Bayesian Differentiable Physics for Cloth Digitalization

Deshan Gong
University of Leeds
Leeds, United Kingdom
scdg@leeds.ac.uk

Ningtao Mao
University of Leeds
Leeds, United Kingdom
N.Mao@leeds.ac.uk

He Wang *
University College London
London, United Kingdom
he_wang@ucl.ac.uk

Figure 1. We introduce a Bayesian Differentiable Physics (BPD) model for digitalizing real cloths by inferring their physical properties from the standard Cusick drape data (a-1, b-1, c-1 left). The digitalized cloths exhibit various drapabilities, faithfully reflecting their diverse mechanical characteristics and materials (a-1, b-1, c-1 middle and right). Further, our model enables the generalization of the learned mechanical characteristics and materials to garments (a-2, b-2, c-2).

Abstract

*We propose a new method for cloth digitalization. Deviating from existing methods which learn from data captured under relatively casual settings, we propose to learn from data captured in strictly tested measuring protocols, and find plausible physical parameters of the cloths. However, such data is currently absent, so we first propose a new dataset with accurate cloth measurements. Further, the data size is considerably smaller than the ones in current deep learning, due to the nature of the data capture process. To learn from small data, we propose a new Bayesian differentiable cloth model to estimate the complex material heterogeneity of real cloths. It can provide highly accurate digitalization from very limited data samples. Through exhaustive evaluation and comparison, we show our method is **accurate** in cloth digitalization, **efficient** in learning from limited data samples, and **general** in capturing material variations. Code and data are available*[1]

*corresponding author, he_wang@ucl.ac.uk

1. Introduction

In the emergence of the Metaverse, being able to build digital replicas of specific real-world objects becomes highly desirable. Despite many efforts to digitalize relatively simple objects such as rigid bodies [69], human bodies [31], challenges still remain for objects with complex behaviors such as cloth, fluid, and gases [48, 66, 76]. One particular challenge is to digitalize cloths, which can greatly benefit multiple application domains, *e.g.* customized fashion design, computer animation, textile manufacturing, *etc.*

The key to cloth digitalization is a model that can capture the complex physical behaviors of a given cloth (not

[1] https : / / github . com / realcrane / Bayesian - Differentiable-Physics-for-Cloth-Digitalization

cloths in general). Broadly speaking, existing research provides several potential avenues. Physics-based approaches employ physics models that explicitly represent the materials of cloths. Replicating a specific cloth then comes down to hand-tuning the material parameters [15] or solving an inverse problem [50, 74]. In parallel, in computer vision, data-driven approaches employ deep learning models to learn the physical behaviors from data, without explicit physics knowledge [42]. While the former requires laborious hand-tuning and slow optimization, which is still difficult to replicate a given cloth [7], the latter requires large amounts of data and still suffers from low accuracy when it comes to mimicking specific cloth samples [30, 49]. Recently, a combination of physics and deep learning, *i.e.* differentiable physics, provides a new direction [28, 46, 50], but the ability to replicate the exact physical behaviors of a given cloth is still under-explored.

We argue that the foremost challenge in digitalizing specific cloth samples is the measurement accuracy in data. Cloth materials and dynamics are subtle, calling for fine-controlled data capture, where there is a notable difference between current deep learning and the textile standards. Compared with textile where variables are strictly controlled, *e.g.* temperature, air moisture, current deep learning learns from the data captured in far less controlled settings, *e.g.* videos [10, 74]. Consequently the learned models are merely sufficient for general motion prediction/simulation [56, 78], and are far from being accurate for detailed simulation, manufacturing, design [23, 26, 53]. We fill this gap by employing Cusick drape testing under the British Standard [1] (Fig. 2). At the high level, Cusick drape testing captures the cloth 'drapability' in images and uses them to characterize the material. It is a vision-based approach which is machine learning friendly and effective in describing the cloth physical properties [16, 18–22, 37, 38].

However, simply applying or adapting the existing methods [28, 47, 50, 64] on Cusick drape data is difficult. First, while the latest differentiable physics methods often assume cloths as homogeneous materials [28, 46, 50, 70], cloths are heterogeneous materials: the mechanical properties in different parts of the same sample are different. In addition, there exists dynamics stochasticity in cloth draping [60]. A deterministic and homogeneous model leads to averaged behaviors hence inaccurate digitalizations. Next, black-box deep learning methods cannot be used either, due to that they need a large amount of data [64]. Collecting such data is prohibitively time-consuming and labor-intensive for accurate tests such as Cusick. This is why there are few public datasets of Cusick drape, compared with hours of videos at the disposal for deep learning. So we collect our own data, but the data size is not even remotely close to videos [62]. Finally, a challenge in learning from Cusick drape is that a standard tester (*i.e.* Cusick drape meter) only provides one drape image (Fig. 2), *i.e.* no 3D geometry or motion, ruling out the methods [62, 64] that require dynamics data.

To address the aforementioned data scarcity, dynamics stochasticity and material heterogeneity, we propose a new Bayesian learning scheme. Starting from the joint probability of a Cusick drape image and the initial state, we model the stochastic draping motion as a series of probabilistic state transitions, with the learnable physical parameters as latent variables. Due to data scarcity (i.e. one image per drape), inferring the latent variables is a formidable task. Therefore, we propose a new differentiable heterogeneous cloth model to govern the dynamics of the draping, and incorporate randomnesses in the material parameters to account for the draping stochasticity. Owning to its high sample efficiency, our model can learn from extremely limited data (*i.e.* merely one image) of a draping sample. Furthermore, to account for the within-type material variations, we impose learnable posterior over the material parameters, leading to a new Bayesian differentiable cloth model, which can learn distributions of plausible physical parameters. Not only does it explicitly model the draping stochasticity, it also enables us to transfer the learned material to arbitrary geometries such as garments.

We show that our method is *accurate* in replicating highly plausible cloth mechanical behaviors, *efficient* in training with limited data, and *general* in capturing material variations. Further, the digitalized cloths can be used to simulate garments made from different materials displaying distinguishable mechanical characteristics. Since there is no existing deep learning method designed for similar tasks, we compare our method with possible alternative solutions including different cloth models and optimization methods. Formally, our contributions include: (1) a new method for cloth digitalization based on limited Cusick drape data, (2) a new Bayesian differentiable cloth model to enable accurate digitalization, and (3) a new dataset collected from the Cusick drape testing.

2. Related Work

Cloth Digitalization aims to create digital replicas of specific cloth samples. Broadly speaking, there are three approaches: supervised learning, self-supervised learning, and physical parameter estimation. Supervised learning mimicks cloth dynamics by learning from cloth [30, 59], which requires a huge amount of training data and often suffers from low generalizability in unseen environments or materials [73]. This is because this type of methods encodes little physics and entirely relies on data. By contrast, self-supervised learning explicitly considers cloth physics and does not rely on data [5, 65]. However, this method usually employs simplified physics models and therefore suffers from issues such as over-smoothing and penetrations that are difficult to solve [6]. Also, both approaches are de-

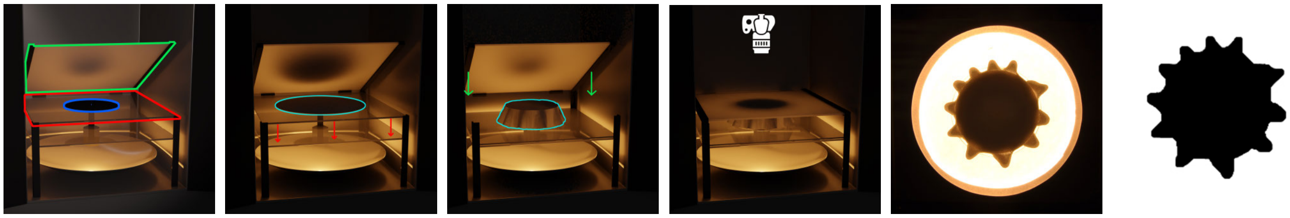

(a) Tester (b) Initial State (c) Drape State (d) Capturing (e) Raw Photo (f) Silhouette

Figure 2. Cusick drape testing. The Tester, *i.e*. Cusick drape meter, has an inner support panel (blue), an outer support panel (red) and a frosted glass lid (green). A round cloth sample is first laid flat on the support panels (light blue in Initial State). Then the outer support panel is lowered to allow the cloth to naturally drape (Drape State). Next, the glass lid is closed so that the light source at the bottom can project the cloth to the lid which is recorded by a camera at the top (Capturing). Finally, the cloth Silhouette is extracted from the Raw Photo. The whole drape meter is in a black chamber so the testing process is not observable. Due to patent restrictions, the images are rendered, not the real device.

signed to capture general behaviors of cloths, rather than replicating specific cloth samples.

Physical parameter estimation aims to infer the cloth simulation parameters so that it can reproduce the dynamics of real cloths. Compared with the aforementioned approaches, this method is the closest to building the exact digital replicas of specific cloth samples. Both model-free and model-based methods are proposed in this approach. Model-free methods learn a mapping function between the observed cloths and the physical parameters without explicitly modeling any physics [10, 39, 78]. Model-based methods employ physics models and optimize the simulation parameters to minimize the difference between the simulation and the observations. While model-free methods are simple to implement, they are data demanding and collecting real cloth data is usually prohibitively time-consuming. By contrast, model-based methods have outstanding data efficiency [17, 55, 74]. Recently, one line of model-based methods called differentiable cloth simulation has demonstrated high learning accuracy and convergence speed [28, 46, 50]. These methods leverage fully differentiable cloth simulators and gradient-based optimization to estimate the parameters. However, they need to learn from 3D geometries which are usually difficult to accurately capture. [36] uses a differentiable renderer [41, 51] so that it can learn from 2D images. However, the digitalized cloths are still significantly different from the real ones.

Our method falls into the category of differentiable cloth models. However, we argue that one common issue in existing research is the data accuracy. Cloth mechanical behaviors are affected by the environment, e.g. temperature and air moisture [8]. Therefore, the casual data collection setups employed in existing research do not actively control these factors. Further, since fabric testing has widely recognized standards [17], we argue standard apparatuses and protocols should be used. In addition to the data accuracy, we argue that the current differentiable cloth models are overly simplified. To be able to digitalize specific cloth samples, the physics model should explicitly consider material heterogeneity and behavioral stochasticity, due to the wide range of materials used in fabrics [27, 63].

To resolve these problems, we introduce a new accurate drape dataset which is collected following widely acknowledged standards [20–22]. Moreover, we propose a novel Bayesian differentiable cloth simulator that can more accurately digitalize real cloth behaviors by modeling the material heterogeneity and dynamics stochasticity through Bayesian inference with outstanding data efficiency.

Physics-based Deep Learning. Our research can be seen as a part of recent attempts to leverage deep learning to solve differential equations, which has spiked research interests to address issues such as noise modeling, finite element mesh generation and high dimensionality [4, 40, 52, 54]. Deep Neural Networks (DNNs) can learn to generate Finite Element meshes for steady state problems [81, 82]. Also, they can be part of Partial Differential Equations (PDEs) for purposes such as reduced-order modeling [32, 67], noise estimation [77] and differentiable simulation [28, 34, 79, 80]. Further, DNNs can replace PDEs completely in physics-informed neural networks (PINNs) [61, 68] where the process of solving PDEs is replaced by inference on trained DNNs. Different from existing work, we propose a Bayesian differentiable physics model for fabrics to explicitly digitalize their stochastic mechanical properties.

3. Methodology

3.1. Cusick Drape Test

Our Cusick drape meter comes with a chamber within which there are two support panels and one frosted glass lid (Fig. 2a). During testing, we first cut a cloth sample into a round shape (Fig. 3(a)) and pin its center to the cen-

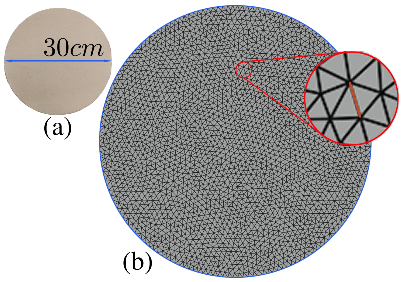

Figure 3. (a) A circular (diameter=30cm) cloth sample for Cusick drape test. (b) Cloth sample mesh has 2699 vertices, 7924 edges (7754 bending edges), and 5226 faces. A **bending edge** (highlighted in orange) is shared between two adjacent triangles, so the edges highlighted in blue (the boundary) are not bending edges.

ter of the blue panel in Fig. 2a (diameter is 18cm), shown in Fig. 2b. Then we lower the transparent panel (the red panel in Fig. 2a) and let the cloth naturally drape until the transparent panel does not contact with the cloth (Fig. 2c). Finally, an image $\mathcal{I} \in \{pix \in \mathbb{Z} : 0 \leq pix \leq 255\}^{L \times L}$ (Fig. 2e) is taken by a DSLR camera from the top (Fig. 2d). To minimize possible external perturbations such light interference, the chamber is closed when capturing.

3.2. A Bayesian Model for Cusick Drape

We discretize a cloth sample into a triangular mesh with v vertices (Fig. 3(b)). Then we define the sample's state as $\mathcal{S}_t = \{\mathbf{x}_t, \dot{\mathbf{x}}_t\}$ where $\mathbf{x}_t \in \mathbb{R}^{3v}$ and $\dot{\mathbf{x}}_t \in \mathbb{R}^{3v}$ are the vertex position and velocity respectively at time t. Therefore, a draping motion with discretized time is $\mathcal{S}_{0:n} = \{\mathcal{S}_t : t \in \mathbb{Z}^+; t \leq n\}$, with a time step size h. Since we only observe the final image $\mathcal{I}$ and the initial state $\mathcal{S}_0$, their joint probability $p(\mathcal{I}, \mathcal{S}_0) =$

$$\int \cdots \int p(\mathcal{I}|\mathcal{S}_n, \tau) \prod_{i=0}^{n-1} p(\mathcal{S}_{i+1}|\mathcal{S}_i, \tau) p(\tau) d\mathcal{S}_{1:n} d\tau \quad (1)$$

where we introduce two sets of latent variables, τ and $\mathcal{S}_{1:n}$. τ is the physical parameters. $\mathcal{S}_{1:n}$ is the intermediate states of the draping motion which we cannot directly observe. Since the draping is a physical process, it is reasonable to assume $\mathcal{S}_t$ is only affected by $\mathcal{S}_{t-1}$ and τ (Markov assumption). Additionally, the captured image $\mathcal{I}$ is only decided by the final state $\mathcal{S}_n$.

Eq. (1) is not easy to estimate due to two challenges. First, unlike the prior works which depend on dense observations on the intermediate state transitions [78] to estimate $p(\mathcal{S}_{i+1}|\mathcal{S}_i, \tau)$, Cusick drape testing does not capture the cloth sample's motion. Also, we do not observe the full $\mathcal{S}_n$, but only its (simplified) 2D representation $\mathcal{I}$.

To this end, we assume two deterministic mappings can be established for $p(\mathcal{S}_{i+1}|\mathcal{S}_i, \tau)$ and $p(\mathcal{I}|\mathcal{S}_n, \tau)$. The determinism assumptions are reasonable as $p(\mathcal{S}_{i+1}|\mathcal{S}_i, \tau)$ can be seen as a quasi-deterministic physical process, subject to minor system stochasticity which is largely mitigated by the rigorous control of Cusick settings. $p(\mathcal{I}|\mathcal{S}_n, \tau)$ can be seen as a rendering process studied in computer graphics. Here we are mainly interested in the silhouette, following the practice in textile in Cusick drape analysis [16]. Therefore, we replace $p(\mathcal{S}_{i+1}|\mathcal{S}_i, \tau)$ with a state transition function $\mathcal{S}_{i+1} = s(\mathcal{S}_i, \tau)$ where s is a deterministic function then $p(\mathcal{S}_{i+1}|\mathcal{S}_i, \tau) = p(s(\mathcal{S}_i, \tau)|\mathcal{S}_i, \tau) = 1$. Similarly, $p(\mathcal{I}|\mathcal{S}_n, \tau) = p(r(\mathcal{S}_n)|\mathcal{S}_n, \tau) = 1$ where r is a rendering function. Therefore, Eq. (1) is transformed to:

$$p(\mathcal{I}, \mathcal{S}_0) = \int p(r(\underbrace{s \ldots s}_{n}(\mathcal{S}_0, \tau))) p(\tau) d\tau \quad (2)$$

Given an initial state $\mathcal{S}_0$ and the observations $D = \{\mathcal{I}_1, \mathcal{I}_2, \ldots\}$, we maximize $p(\mathcal{I}, \mathcal{S}_0)$ which is equivalent to finding the posterior distribution of cloth material parameters: $p(\tau|D) \propto p(D|\tau)p(\tau) = p(D|r \circ s \circ s \ldots s(\mathcal{S}_0, \tau))p(\tau) = p(D|\hat{\mathcal{I}})p(\tau)$ where the composite function $\hat{\mathcal{I}} = r \circ s \circ s \ldots s(\mathcal{S}_0, \tau)$ is deterministic. Overall, the corresponding Probabilistic Graphical Model (PGM) is illustrated in Fig. 4 a.

3.3. Model Specification

To infer $p(\tau|D)$, we need to instantiate s and r. For r, we use differentiable renderer DIR-B [14]. Given a final draped state $\mathcal{S}_n$ which is a 3D mesh, and the virtual camera pose, DIR-B converts it to a 2D image. We set up the virtual camera pose (relative to the cloth drape) according to the real camera in our Cusick drape meter, so that the images captured in the Cusick drape test can be directly used as training data. Additionally, we only use drape silhouettes and ignore other information such as textures. This is because it is less reliable during capture and irrelevant to cloth drapability [20–22].

3.3.1 A Bayesian Differentiable Cloth Model

The instantiation of s is more complex than r. We propose to use a differentiable cloth model for s so that we can use back-propagation for learning. Unlike the previous differentiable models, we also want to account for the material heterogeneity, and dynamics stochasticity. Therefore, we build a stochastic heterogeneous model, where we coin the term Bayesian differentiable cloth model. We give the key equations below and leave the details in the supplementary material (SM). By using implicit Euler method [3], the physical governing equation is defined as:

$$\left(\mathbf{M} - h\frac{\partial \mathbf{f}}{\partial \dot{\mathbf{x}}} - h^2 \frac{\partial \mathbf{f}}{\partial \mathbf{x}}\right) \Delta \dot{\mathbf{x}} = h \left(\mathbf{F}_{t-1} + h\frac{\partial \mathbf{f}}{\partial \mathbf{x}} \dot{\mathbf{x}}_{t-1}\right) \quad (3)$$

where $\mathbf{M}$ is the general mass matrix, function $\mathbf{f}$ takes as input the vertex position $\mathbf{x}$ and velocity $\dot{\mathbf{x}}$ to compute the

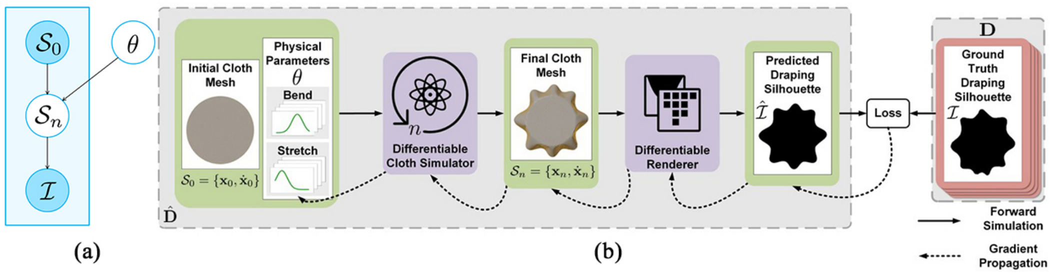

Figure 4. (a) The **Probabilistic Graphical Model** (PGM) of our Bayesian Differentiable Simulator. (b) **Model overview.** The physical parameters (stretching stiffness, bending stiffness) are first drawn from their learnable posteriors. Then the parameters and the cloth initial state are fed to a differentiable cloth simulator to run and predict cloth's final state $\mathcal{S}_n = \{\mathbf{x}_n, \dot{\mathbf{x}}_n\}$. The cloth in the final state is passed to a differentiable renderer. The rendered cloth silhouette is compared with the ground truth to compute the loss for back-propagation to update the parameters in the posteriors.

resultant force $\mathbf{F}$: $\mathbf{F}_{t-1} = \mathbf{f}(\mathbf{x}_{t-1}, \dot{\mathbf{x}}_{t-1})$ where $\mathbf{F}_{t-1}$ is the resultant force at time $t-1$. The resultant force consists of all the internal and external forces: $\mathbf{F} = \mathbf{F}_{gravity} + \mathbf{F}_{handle} + \mathbf{F}_{stretch} + \mathbf{F}_{bend}$. $\mathbf{F}_{gravity}$ is simply gravity and $\mathbf{F}_{handle}$ is the force for pinning and supporting a cloth sample, *e.g.* simulating the inner support panel (Fig. 2a).

Deviating from previous methods [28, 50], we model a material variation across the mesh which is discretized by finite elements. This allows us to localize the learning to each element, *i.e.* making the learning of $\mathbf{F}_{stretch}$ and $\mathbf{F}_{bend}$ dependent on local deformation. The stretching force [72] on face j is:

$$\mathbf{F}^{(j)}_{stretch} = -A^{(j)} \left(\sum_{m \in (uu,vv,uv)} \sigma^{(j)}_m \left(\frac{\partial \varepsilon^{(j)}_m}{\partial \mathbf{x}_i} \right) \right) \quad (4)$$

where $A^{(j)}$ is the rest area of the mesh face j, $\varepsilon^{(j)}$ denotes the stretching strain, and the stretching stress $\sigma^{(j)}_m = \mathbf{C}^{(j)} \varepsilon^{(j)}_m$ where the stretching stiffness $\mathbf{C}^{(j)} \in \mathbb{R}^{6\times4}$. The subscripts uu, vv, and uv denote strain/stress along cloth warp/wale, weft/course and diagonal direction respectively. Further, the bending force on a bending edge w (Fig. 3) is defined as:

$$\mathbf{F}^{(w)}_{bend} = \mathbf{B}^{(w)} \frac{|\mathbf{e}^{(w)}|}{\psi^{(w)}_1 + \psi^{(w)}_2} \sin\left(\frac{\gamma^{(w)}}{2} - \frac{\bar{\gamma}^{(w)}}{2}\right) u_i \quad (5)$$

where $\mathbf{B}^{(w)} \in \mathbb{R}^{3\times5}$ is the bending stiffness, $|\mathbf{e}^{(w)}|$ is the rest length of the bending edge, $\psi^{(w)}_1$ and $\psi^{(w)}_2$ are the height of the two triangular faces sharing the edge w. $\gamma^{(w)}$ and $\bar{\gamma}^{(w)}$ are the current and the rest dihedral angles between two faces [71]. Refer to [11] for the u's. Overall, the learnable parameters $\tau = \{\mathbf{C}, \mathbf{B}\}$ is high dimensional (elaborated in the SM).

Now we replace the deterministic mapping s in Eq. (2) by solving Eq. (3) forward in time, so that Eq. (2) considers physical parameter variation. The prior $p(\tau)$ (Eq. (2)) is used as a belief of the parameter distribution and a posterior $p(\tau|D)$ is learned through inference.

3.4. Model Inference

Directly estimating $p(\tau|D)$ is computationally intractable. We adopt variational inference [33] to seek an variational distribution $q_\theta(\tau)$ parameterized by θ, to approximate the true posterior $p(\tau|D)$ by minimizing the Kullback-Leibler divergence between them:

$$\begin{aligned}
\theta &= \underset{\theta}{\operatorname{argmin}}\, D_{\mathbb{KL}}(q_\theta(\tau) \| p(\tau|D)) \\
&= \underset{\theta}{\operatorname{argmin}}\, \mathbb{E}_{q_\theta(\tau)} \left[\log q_\theta(\tau) - \log \left(\frac{p(D|\tau)p(\tau)}{p(D)} \right) \right] \\
&= \underset{\theta}{\operatorname{argmin}}\, \mathbb{E}_{q_\theta(\tau)} [\log q_\theta(\tau) - \log p(D|\tau)p(\tau)] + \underbrace{\log p(D)}_{\text{const}} \\
&\equiv \underset{\theta}{\operatorname{argmin}}\, \underbrace{\mathbb{E}_{q_\theta(\tau)} [\log q_\theta(\tau) - \log p(D|\tau)p(\tau)]}_{\mathcal{L}(\theta|\mathcal{D},\tau)}
\end{aligned} \quad (6)$$

Calculating $\mathcal{L}(\theta|\mathcal{D},\tau)$, the negative evidence lower bound, is computationally prohibitive. So we approximate it by Monte Carlo sampling:

$$\mathcal{L}(\theta|\mathcal{D},\tau) \approx \sum_{i=1}^{m} \big(\underbrace{\log q_\theta(\tau_i) - \log p(D|\tau_i)p(\tau_i)}_{l(\theta,\tau)} \big) \quad (7)$$

where τ_i denotes the ith Monte Carlo sample from the variational posterior distribution $q_\theta(\tau)$. Moreover, we assume that the cloth physical parameters are distributed as a Gaussian $\tau \sim \mathcal{N}(\mu, \Sigma)$ where Σ is a diagonal matrix. To enable stochastic gradient back-propagation, we adopt the re-parameterization trick [9, 44] to sample physical parameters $\tau = t(\epsilon, \theta)$ by shifting a stochastic parameter-free noise $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ through the deterministic function $t(\epsilon, \theta) = \mu + \log(1 + exp(\eta)) \odot \epsilon$ where variational parameters $\theta = \{\mu, \eta\}$. Consequently, the variational distribution q_θ is sought within the Gaussian family and the prior

$p(\tau)$ is an isotropic Gaussian distribution with fixed parameters. Additionally, the output distribution is also a Gaussian $\mathcal{N}(\mu_I, \sigma^2)$ whose mean depends on predicted image $\hat{\mathcal{I}}$, i.e. $\mu_I = \hat{\mathcal{I}}$. The variance, σ^2, is fixed and used to control the tolerance to residual error. Therefore, the negative log likelihood $-\log p(D|\tau)$ is:

$$-\sum_{i=1}^{L}\sum_{j=1}^{L} \log\left[\left(\frac{1}{2\pi\sigma^2}\right)^{\frac{1}{2}} e^{-\frac{1}{2\sigma^2}(\mathcal{I}_{ij}-\hat{\mathcal{I}}_{ij})^2}\right] \tag{8}$$

which is essentially proportional to the Mean Squared Error (MSE). In back-propagation, the gradients of the variational distribution parameters are calculated by[9]:

$$\frac{\partial l(\theta,\tau)}{\partial\mu} = \frac{\partial l(\theta,\tau)}{\partial\tau} + \frac{\partial l(\theta,\tau)}{\partial\mu} \tag{9}$$

$$\frac{\partial l(\theta,\tau)}{\partial\eta} = \frac{\partial l(\theta,\tau)}{\partial\tau}\frac{\epsilon}{1+e^{-\eta}} + \frac{\partial l(\theta,\tau)}{\partial\eta} \tag{10}$$

The training/inference algorithms are detailed in the SM.

3.5. Implementation

Our differentiable cloth simulator is implemented in Pytorch's C++ frontend [58]. We exploit vectorization and CUDA GPU parallel computing for fast simulation and learning (in the SM). We use Eigen's sparse solver [29] to solve the governing equations Eq. (3). To reduce the memory consumption, we use sparse matrices whenever possible. Moreover, we do in-place gradient update for every time step for back-propagation, so our memory usage does not increase with simulation steps [13]. We use the Kaolin differentiable rendering package [35] for image rendering.

4. Data Collection

The Cusick drape data is collected following the BS EN ISO 9073-9:2008 [1]. In every test, the warp/wale and weft/course directions are aligned across samples to ensure the same initial condition. In addition, we are able to reconstruct the 3D meshes of cloth drapes thanks to our patented drape meter. However, the 3D data are mainly for evaluation and the silhouette images are for learning. This is because not every Cusick drape meter can reconstruct 3D meshes out of cloth drapes. Being able to learn from the 2D silhouette images only is crucial in making our method applicable in real-world settings. In addition, we also measure the sample weight and calculate the average area density $\rho = (\sum_{i<13}^{i=1} m_i)/(12\pi R^2)$ where m_i denotes the measured weight of cloth sample i and $R = 0.15m$. We also measure the thickness and include them in our dataset (refer to SM).

5. Experiments

Due to the nature of limited data, we conduct our experiments on small datasets. We use 5 representative types of cloths (each with 12 samples) that exhibit visually distinguishable drapability. For convenience, we name cloths as "material color", e.g. Cotton Blue, Cotton Pink. To digitalize a cloth, we randomly select 1 out of 12 samples for training to learn their parameter distributions. During learning, we run 100 steps ($n = 100$) for the forward simulation, with time step size $h = 0.05s$ and the total time lapse is $5s$. To qualitatively evaluate the digitalization, we simulate garments made from different digitalized cloths and visually compare the simulated garments with their corresponding real cloths. To evaluate the model (i.e. drape fitting), we employ several metrics. For fitting capability, we use mean squared error (MSE) between the fitted and the ground truth drape images. We also use Hausdorff distance (H.Dis) between the simulated 3D mesh and the ground-truth (GT) mesh. Further, we use radius-angle graph (Fig. 8a) [43] which is widely adopted for describing Cusick drape waves and comparing drape shapes. As generalization, we also test if our model learned on one sample can predict the physics of unseen cloth samples of the same type. https://youtu.be/ProN0y1bURY has more visual results.

5.1. Cloth Digitalization for Garments

Our model can digitalize real cloths for garment simulation, by sampling from the learned parameter distributions. We show three representative cloths (Fig. 1 (a-1, b-1, c-1)) with distinctive silhouette and 3D drape shapes due to their diversified drapability. Given a skirt geometry, we sample from the learned parameter distributions for the stretching and bending stiffness for each mesh triangle for simulation, shown in Fig. 1 (a-2, b-2, c-2). Within each cloth, both static (left) and dynamic (right) drapes are shown. Visually, the drapability of the skirts are similar to their corresponding cloths above. For instance, the Viscose White has a small drape shape because it is heavy and soft (small bending stiffness). By contrast, the Wool Red has a much larger drape shape because it is thick and stiff (large bending stiffness). Correspondingly, the simulated Wool Red skirt looks wider in the static and has larger folds in motions. Here, being able to learn the possible distributions of physical parameters enables us to apply the learned material to garments with arbitrary geometries.

Apart from cross-material differences, our Bayesian Differentiable Physics (BDP) can also capture the material heterogeneity and dynamics stochasticity within the same cloth type. Fig. 5 left shows the simulated skirts made from the same three cloths as in Fig. 1. Three sets of parameters are sampled for each cloth type and we also show the silhouette of the drape of the skirt in Fig. 5 right. Within each cloth type, the material heterogeneity and dynamics stochasticity can be observed in both the 3D drapes and the 2D silhouette, while the overall distributions of them across types are

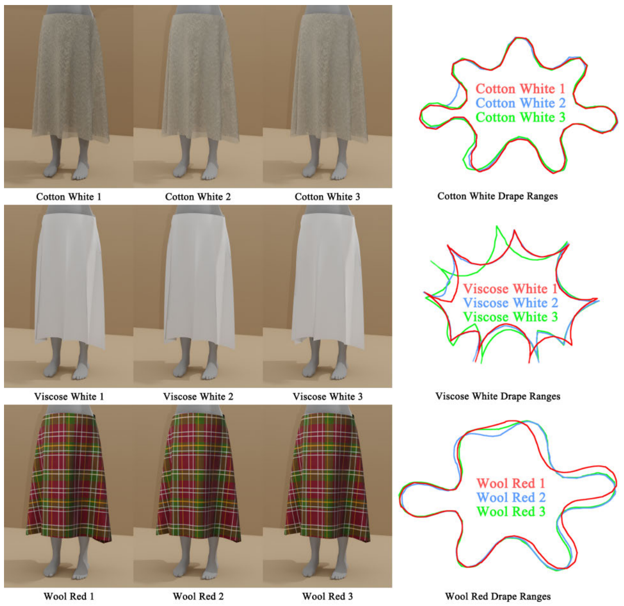

Figure 5. Given the digitialized cloths, our BDP model can simulate the skirts made from these cloths and reflect cloth material heterogeneity and draping stochasticity.

significantly different each other. This shows BDP captures both the cross-material and within-type heterogeneity, and dynamics stochasticity well.

5.2. Comparison

Alternative Cloth Models To our best knowledge, there is no similar method designed for exactly the same setting as ours. We therefore adopt [50] the closest methods as baselines because their method can also digitalize cloth by learning cloth physical parameters, albeit only taking simulation data as input. Since they only model homogeneous material, we refer to their model as HOMO. Further, we augment their model by making the cloth physical parameter learning element-wise, to enable learning heterogeneous materials. We call this model HETER for short. Both HOMO and HETER are *deterministic* models. So we use MSE as the loss function: $\mathcal{L}(\tau) = \sum_{i=1}^{L} \sum_{j=1}^{L} (\hat{\mathcal{I}}_{ij} - \mathcal{I}_{ij})^2$, for training. For comparison, we train HOMO and HETER with the same data (1 for training and the rest 11 for testing) as BDP. For BDP, we draw parameters from the learned parameter distributions for 1000 times and run Cusick drape simulations. Then, we select the best result to calculate the MSE and H. Dis.

As shown in Tab. 1 (and the corresponding rendered cloths in Fig. 7), our BDP outperforms HOMO and HETER. Visually, our method not only can accurately fit the training sample (Fig. 6 (b)), but can also predict the testing sample (Fig. 6 (f)), demonstrating that our model can capture cloth within-type variations. Unsurprisingly, HOMO and

Metrics	HOMO	HETER	BDP
Avg MSE	5.72×10^{-2}	5.51×10^{-2}	3.84×10^{-2}
Avg H. Dis	3.40×10^{-2}	3.28×10^{-2}	2.17×10^{-2}

Table 1. MSE and H. Dis (meter) of HOMO, HETER and BDP.

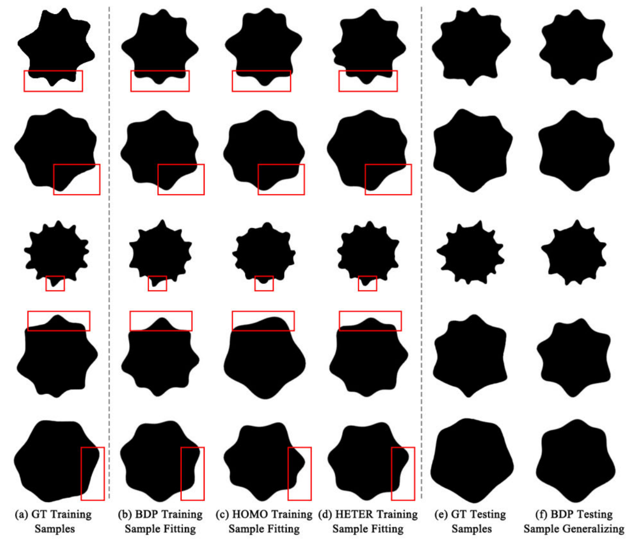

Figure 6. (a), (b), (e), and (f) show that BDP can fit the training sample and generalize to the unseen testing samples. By contrast, the learned drape shapes by HOMO (c) and HETER (d) are different from the testing sample. As such, only BDP can capture/reproduce cloth within-type variations and is superior in cloth digitalization. Additionally, by modeling material heterogeneity, HETER fit the training sample more accurately than HOMO.

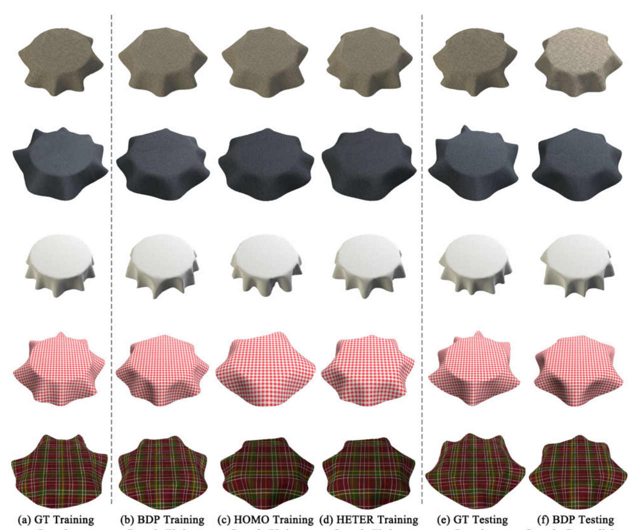

Figure 7. The rendered drape shapes in Fig. 6.

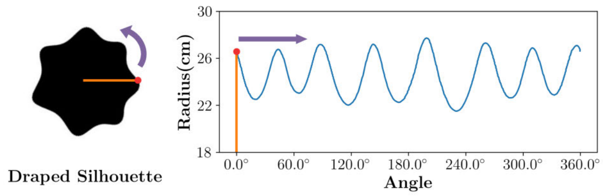

(a) An example of draping radius

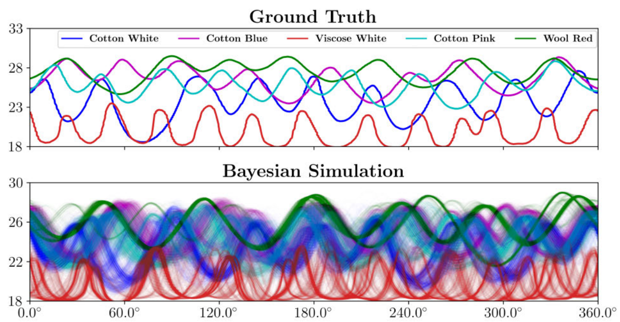

(b) Clustered drape radius

Figure 8. (a): an drape radius-angle graph illustrates the varied radius of an draped cloth sample's boundary w.r.t. angle [43]. (b): the ground truth from five real cloth samples (top) and the simulated 1000 samples from our BDP (bottom) for the corresponding samples. It demonstrates our trained BDP models can distinguish different cloths and are not overly generalized.

HETER are deterministic models and can only learn from and reproduce the training sample. Furthermore, compared with HOMO, HETER fits the training sample better and this demonstrates the importance of modeling cloth material heterogeneity. However, HETER alone is inadequate for cloth digitialization because the estimated physical parameters are tied to the sample geometry and it is difficult to generalize it to garments. Although HOMO can theoretically generalize to garments, it can only simulate the material specific to the training sample, not being able to generalize to similar materials within the same cloth type.

Although BDP can learn within-type material variations (Fig. 6), one question is whether it overly generalizes, *i.e.* unable to distinguish different cloth types. To this end, we use radius-variation [43] to demonstrate that BDP does not overly generalize. In Fig. 8b, the top figure shows that the five real cloth samples have distinctive drape shapes. The bottom figure shows the drape shapes of samples from our BDP which largely follow the same patterns as the ground truth. For example, Wool Red is obviously stiffer than Viscose White, as shown in the Ground Truth. Likewise, it is also observed in our BDP's simulation at the bottom. Refer to the SM for more results.

Alternative Learning Methods While our method is built on differentiable physics models for learning, *i.e.* derivative-based optimization, the traditional methods widely used in material science and physics are usually based on derivative-free optimization. So we also compare different learning strategies. Bayesian Optimization (BO) is a representative derivative-free optimizer which usually uses the Gaussian Process to approximate an unknown optimized objective function [12, 25]. However, the performance of vanilla BO drops drastically when the number of parameters is above 20 [24, 45]. There are over 200,000 parameters in the HETER cloth model. So we employ Random Embedding Bayesian Optimization (REMBO) [75] and Hashing-enhanced Subspace Bayesian Optimization (HeSBO) [57] as baselines. We use BoTorch's [2] REMBO and HeSBO implementation and compare them with our BDP. Given the same drape silhouette, we run our method 1000 epochs (gain the best optimization result within 200 epochs), and the REMOB and HeSBO for 500 trials. Tab. 2 shows that, our derivative-based method is better with fewer optimization steps. More results can be found in the SM.

Metrics	BDP	REMBO	HeSBO
Avg MSE	4.52×10^{-2}	6.03×10^{-2}	5.90×10^{-2}
Avg H. Dis	2.01×10^{-2}	4.25×10^{-2}	4.35×10^{-2}

Table 2. The average MSE and H. Dis of ours, REMBO, and HeSBO optimizer. The results shows that our gradient-based optimization achieves better results.

6. Conclusion and Future Work

We have proposed a new method for cloth digitalization by estimating detailed cloths physical properties. To our best knowledge, this is the first Bayesian differentiable cloth model that can work seamlessly with standard Cusick drape data. Our model has been proven to be highly accurate and generalizable. Despite focused on cloth digitalization, we believe BDP as a framework has the potential to generalize to more generic digitalization tasks and itself is a methodological extension of current DP research in computer vision. Compared with black-box deep learning methods, our limitation is that it requires prior knowledge of the underlying physics and cannot simply *plug and play* on data. However, we argue that this is a reasonable trade-off when data collection is expensive and slow. In future, we plan to model more dynamics stochasticity, *e.g.* buckling, and also compare our simulated garments with real ones by using accurate motion capture and 3D reconstruction systems.

Acknowledgements

The project is partially supported by the Art and Humanities Research Council (AHRC) UK under project FFF (AH/S002812/1).

References

[1] Bs en iso 9073-9:2008: Textiles. test methods for nonwovens: Determination of drapability including drape coefficient, 2008. 2, 6

[2] Maximilian Balandat, Brian Karrer, Daniel Jiang, Samuel Daulton, Ben Letham, Andrew G Wilson, and Eytan Bakshy. Botorch: a framework for efficient monte-carlo bayesian optimization. *Advances in neural information processing systems*, 33:21524–21538, 2020. 8

[3] David Baraff and Andrew Witkin. Large steps in cloth simulation. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, pages 43–54, 1998. 4

[4] Christian Beck, Martin Hutzenthaler, Arnulf Jentzen, and Benno Kuckuck. An overview on deep learning-based approximation methods for partial differential equations. *arXiv preprint arXiv:2012.12348*, 2020. 3

[5] Hugo Bertiche, Meysam Madadi, and Sergio Escalera. Pbns: physically based neural simulator for unsupervised garment pose space deformation. *arXiv preprint arXiv:2012.11310*, 2020. 2

[6] Hugo Bertiche, Meysam Madadi, and Sergio Escalera. Neural cloth simulation. *ACM Transactions on Graphics (TOG)*, 41(6):1–14, 2022. 2

[7] Kiran S Bhat, Christopher D Twigg, Jessica K Hodgins, Pradeep Khosla, Zoran Popovic, and Steven M Seitz. Estimating cloth simulation parameters from video. 2003. 2

[8] Dennis P Bishop. Fabrics: sensory and mechanical properties. *Textile Progress*, 26(3):1–62, 1996. 3

[9] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In *International conference on machine learning*, pages 1613–1622. PMLR, 2015. 5, 6

[10] Katherine L Bouman, Bei Xiao, Peter Battaglia, and William T Freeman. Estimating the material properties of fabric from video. In *Proceedings of the IEEE international conference on computer vision*, pages 1984–1991, 2013. 2, 3

[11] Robert Bridson, Sebastian Marino, and Ronald Fedkiw. Simulation of clothing with folds and wrinkles. In *ACM SIGGRAPH 2005 Courses*, pages 3–es. 2005. 5

[12] Eric Brochu, Vlad M Cora, and Nando De Freitas. A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*, 2010. 8

[13] Tianqi Chen, Bing Xu, Chiyuan Zhang, and Carlos Guestrin. Training deep nets with sublinear memory cost. *arXiv preprint arXiv:1604.06174*, 2016. 6

[14] Wenzheng Chen, Huan Ling, Jun Gao, Edward Smith, Jaakko Lehtinen, Alec Jacobson, and Sanja Fidler. Learning to predict 3d objects with an interpolation-based differentiable renderer. *Advances in Neural Information Processing Systems*, 32, 2019. 4

[15] Kwang-Jin Choi and Hyeong-Seok Ko. Research problems in clothing simulation. *Computer-aided design*, 37(6):585–592, 2005. 2

[16] Chauncey C Chu, Clinton L Cummings, and Newton A Teixeira. Mechanics of elastic performance of textile materials: Part v: a study of the factors affecting the drape of fabrics—the development of a drape meter. *Textile Research Journal*, 20(8):539–548, 1950. 2, 4

[17] David Clyde, Joseph Teran, and Rasmus Tamstorf. Modeling and data-driven parameter estimation for woven fabrics. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 1–11, 2017. 3

[18] Billie J Collier. Measurement of fabric drape and its relation to fabric mechanical properties and subjective evaluation. *Clothing and Textiles Research Journal*, 10(1):46–52, 1991. 2

[19] Billie J Collier, Virginia A Paulins, and John R Collier. Effects of interfacing type on shear and drape behavior of apparel fabrics. *Clothing and Textiles Research Journal*, 7(3): 51–56, 1989.

[20] GE Cusick. 30—the resistance of fabrics to shearing forces: A study of the experimental method due to mörner and eegolofsson. *Journal of the Textile Institute Transactions*, 52(9): T395–T406, 1961. 3, 4

[21] GE Cusick. 46—the dependence of fabric drape on bending and shear stiffness. *Journal of the Textile Institute Transactions*, 56(11):T596–T606, 1965.

[22] GE Cusick. 21—the measurement of fabric drape. *Journal of the Textile Institute*, 59(6):253–260, 1968. 2, 3, 4

[23] Koorosh Delavari and Hadi Dabiryan. Mathematical and numerical simulation of geometry and mechanical behavior of sandwich composites reinforced with $1\times$ 1-rib-gaiting weft-knitted spacer fabric; compressional behavior. *Composite Structures*, 268:113952, 2021. 2

[24] David Eriksson and Martin Jankowiak. High-dimensional bayesian optimization with sparse axis-aligned subspaces. In *Uncertainty in Artificial Intelligence*, pages 493–503. PMLR, 2021. 8

[25] Peter I Frazier. A tutorial on bayesian optimization. *arXiv preprint arXiv:1807.02811*, 2018. 8

[26] Ziyue Gao and Li Chen. A review of multi-scale numerical modeling of three-dimensional woven fabric. *Composite Structures*, 263:113685, 2021. 2

[27] Salvatore Gazzo. *Characterisation of the mechanical behaviour of networks and woven fabrics with a discrete homogenization model*. PhD thesis, Université de Lyon; Università degli studi (Catane, Italie), 2019. 3

[28] Deshan Gong, Zhanxing Zhu, Bulpitt. Andrew, and He Wang. Fine-grained differentiable physics: a yarn-level model for fabrics. In *International Conference on Learning Representations*, 2022. 2, 3, 5

[29] Gaël Guennebaud, Benoît Jacob, et al. Eigen v3. http://eigen.tuxfamily.org, 2010. 6

[30] Erhan Gundogdu, Victor Constantin, Shaifali Parashar, Amrollah Seifoddini, Minh Dang, Mathieu Salzmann, and Pascal Fua. Garnet++: Improving fast and accurate static 3d cloth draping by curvature loss. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(1):181–195, 2020. 2

[31] Erik Gärtner, Mykhaylo Andriluka, Erwin Coumans, and Cristian Sminchisescu. Differentiable dynamics for articulated 3d human motion reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13190–13200, 2022. 1

[32] Jiequn Han, Arnulf Jentzen, and E Weinan. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115 (34):8505–8510, 2018. 3

[33] Matthew D Hoffman, David M Blei, Chong Wang, and John Paisley. Stochastic variational inference. *Journal of Machine Learning Research*, 2013. 5

[34] Philipp Holl, Vladlen Koltun, Kiwon Um, and Nils Thuerey. phiflow: A differentiable pde solving framework for deep learning via physical simulations. In *NeurIPS Workshop*, 2020. 3

[35] Krishna Murthy Jatavallabhula, Edward Smith, Jean-Francois Lafleche, Clement Fuji Tsang, Artem Rozantsev, Wenzheng Chen, Tommy Xiang, Rev Lebaredian, and Sanja Fidler. Kaolin: A pytorch library for accelerating 3d deep learning research. *arXiv preprint arXiv:1911.05063*, 2019. 6

[36] Krishna Murthy Jatavallabhula, Miles Macklin, Florian Golemo, Vikram Voleti, Linda Petrini, Martin Weiss, Breandan Considine, Jerome Parent-Levesque, Kevin Xie, Kenny Erleben, et al. gradsim: Differentiable simulation for system identification and visuomotor control. *arXiv preprint arXiv:2104.02646*, 2021. 3

[37] YJ Jeong. A study of fabric-drape behaviour with image analysis part i: Measurement, characterisation, and instability. *Journal of the Textile Institute*, 89(1):59–69, 1998. 2

[38] YJ Jeong and DG Phillips. A study of fabric-drape behaviour with image analysis. part ii: the effects of fabric structure and mechanical properties on fabric drape. *Journal of the Textile Institute*, 89(1):70–79, 1998. 2

[39] Eunjung Ju and Myung Geol Choi. Estimating cloth simulation parameters from a static drape using neural networks. *IEEE Access*, 8:195113–195121, 2020. 3

[40] George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nat Rev Phys*, (3):422–440, 2021. 3

[41] Hiroharu Kato, Deniz Beker, Mihai Morariu, Takahiro Ando, Toru Matsuoka, Wadim Kehl, and Adrien Gaidon. Differentiable rendering: A survey. *arXiv preprint arXiv:2006.12057*, 2020. 3

[42] Doyub Kim, Woojong Koh, Rahul Narain, Kayvon Fatahalian, Adrien Treuille, and James F O'Brien. Near-exhaustive precomputation of secondary cloth effects. *ACM Transactions on Graphics (TOG)*, 32(4):1–8, 2013. 2

[43] Sungmin Kim. Determination of fabric physical properties for the simulation of cusick drapemeter. *Fibers and Polymers*, 12(1):132–136, 2011. 6, 8

[44] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 5

[45] Ben Letham, Roberto Calandra, Akshara Rai, and Eytan Bakshy. Re-examining linear embeddings for high-dimensional bayesian optimization. *Advances in neural information processing systems*, 33:1546–1558, 2020. 8

[46] Yifei Li, Tao Du, Kui Wu, Jie Xu, and Wojciech Matusik. Diffcloth: Differentiable cloth simulation with dry frictional contact. *arXiv preprint arXiv:2106.05306*, 2021. 2, 3

[47] Yifei Li, Tao Du, Kui Wu, Jie Xu, and Wojciech Matusik. Diffcloth: Differentiable cloth simulation with dry frictional contact. *ACM Trans. Graph.*, 42(1), 2022. 2

[48] Yifei Li, Tao Du, Kui Wu, Jie Xu, and Wojciech Matusik. Diffcloth: Differentiable cloth simulation with dry frictional contact. *ACM Transactions on Graphics (TOG)*, 42(1):1–20, 2022. 1

[49] Yudi Li, Min Tang, Yun Yang, Zi Huang, Ruofeng Tong, Shuangcai Yang, Yao Li, and Dinesh Manocha. N-Cloth: Predicting 3D cloth deformation with mesh-based networks. *Computer Graphics Forum (Proceedings of Eurographics)*, 41(2):547–558, 2022. 2

[50] Junbang Liang, Ming Lin, and Vladlen Koltun. Differentiable cloth simulation for inverse problems. *Advances in Neural Information Processing Systems*, 32, 2019. 2, 3, 5, 7

[51] Shichen Liu, Tianye Li, Weikai Chen, and Hao Li. Soft rasterizer: A differentiable renderer for image-based 3d reasoning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7708–7717, 2019. 3

[52] Lu Lu, Xuhui Meng, Zhiping Mao, and George Em Karniadakis. Deepxde: A deep learning library for solving differential equations. *SIAM Review*, 63(1):208–228, 2021. 3

[53] Christiane Luible and Nadia Magnenat-Thalmann. The simulation of cloth using accurate physical parameters. *CGIM 2008, Insbruck, Austria*, 2008. 2

[54] Chuizheng Meng, Sungyong Seo, Defu Cao, Sam Griesemer, and Yan Liu. When physics meets machine learning: A survey of physics-informed machine learning. *arXiv preprint arXiv:2203.16797*, 2022. 3

[55] Eder Miguel, Rasmus Tamstorf, Derek Bradley, Sara C Schvartzman, Bernhard Thomaszewski, Bernd Bickel, Wojciech Matusik, Steve Marschner, and Miguel A Otaduy. Modeling and estimation of internal friction in cloth. *ACM Transactions on Graphics (TOG)*, 32(6):1–10, 2013. 3

[56] Rahul Narain, Armin Samii, and James F O'brien. Adaptive anisotropic remeshing for cloth simulation. *ACM transactions on graphics (TOG)*, 31(6):1–10, 2012. 2

[57] Amin Nayebi, Alexander Munteanu, and Matthias Poloczek. A framework for bayesian optimization in embedded subspaces. In *International Conference on Machine Learning*, pages 4752–4761. PMLR, 2019. 8

[58] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 6

[59] Chaitanya Patel, Zhouyingcheng Liao, and Gerard Pons-Moll. Tailornet: Predicting clothing in 3d as a function of human pose, shape and garment style. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7365–7375, 2020. 2

[60] P Pratihar. Static and dynamics drape of fabric: An emerging arena of fabric evaluation. *International Journal of Engi-*

neering Research and Applications*, 3(5):1007–1011, 2013. 2

[61] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019. 3

[62] Tom FH Runia, Kirill Gavrilyuk, Cees GM Snoek, and Arnold WM Smeulders. Cloth in the wind: A case study of physical measurement through simulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10498–10507, 2020. 2

[63] Shunsuke Saito, Nobuyuki Umetani, and Shigeo Morishima. Macroscopic and microscopic deformation coupling in upsampled cloth simulation. *Computer Animation and Virtual Worlds*, 25(3-4):435–444, 2014. 3

[64] Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter Battaglia. Learning to simulate complex physics with graph networks. In *Proceedings of the 37th International Conference on Machine Learning*, pages 8459–8468. PMLR, 2020. 2

[65] Igor Santesteban, Miguel A Otaduy, and Dan Casas. Snug: Self-supervised neural dynamic garments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8140–8150, 2022. 2

[66] Connor Schenck and Dieter Fox. Spnets: Differentiable fluid dynamics for deep neural networks. In *Conference on Robot Learning*, pages 317–335. PMLR, 2018. 1

[67] Siyuan Shen, Yin Yang, Tianjia Shao, He Wang, Chenfanfu Jiang, Lei Lan, and Kun Zhou. High-order differentiable autoencoder for nonlinear model reduction. *ACM Trans. Graph.*, 40(4), 2021. 3

[68] Yanjie Song, He Wang, He Yang, Maria Luisa Taccari, and Xiaohui Chen. Loss-attentional physics-informed neural networks. *Journal of Computational Physics*, page 112781, 2024. 3

[69] Michael Strecke and Joerg Stueckler. Diffsdfsim: Differentiable rigid-body dynamics with implicit shapes. In *2021 International Conference on 3D Vision (3DV)*, pages 96–105. IEEE, 2021. 1

[70] Tuur Stuyck and Hsiao-yu Chen. Diffxpbd: Differentiable position-based simulation of compliant constraint dynamics. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 6(3):1–14, 2023. 2

[71] Rasmus Tamstorf and Eitan Grinspun. Discrete bending forces and their jacobians. *Graphical models*, 75(6):362–370, 2013. 5

[72] Pascal Volino, Nadia Magnenat-Thalmann, and Francois Faure. A simple approach to nonlinear tensile stiffness for accurate cloth simulation. *ACM Transactions on Graphics*, 28(4):Article–No, 2009. 5

[73] Huamin Wang. Gpu-based simulation of cloth wrinkles at submillimeter levels. *ACM Transactions on Graphics (TOG)*, 40(4):1–14, 2021. 2

[74] Huamin Wang, James F O'Brien, and Ravi Ramamoorthi. Data-driven elastic models for cloth: modeling and measurement. *ACM transactions on graphics (TOG)*, 30(4):1–12, 2011. 2, 3

[75] Ziyu Wang, Frank Hutter, Masrour Zoghi, David Matheson, and Nando De Feitas. Bayesian optimization in a billion dimensions via random embeddings. *Journal of Artificial Intelligence Research*, 55:361–387, 2016. 8

[76] Zhou Xian, Bo Zhu, Zhenjia Xu, Hsiao-Yu Tung, Antonio Torralba, Katerina Fragkiadaki, and Chuang Gan. Fluidlab: A differentiable environment for benchmarking complex fluid manipulation. *arXiv preprint arXiv:2303.02346*, 2023. 1

[77] Liu Yang, Xuhui Meng, and George Em Karniadakis. B-pinns: Bayesian physics-informed neural networks for forward and inverse pde problems with noisy data. *Journal of Computational Physics*, 425:109913, 2021. 3

[78] Shan Yang, Junbang Liang, and Ming C Lin. Learning-based cloth material recovery from video. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4383–4393, 2017. 2, 3, 4

[79] Jiangbei Yue, Dinesh Manocha, and He Wang. Human trajectory prediction via neural social physics. In *The European Conference on Computer Vision (ECCV)*, pages 376–394, 2022. 3

[80] Jiangbei Yue, Dinesh Manocha, and He Wang. Human trajectory forecasting with explainable behavioral uncertainty. *arXiv preprint arXiv:2307.01817*, 2023. 3

[81] Zheyan Zhang, Yongxing Wang, Peter K. Jimack, and He Wang. Meshingnet: A new mesh generation method based on deep learning. In *Computational Science – ICCS 2020*, pages 186–198, Cham, 2020. Springer International Publishing. 3

[82] Zheyan Zhang, Peter K. Jimack, and He Wang. MeshingNet3D: Efficient generation of adapted tetrahedral meshes for computational mechanics. *Advances in Engineering Software*, 157-158, 2021. 3