages the EC to predict F^{t+1}_k, given the action a^t performed by the end-effector.

For the simulation experiments, we model the force dynamics through a Graph Neural Network [17], [18] as we can easily access the ground truth graph representing the textile. We can thus formulate the problem of learning the force dynamics of textiles as learning the parameters θ of a Graph Neural Network (GNN) Λ_θ. We employ a standard message-passing architecture for the GNN model Λ_θ [19]. Specifically, let $G_k = (V_k, E_k)$ be a graph representing a textile sample $k \in K$, where V is the set of nodes, E the set of edges and K is the set of all possible elastic samples. We define the features of each node $v \in V_k$ by their position $x^t_v \in \mathbb{R}^3$ in the Euclidean space at time step t. The edges $e \in E_k$ instead describe the elastic relation between two connected nodes. We propose to define the features of the edges as $EC_k \in \mathbb{R}^{n_{EC}}$ evaluated for the specific elastic textile k. EC is therefore a feature shared among all the edges encoding the elastic property of the textile into G_k. Thus, the input graph is composed by the node's features, the action concatendated to the features of the grasped nodes, and the edge features, represented by the EC. The output of the model is the prediction of F^{t+1}_k, which can be useful for applications requiring a constraint over the forces exerted on the deformable object, e.g., robotic dressing or bathing assistance. The overall model Λ can be learned using a dataset $\mathscr{D} = \{(G^t_k, F^{t+1}_k, a^t_k, EC_k)\}_{\forall k \in K}$, optimising the parameters θ using a supervised loss on the prediction of the nodes position and the force exerted at the grasp-nodes:

$$\mathscr{L} = \mathbb{E}_{\mathscr{D}} \left[\, d(\Lambda_\theta(G_k^t, a_k^t, EC_k), \; F_k^{t+1} \, \right], \tag{1}$$

where d is a measure of the distance between the prediction and the ground-truth of graphs and forces. In our case, d is implemented as the sum of Mean-Squared Error (MSE) of the graph's position and of the force.

For the real-world experiments we rely on a dataset $\mathscr{D} = \{(F_k^{t+1}, a_k^t, EC_k)\}_{k \in K}$ which does not use the ground truth graph G_k^t as it is not easily accessible in the real world. We leave the collection of the real-world ground-truth graphs G_k^t for future work, which could be performed leveraging recent approaches for state estimation and dynamics prediction for cloth [20]–[22]. As we do not have graphs, for modeling the force dynamics we resort to a standard Multi-layer Perception (MLP) which receives as input the concatenation of the action and the EC, and outputs the the prediction of F_k^{t+1}. The MLP is trained as the GNN with a MSE loss evaluated between the model predictions and the ground truth forces.

III. Experimental Evaluation

In this section, we evaluate the performance of the EC in modeling linear and non-linear dynamics of textile objects for robotic manipulation tasks. In particular, we show using simulation that EC with GNNs leads to more accurate force-feedback predictions of unseen elastic textiles. Furthermore, we analyze the role of the dimensionality of EC both in simulation and real-world scenarios, highlighting the importance of increasing the EC dimension (n_{EC}) in presence of non-linear force dynamics.

A. Experimental Setup

Task: We evaluate the performance of EC in predicting forces perceived by a robot when manipulating *unseen* elastic samples. To this end, we devised a two-stage simplified assistive dressing task. In the first stage, a dual-arm robot is tasked to stretch a textile up to a maximum stress $\sigma_{max} = 3 \times 10^4$ Pa, corresponding to a $F_{max} = \sigma_{max} \times A_0$ [N] force perceived at the end-effector for a textile with cross-sectional area A_0 [m^2]. The forces recorded during this interaction are used to recover the EC, following the procedure defined in Section II. In the second stage, the same dual-arm robot is tasked with pulling the sample over a sphere, resembling the head of a person, up to a cumulative gripper displacement of a_{max} [cm] from the instant the textile starts to induce force on the end-effector.

The goal of our model is to perform a *force-forecasting* task, which consists of predicting the force response of the textile from the moment of contact with the object until the goal displacement is reached. An overview of the setup is presented in Fig. 3. We reproduced the aforementioned scenario both in the real world and in the PyBullet simulator [15], [23]. To manipulate the textile in the real-world setting we used two Franka-Emika Panda equipped with Optoforce Force/Torque sensors, while in the simulation we only used free-floating end-effectors and their force sensors.

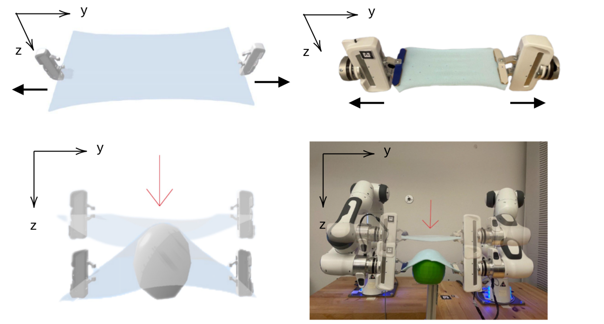

Fig. 3: Simulated (left) and real-world (right) environments to instantiate the simplified assistive dressing task: the y-axis corresponds to the pulling direction used to collect the context, while the z-axis is related to the task execution.

Data Collection: The simulation dataset $\mathscr{D}_{SIM} = \{(G_k^t, F_k^{t+1}, a_k^t, EC_k)\}_{k \in K}$ was collected performing the aforementioned task where we varied for each execution the elastic properties $k \in K$ of the simulated textile. $K = [20, 119]$ was defined empirically by selecting the *elasticity* object parameter to avoid unstable behaviors of the mesh during the collision with the sphere. We uniformly sampled k with a step size of 1 while keeping the *bending* and *damping* properties fixed to 0.1 and 1.5 respectively, obtaining a total of 100 different elastic samples.

For the real-world experiments, where the ground truth elastic properties of the textile are not easily accessible, we collected a dataset $\mathscr{D}_{RW} = \{(F_k^{t+1}, a_k^t, EC_k)\}_{k \in K}$ defining K according to a proxy categorization of textile properties represented by the taxonomy proposed in [13]. We chose 40 different combinations of yarn material to maximize the variance of the elastic responses, while keeping the construction technique fixed to knitted as the one leading to more elastic behaviors.

B. Elastic Context Evaluation

We start by evaluating the role of the Elastic Context in the simulation environment as the ground truth properties of the simulated textile model are easily accessible, allowing us to compare against an oracle baseline. We carry out the evaluation on the *force-forecasting* task, where we compare the force predictions of our proposed GNN informed with the EC to a *baseline* GNN that has no access to the EC information, and to an *Oracle* model that knows the ground truth properties of the textile. Furthermore, we evaluate the effect the dimension of the EC has on simulated force dynamics, where we considered $n_{EC} = \{1, 2, 5\}$.

Table I presents the test MSE of the force predictions evaluated on 20 *unseen* elastic samples. The results shows that the baseline model performs worse than the GNNs that have access to elastic information. Moreover, the models that levarage the EC have comparable performances to the oracle, supporting our hypothesis that EC enables the models to leverage elastic information. These findings are further confirmed by the qualitative results of the *force forecasting* prediction visualized in Fig. 4. The baseline model (Fig, 4a)

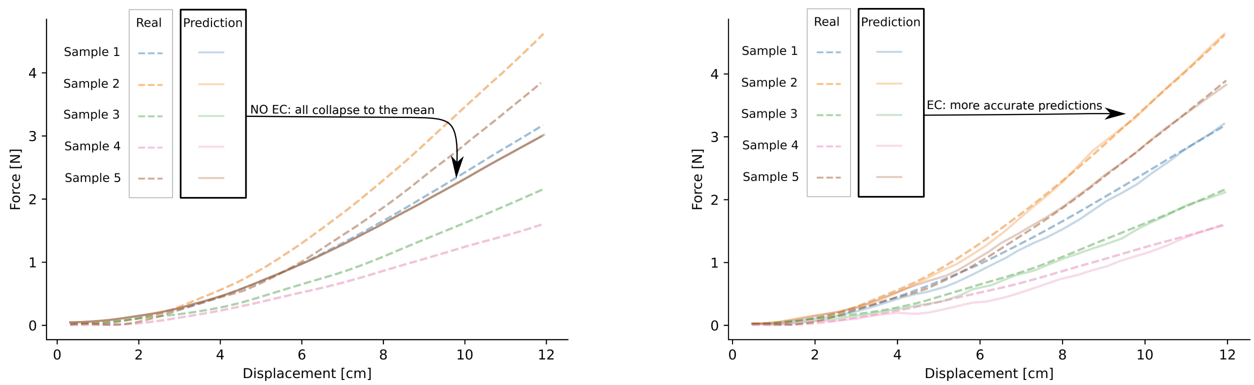

(a) Baseline

(b) GNN + EC ($n_{EC} = 1$)

Fig. 4: Force-forecasting predictions of the baseline model (a) and the GNNs + EC with dimension $n_{EC} = 1$ model (b) evaluated on 5 test elastic samples. The results show that the GNN + EC model generalizes to unseen elastic textiles, leveraging the information provided by the EC.

predicts the same force evolution for all test samples, while the GNN with EC $n_{EC} = 1$ (Fig, 4b) successfully covers a larger spectrum of the test set, improving the accuracy of the predictions of different elastic behaviors. These results highlight the relevance of encoding elastic behavior of textiles, such as the EC, especially for real-world applications requiring a constraint over the forces exerted on the deformable object (e.g., assistive dressing or bathing assistance). We further observe from Table I that increasing the dimension of the EC does not lead to improvements in the model performance, suggesting that $n_{EC} = 1$ is enough to describe the variety of simulated elastic behaviors that assume a linear force dynamics.

TABLE I: Mean and standard deviation of the MSE of the *force-forecasting* task rollout evaluated on 20 different test samples.

Model	FORCE error [N]
Oracle	0.207 ± 0.239
Baseline	3.169 ± 2.434
GNN $n_{EC} = 1$	0.269 ± 0.301
GNN $n_{EC} = 2$	0.235 ± 0.108
GNN $n_{EC} = 5$	0.246 ± 0.247

C. Real-world Evaluation

In this section we showcase the relevance of increasing the dimensionality of EC in the presence of non-linear real-world force dynamics. To fulfill this goal, we trained different variants of the dynamic model by varying the dimensionality of the EC in the range $n_{EC} \in [0,5]$. Fig. 5 presents the force prediction MSE of each model on 8 test elastic samples averaged over 6 randomly-selected seeds, where the test samples were randomly selected from the dataset for each seed. Contrarily to what was observed for the simulation experiments, these results show that increasing the dimensionality of the EC leads to more accurate predictions of real-world non-linear force profiles with the respect to the models trained with an EC with dimension equal to

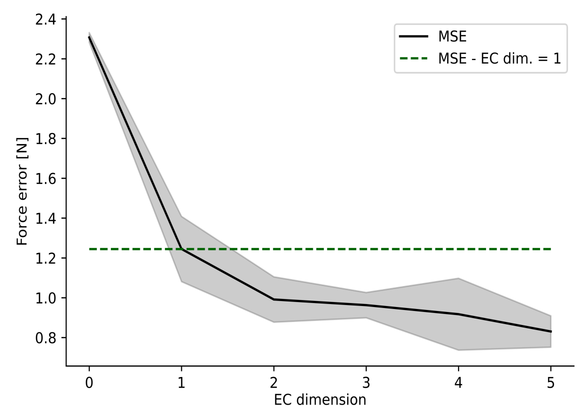

Fig. 5: Role of the EC dimension in real world

1 (dashed line). This outcome highlights the gap between simulated and real-world force dynamics. An interesting future research direction is to explore how this gap hinders the model performance when trained using simulated data and then directly applied to real-world textiles.

IV. Conclusions

In this work, we presented and evaluated *Elastic Context* (EC), an approach to encode elasticity in data-driven models of textiles. We have shown the role of the EC in a *force forecasting* prediction task on both simulated and real world data. The EC can be easily employed in real-world robotic platforms, providing a simple way to understand elastic properties of textiles in scenarios where no labels are provided. Such understanding can be of great importance in assistive robotics and human-robot interaction scenarios, where the manipulation of textile and other deformable objects is considered. In fact, the proposed approach has potential applications in manipulation tasks enabling planning through contacts. Additionally, the work may benefit from the use of tactile sensing and sensor fusion techniques to improve the accuracy of the learned models.

REFERENCES

[1] B. S. Peters, P. R. Armijo, C. Krause, S. A. Choudhury, and D. Oleynikov, "Review of emerging surgical robotic technology," *Surgical endoscopy*, vol. 32, no. 4, pp. 1636–1655, 2018.

[2] Z. Erickson, H. M. Clever, G. Turk, C. K. Liu, and C. C. Kemp, "Deep haptic model predictive control for robot-assisted dressing," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 4437–4444.

[3] A. Verleysen, M. Biondina, and F. Wyffels, "Video dataset of human demonstrations of folding clothing for robotic folding," *The International Journal of Robotics Research*, vol. 39, no. 9, pp. 1031–1036, 2020.

[4] S. D. Klee, B. Q. Ferreira, R. Silva, J. P. Costeira, F. S. Melo, and M. Veloso, "Personalized assistance for dressing users," in *International Conference on Social Robotics*. Springer, 2015, pp. 359–369.

[5] A. Kapusta, Z. Erickson, H. M. Clever, W. Yu, C. K. Liu, G. Turk, and C. C. Kemp, "Personalized collaborative plans for robot-assisted dressing via optimization and simulation," *Autonomous Robots*, vol. 43, no. 8, pp. 2183–2207, 2019.

[6] C. Chi, B. Burchfiel, E. Cousineau, S. Feng, and S. Song, "Iterative residual policy: for goal-conditioned dynamic manipulation of deformable objects," *arXiv preprint arXiv:2203.00663*, 2022.

[7] H. Ha and S. Song, "Flingbot: The unreasonable effectiveness of dynamic manipulation for cloth unfolding," in *Conference on Robot Learning*. PMLR, 2022, pp. 24–33.

[8] S. Poincloux, M. Adda-Bedia, and F. Lechenault, "Geometry and elasticity of a knitted fabric," *Physical Review X*, vol. 8, no. 2, p. 021075, 2018.

[9] D. Clyde, J. Teran, and R. Tamstorf, "Modeling and data-driven parameter estimation for woven fabrics," in *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, 2017, pp. 1–11.

[10] G. Sperl, R. Narain, and C. Wojtan, "Homogenized yarn-level cloth." *ACM Trans. Graph.*, vol. 39, no. 4, p. 48, 2020.

[11] B. Eberhardt, A. Weber, and W. Strasser, "A fast, flexible, particle-system model for cloth draping," *IEEE Computer Graphics and Applications*, vol. 16, no. 5, pp. 52–59, 1996.

[12] M. I. Yousef and G. K. Stylios, "Investigating the challenges of measuring combination mechanics in textile fabrics," *Textile Research Journal*, vol. 88, no. 23, pp. 2741–2754, 2018.

[13] A. Longhini, M. C. Welle, I. Mitsioni, and D. Kragic, "Textile taxonomy and classification using pulling and twisting," *arXiv preprint arXiv:2103.09555*, 2021.

[14] S. Kawabata and M. Niwa, "Fabric performance in clothing and clothing manufacture," *Journal of the Textile Institute*, vol. 80, no. 1, pp. 19–50, 1989.

[15] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," http://pybullet.org, 2016–2021.

[16] V. E. Arriola-Rios, P. Guler, F. Ficuciello, D. Kragic, B. Siciliano, and J. L. Wyatt, "Modeling of deformable objects for robotic manipulation: A tutorial and review," *Frontiers in Robotics and AI*, p. 82, 2020.

[17] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," *arXiv preprint arXiv:1609.02907*, 2016.

[18] A. Sanchez-Gonzalez, J. Godwin, T. Pfaff, R. Ying, J. Leskovec, and P. Battaglia, "Learning to simulate complex physics with graph networks," in *International Conference on Machine Learning*. PMLR, 2020, pp. 8459–8468.

[19] Y. Li, J. Wu, J.-Y. Zhu, J. B. Tenenbaum, A. Torralba, and R. Tedrake, "Propagation networks for model-based control under partial observation," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 1205–1211.

[20] X. Lin, Y. Wang, Z. Huang, and D. Held, "Learning visible connectivity dynamics for cloth smoothing," in *Conference on Robot Learning*. PMLR, 2022, pp. 256–266.

[21] Z. Huang, X. Lin, and D. Held, "Mesh-based dynamics with occlusion reasoning for cloth manipulation," *arXiv preprint arXiv:2206.02881*, 2022.

[22] A. Longhini, M. Moletta, A. Reichlin, M. C. Welle, D. Held, Z. Erickson, and D. Kragic, "Edo-net: Learning elastic properties of deformable objects from graph dynamics," 2022. [Online]. Available: https://arxiv.org/abs/2209.08996

[23] Z. Erickson, V. Gangaram, A. Kapusta, C. K. Liu, and C. C. Kemp, "Assistive gym: A physics simulation framework for assistive robotics," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 10 169–10 176.