A Visual Approach to Measure Cloth-Body and Cloth-Cloth Friction

Abdullah Haroon Rasheed ⓘ, Victor Romero ⓘ, Florence Bertails-Descoubes, Stefanie Wuhrer, Jean-Sebastien Franco ⓘ, and Arnaud Lazarus

Abstract—Measuring contact friction in soft-bodies usually requires a specialised physics bench and a tedious acquisition protocol. This makes the prospect of a purely non-invasive, video-based measurement technique particularly attractive. Previous works have shown that such a video-based estimation is feasible for material parameters using deep learning, but this has never been applied to the friction estimation problem which results in even more subtle visual variations. Because acquiring a large dataset for this problem is impractical, generating it from simulation is the obvious alternative. However, this requires the use of a frictional contact simulator whose results are not only visually plausible, but physically-correct enough to match observations made at the macroscopic scale. In this paper, which is an extended version of our former work A. H. Rasheed, V. Romero, F. Bertails-Descoubes, S. Wuhrer, J.-S. Franco, and A Lazarus, "Learning to measure the static friction coefficient in cloth contact," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 9909–9918, we propose to our knowledge the first non-invasive measurement network and adjoining synthetic training dataset for estimating cloth friction at contact, for both cloth-hard body and cloth-cloth contacts. To this end we build a protocol for validating and calibrating a state-of-the-art frictional contact simulator, in order to produce a reliable dataset. We furthermore show that without our careful calibration procedure, the training fails to provide accurate estimation results on real data. We present extensive results on a large acquired test set of several hundred real video sequences of cloth in friction, which validates the proposed protocol and its accuracy.

Index Terms—Friction estimation, cloth simulation, deep learning, material estimation, inverse problem

◆

1 INTRODUCTION

COMPUTER vision techniques are powerful in analysing the shape and motion of physical objects from mere images. Yet, accurately inferring physical properties from visual data is still a challenging problem [40]. Our work aims at estimating physical parameters from the rich dynamic effects that can be observed in cloth motion. Applications range from non-invasive estimation of cloth properties in textile engineering to realistic clothing synthesis for virtual and augmented reality. So far, most existing methods for cloth acquisition have focused on estimating cloth material parameters, that is, its stiffness and mass [3], [14], [42]. Friction at contact is a much less studied problem due to the difficulty of modelling, measuring, and simulating dry frictional contact accurately. However, friction has a high impact on the overall cloth dynamics, as illustrated in Fig. 1.

Measuring a friction coefficient accurately in cloth is a long-standing problem, mainly studied in textile engineering [28], [36]. Better accuracy often comes at the expense of tedious and generally invasive protocols with specialised mechanical

- Abdullah Haroon Rasheed, Victor Romero, Florence Bertails-Descoubes, Stefanie Wuhrer, and Jean-Sebastien Franco are with University Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38400 Grenoble, France.
  E-mail: {abdullah-haroon.rasheed, victor.romero-gramegna, florence.descoubes, stefanie.wuhrer, jean-sebastien.franco}@inria.fr.
- Arnaud Lazarus is with Sorbonne Universit', CNRS, UMR 7190, Institut Jean le Rond d'Alembert, 75005 Paris, France.
  E-mail: Arnaud.Lazarus@upmc.fr.

sensors and hardware. Yet, the visual variability induced by friction shown in Fig. 1 hints at the possibility of a purely vision-based protocol, discriminating the different motions and folding patterns observed under friction variation. This general idea has recently been leveraged by a few studies in physics to infer friction coefficients from contacting slender structures like hair strands [9] or stiff ribbons [32]. While such methods effectively exploit the connection between observed geometry and friction, they are only valid for very specific classes of materials and conditions – such as setups where gravity plays a negligible role. A major challenge for cloth is that, larger friction coefficient differences are easily distinguished, but smaller differences yield increasingly subtle visual differences (Fig. 1, centre and bottom rows).

In this work, we propose a first step towards a generally applicable vision-based method to estimate the dry friction coefficient between cloth and a contacting surface, and between two cloth samples.

To this goal, we use a Long-term Recurrent Convolutional Network (LRCN) [15] to regress from an input video showing a cloth motion under contact to both material parameters and a dry friction coefficient. Our network is trained on sequences of simulated cloth deformations. We leverage the recent simulator ARGUS [22] for capturing dry frictional contact in cloth, and search parameter settings that lead to sufficiently accurate results for our needs. Our results indicate that the network, trained exclusively on simulated data, does generalise to real videos showing similar cloth motions; and this in spite of the differences between the renderings of the simulations and the real captures.

In scenarios where capturing significant amounts of real data is prohibitive, many works have explored the idea of training on simulated data with the aim to generalise
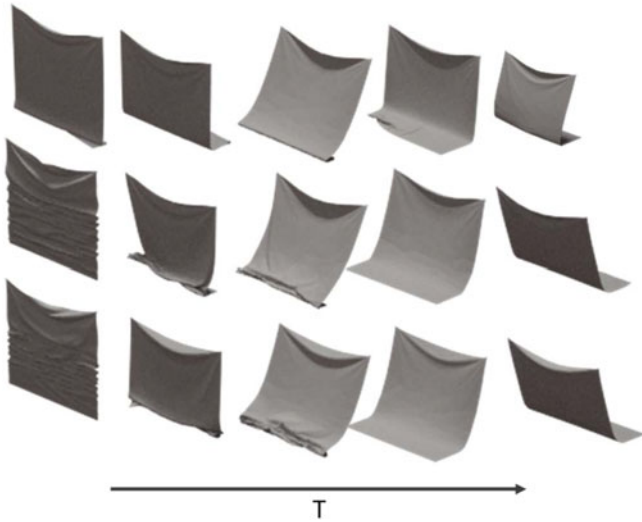
Fig. 1. Three cloth motion sequences simulated with the same material but different friction coefficients at contact (μ). Top: $\mu = 0.0$, Centre: $\mu = 0.5$, Bottom: $\mu = 0.6$. Differences are significant between $\mu = 0.0$ and $\mu = 0.5$, but more subtle between $\mu = 0.5$ and $\mu = 0.6$.

inference on real data. These include approaches for cross domain transfer learning [19], [37] and interpretable low dimensional representation learning [18]. Our approach can be considered orthogonal to such efforts as we present a protocol to calibrate simulated data generation by experimental verification, before learning a model from this data.

This paper is an extension to our previously presented work [31], which trains a model on synthetic data, generated from a verified simulator, and tests it on real captured data to estimate material and friction parameters between a cloth and a substrate. In addition to [31], we show here that calibrating the simulator properly before generating the training data is a mandatory step for achieving good estimates on real data. Furthermore, we extend our framework [31] to estimate friction in case of cloth on cloth contact, a harder measurement problem which benefits from the visual nature of our measurement protocol. Finally, we show examples of applications of our measurement protocol by creating simulations of cloth materials and substrates from our real captured test data.

Altogether, the combined contributions of our previously presented work [31] and that of our above extensions, merged in this paper, are the following:

- We present a deep learning based algorithm to solve the inverse problem of parameter estimation for cloth simulation. Our algorithm measures the material and friction properties of cloth given a video sequence.
- To generate physically correct training data, we validate the ARGUS simulator physically and calibrate its accuracy against a constrained, measurable real-world physics experiment.
- We demonstrate the usefulness of physically validating the simulator before learning data from it; to achieve this we generate 3 datasets with a variable level of physical accuracy, and show that only the dataset produced by the calibrated simulator leads to satisfactory prediction results of our learning-based method.
- We release 2 real captured test datasets of 315 and 225 real videos in controlled conditions closely matching

the simulation setup, and make baseline estimates of friction values within a measurement error range.
- We demonstrate the validity of our learned model by achieving an error <0.1 from the baseline range on 93.6 percent of real test data capturing friction between a cloth and substrate.
- We extend our full method to cloth on cloth contact. We present friction parameter estimation results on this test dataset and achieve an error of <0.1 from the baseline range on 81.1 percent of the test dataset.

2   RELATED WORKS

Inferring material properties of an object from geometric or visual data has been studied in various communities including computer vision, computer graphics, textile engineering and physics. We first introduce the general model of friction we consider and existing techniques to measure it, before focusing on slender elastic structures such as cloth.

Dry Friction Models. Dry friction is a force that opposes the relative motion of two solid objects. The way in which the interface of two compliant objects in contact evolves when subjected to load is complex, and finding the precise law that describes this phenomenon remains an open problem in physics and mechanics. Our work is based on Amontons-Coulomb's law [2] for friction, which is a commonly used model that successfully approximates this complex scenario at the macroscopic scale. In this model, surfaces in contact interact throughout normal and shear forces, and sliding occurs when the ratio between the shear and normal force reaches a threshold value, called the *static friction coefficient*, which is independent of the area of contact and depends only on the roughness of the interacting surfaces. While cloth simulators usually rely on an isotropic Coulomb friction model [6], [33], some recent works have explored anisotropic variations of the Coulomb model when simulating interacting rigid bodies and cloth [8], [17], [30]. Interestingly, Chen *et al.* [8] perform extensive cloth-solid experiments and report a few frictional measurements that exhibit either some anisotropic behaviour, or a non-constant friction coefficient (i.e., a nonlinear relationship between the tangential and normal contact forces). Furthermore, in our cloth-cloth experiments, we noted that in a few cases, departures from the Coulomb model would arise from cohesive effects due to interlocking fibres. However, despite such discrepancies, overall we observed that the isotropic Coulomb model could already capture most of the cloth-solid and cloth-cloth interactions faithfully. Moreover, although simple in appearance, the isotropic Coulomb model actually proves difficult to be simulated numerically, as it is a non-smooth model characterised by a non-constant threshold depending (linearly) on the normal contact force. One strength of the ARGUS simulator precisely lies in its capability to capture this threshold accurately, without relying on any regularisation. For these reasons, we stick to the isotropic Coulomb model in our work, and show the validity of this model through our results. In the future, it would be interesting to refine our study by considering non-constant friction coefficients, anisotropy, and cohesion.

Friction Measurement. Estimating friction and material parameters jointly from visual data has recently become a topic of study in computer vision. Miguel *et al.* [27] propose

a reparameterization of Dahl's friction model [13] for estimating internal friction in cloth from geometric information. Internal friction is however different from static friction at contact since it models internal dissipation within the cloth, hence this technique does not apply to our case. Wu *et al.* [41] propose to use a generative model to estimate friction and material parameters. Unlike our work, this method is targeted at rigid objects. Zhang *et al.* [46] analyse visible reflections, while Yuan *et al.* [44] combine visual and haptic data to estimate friction information. Both works present a static joint estimation of material and friction based on visual attributes of the material, whereas we focus on the dynamic behaviour of cloth under frictional contact and wish to estimate the friction coefficient directly.

In textile engineering, friction estimation has been studied using invasive techniques [16], [25], [28], [36]. Some studies in other fields connect the perception of friction (i.e., feeling) to actual measurements [7], [24] and visual features [5]. These works show that visual features correlate to friction information using perceptual studies.

In physics, a few isolated studies consider inferring dry static friction coefficients from contacting elastic slender structures. In particular, the friction of a hair fibre can be inferred by the geometric configuration of a relaxed knot formed by the hair [9], and the friction of a relatively stiff isotropic ribbon can be inferred by the geometric configuration of the ribbon's bending and slipping behaviour when pushed onto a substrate [32]. While these studies only hold in specific scenarios with strict boundary conditions, they demonstrate the relationship between geometric configurations and friction for elastic slender structures. Our work of estimating cloth friction based on videos is inspired by such intriguing studies. Note that in order to obtain reference friction measurements for evaluation, we use a classical inclined-plane protocol [12], [36].

Material Parameter Estimation. The problem of estimating material parameters from geometric or visual data has been studied in different scenarios and for various types of materials, including soft deformable materials [35], [38], [43]. To our knowledge, recovering material information from cloth has always been formulated as a fitting problem, consisting of three steps: first, considering a material model as a function of deformation of cloth, then capturing deformations, and finally fitting the function to captured data. One approach is to control boundary conditions by setting up a physical apparatus for yielding deformations [11], [26], [39]. To avoid capturing deformations under controlled conditions, some works have explored extracting material parameters from casually captured videos of cloth [3], [21]. Bouman *et al.* [4] capture stiffness and density of fabrics from video using handcrafted discriminatory features. Davis *et al.* [14] estimate material parameters based on videos that show small vibrating motions, and apply this method to fabrics. Recently, Liang *et al.* [23] propose a computationally efficient way of estimating material parameters on synthetic data with a differentiable cloth simulator.

The closest methods to ours identify cloth material parameters from videos [14], [23], [42]. Yang *et al.* [42] train a neural network on simulated cloth deformations and use the resulting architecture to infer bending and stiffness parameters. This approach follows a recent trend in computer vision and machine learning, where deep neural networks are trained with simulated input data and tested on real data. Training with simulated input offers the advantage of yielding sufficiently large training datasets, hence we follow the same approach here. However, unlike previous methods, we carefully validate and finely tune our simulator for physical realism, through an experimental protocol described in Section 4.2.1.

Note that none of the works discussed in this section consider estimation of friction at contact.

3 BACKGROUND

In this section we describe the elastic model used for our material properties, and the Amontons-Coulomb's law for friction. We chose these models as they are known to be realistic in the range of macroscopic deformations.

3.1 Material Parameter Space

To obtain physically accurate material parameter settings in the simulator, we leverage the work of Wang *et al.* [39] that encodes the material properties of ten representative classes of cloth ranging from very soft to stiffer materials. Material parameters are encoded as three parameters of the cloth simulator ARCSIM [29]: a linear mass density, the coefficients of the strain-stress matrix and the coefficients of bending stiffness. This model has been previously used by Yang *et al.* [42] for recovering cloth material parameters.

3.2 Amontons-Coulomb's Law for Friction

Our work is based on Amontons-Coulomb's law for dry friction as this model successfully approximates the macroscopic behaviour of two solid surfaces at contact. In a simplified version that we shall use here (no distinction between static and dynamic friction coefficient), this law defines the friction coefficient μ as a threshold value for the stick to slip transition for two contacting surfaces, and as the coefficient relating normal and tangential forces during sliding. More specifically, in our scenario, let $\vec{R}$ denote the reaction force on the surface in contact with a piece of cloth. We can divide $\vec{R}$ into two components: the force component normal to the surface and the one tangential to the surface, denoted by $\vec{P}$ and $\vec{Q}$, respectively. The force $\vec{P}$ keeps the two contacting surfaces from interpenetrating, and $\vec{Q}$ opposes relative displacements between the two surfaces in contact; more precisely, the two surfaces stick if $\vec{Q} \leq \mu\vec{P}$ and they slip if $\vec{Q} = \mu\vec{P}$ (the case $\vec{Q} > \mu\vec{P}$ is not admissible). Fig. 3-right illustrates these forces for a strip of material that is pushed onto a substrate.

4 DATA GENERATION

Our first contribution consists in generating a dataset of closely matching captured videos and simulations with, for each video, corresponding material classes and friction coefficients. We choose a simple motion that can easily be replicated with a real piece of cloth while containing representative material classes and friction parameters. In particular, we consider a *drop and drag* motion, in which a square of cloth of side length 20 cm, suspended by its corners, is dropped vertically on a substrate floor and then dragged back and forth as shown in Fig. 4.

4.1 Real Data Capture

We now present our experimental data capture setup. Since we aim at predicting material and frictional properties from videos, we constrain the setup of generating real data to a controlled environment in order to remove sources of variation other than material properties and friction.

Cloth Materials. We use 9 materials in bright colours, out of which 8 are close in composition and density to material classes defined in [39], and one (silk) is not covered by these material classes. Details on the material features can be found in Table S1 of the supplemental material, which can be found on the Computer Society Digital Library at http://doi. ieeecomputersociety.org/TPAMI.2021.3097547, for this paper. From each material we laser cut 20x20 cm pieces with 1x1cm holders at two corners of one edge for clamping purposes.

Substrates. We choose 7 substrates that, combined with our material samples, allow to cover a wide range of friction behaviours. In particular, chosen substrate materials are aluminium, aluminium-PET, ceramic, rough glass, smooth glass, polyester-mirror and stainless steel.

Frictional behaviour of a single substrate can be very complicated and difficult to characterise, and through our experiments we have learnt that intuition is not reliable enough to describe the frictional behaviour of a given cloth-substrate pair. For this reason, we test several materials that can be used as a substrate, and select 7 of them which give us the broadest range of frictional coefficients with the cloth materials we have chosen.

Experimental Setup. Using two translational stages (Thorlabs LTS-300M) in front of a black background, we control the drop and drag movement of the cloth for repeatability. The total movement is characterised as follows. First, the material is held 1 cm above the substrate and dropped by 6 cm in a movement, which accelerates from rest at $10\,\mathrm{mm/s^2}$ until it reaches $10\,\mathrm{mm/s}$. Subsequently, the system decelerates to reach 6 cm of total displacement. Immediately, the second motorised stage, for the horizontal displacement, starts the horizontal drag by accelerating at $10\,\mathrm{mm/s^2}$ until it reaches $10\,\mathrm{mm/s}$. This velocity is kept constant until the system starts decelerating to achieve a total displacement of 30 cm. We repeat this process forward and backward twice per experiment to allow the observation of a sufficient number of drag and wrinkling patterns.

Video Acquisition. With a calibrated camera, we record the cloth motion from the viewpoint shown in Fig. 4, which allows to observe wrinkling patterns. The whole video contains around 300 frames, corresponding to 2.4 frames per second. For each material/substrate pair, we repeat the experiment 5 times, leading to 5 synchronised videos.

Reference Friction Measurements. A quantitative evaluation of the predicted friction coefficients requires measurements of the friction coefficient μ for each material/substrate pair. Accurately measuring μ is an involved process in physics and mechanics which is often performed through invasive protocols. For simplicity, we capture merely a reference measurement for μ using the non-invasive inclined plane technique [12], [36]. We believe that the friction in the inclined plane scenario is close to the one in the drag situation, and hence its quantification provides a good reference estimation for our purposes. The inclined plane protocol measures friction by placing an object on an inclined plane, and by increasing the slope of the plane until the object starts slipping. The
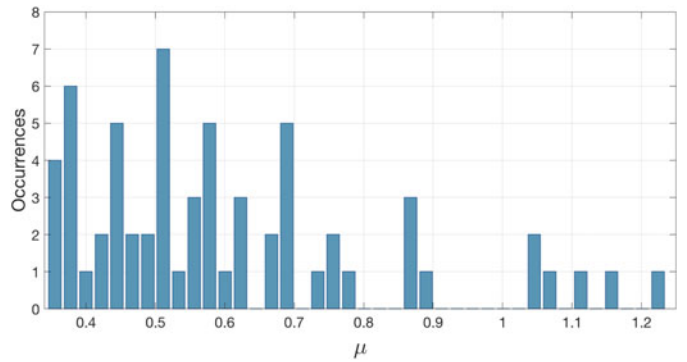

Fig. 2. Histogram of μ for all material / substrate pairs.

friction coefficient is then computed based on the slope of the plane at the point where the object slips.

To ensure that the reference measurement is robust, we test for each material/substrate pair different locations and orientations of the cloth on the substrate and do not find a noticeable difference in the slippage angle. Our physical setup of the inclined plane gives rise to a measurement error due to the discrete motion of the plane. More details are provided in Section S1.2 of the supplemental material, available online. Fig. 2 shows a histogram of the values of μ that were measured for all material/substrate pairs. Note that our dataset covers a wide range of friction coefficients.

4.2 Simulated Data Generation

For training and testing, we simulate a physically accurate dataset that closely resembles the videos captured using the experimental protocol. Implementing Amontons-Coulomb's law poses difficulties in practice because the force response is nonsmooth. This is further complicated by the requirement for discrete representations in both space and time. We use the ARGUS implementation [22] to simulate cloth deformation subject to frictional contact, for two reasons. First, this state-of-the-art simulator uses an efficient nonsmooth solver to model the dry frictional behaviour for mesh-based systems. Second, a free implementation of ARGUS is available on github.

4.2.1 Physical Validation of the ARGUS Simulator

To our knowledge, most frictional contact solvers for cloth, including ARGUS, have never been validated against real experiments. To verify that ARGUS produces physically accurate simulations, we compare simulations produced by ARGUS against physical experiments. This comparison is performed in a constrained setting, in which the friction behaviour can be derived analytically. The reason for restraining the experimental setting is that a verification against the ground truth physical behaviour is not possible in more complex scenarios, as measuring static friction accurately remains a challenging problem in physics and mechanics.

The constrained experiment we use has recently been proposed by Sano *et al.* [32] in the physics community, and is depicted in Fig. 3 and further illustrated in the supplemental material, available online for this paper. The experiment considers the deflection of a strip, clamped at its top, that is vertically pushed against a substrate with a vertical strain ϵ_y, as illustrated in Fig. 3-right. Because of friction, this strip remains *pinned* at its bottom for small ϵ_y, which is shown as the red region in Fig. 3-left. The more the strip is pushed down, the
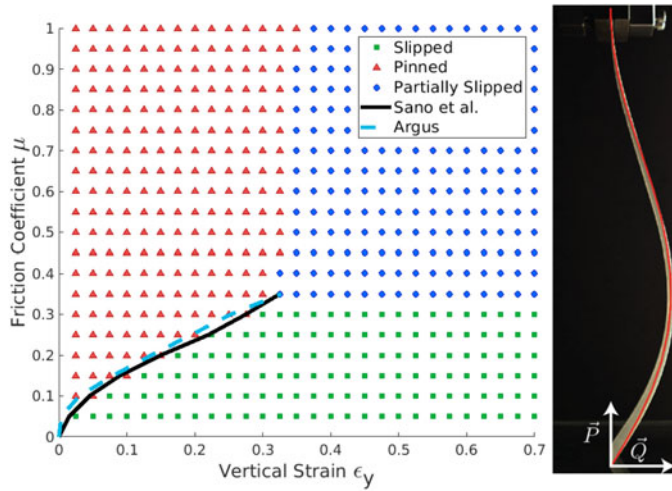
Fig. 3. Physical validation of Argus [22] under a constrained setting that is well understood in physics [32]. After proper calibration of the simulator, we observe simulations (dotted curve) that are in very good agreement with the theory (black curve).

higher the frictional force holding the strip pinned. The strip is geometrically constrained, hence depending on the value of the friction coefficient, this system has two different outcomes. First, the elastic forces can overcome friction, and the strip abruptly slips. This regime is called *slipped*, and illustrated by the green region in Fig. 3-left. Second, for rougher surfaces, the frictional forces are high enough to prevent the strip from slipping and it deforms until there is a real contact between the strip and the substrate. This regime is called *partially slipped*, and illustrated in blue in Fig. 3-left. Sano *et al.* show analytically that in the case where gravity is negligible, these outcomes are purely regulated by the friction coefficient, and independent of the strip's material parameters. That is, the deformed shape of the strip only depends on the vertical strain ϵ_y. The analytical solution allows to calculate at what vertical strain the strip will slip (and thus will move from the pinned to the slipped region), which is depicted as a black curve in Fig. 3-left.

When simulating the experiment of Sano *et al.*, a physically accurate simulator should simulate strips that slip when the vertical strain lies on the black curve. We succeeded in matching this theoretical result with the output of Argus, but this required carefully setting different numerical parameters, as explained in the following. First, the temporal resolution of the simulation can be controlled by a timestep parameter between two adjacent frames. Second, the spatial resolution can be controlled by either forcing a static discretisation to be used (whose resolution needs to be provided), or by allowing an adaptive remeshing of the geometry, in which case a rate of refinement needs to be provided. Finally, Argus also provides contact solver tolerance values that can be controlled by a residual value and a maximum iteration count of the solver, and damping forces are allowed to stabilise the simulations. In our experiments, we found that all of these numerical parameters heavily influence the geometry of the resulting simulations, and thus should be chosen carefully. Our rationale for finding an adequate set of numerical parameters was, first, to set a fixed and high resolution for the cloth mesh, avoiding small energy perturbations due to remeshing; then, mimic a quasi-static experiment by increasing damping forces, thus

removing spurious dynamical vibrations of the cloth; finally, find the right level of accuracy by sufficiently decreasing the timestep of the simulation, as well as the solver tolerance, to the point where further refinement would lead to indistinguishable results at our observation scale. In our comparisons, we use a timestep of 0.5 milliseconds, a mesh resolution of 149 vertices for a 20 cm long ribbon, a maximum number of iterations of 50000 and a damping value set to 0.002. We repeat the simulation using 2 materials having different Young's moduli. Our simulations yield the aqua curve shown in Fig. 3-left, which matches the black analytic curve almost perfectly. Fig. 3-right shows an overlay of the physical strip with our simulated result shown as a red curve, and they are also in agreement. Animated illustrations of the same is provided in the supplemental material, available online for this paper.

These results certify that under correct numerical calibration, Argus generates simulations that match macroscopic observations in a constrained setting, under negligible effect of gravity. Note that while generating our training dataset, we work under the hypothesis that the same numerical calibration will guarantee highly realistic results while simulating cloth in a more general setting, i.e., when gravity is no more negligible, and where both friction and material parameters influence the geometry of the cloth. Our good prediction results presented in Section 6.2 confirm that this hypothesis is reasonable. As the calibrated parameter settings make the simulation run-time significantly slower, conversely, one could argue that our calibration process imposes an overly strong constraint for our solver, hence an unnecessarily large training cost. By generating training datasets with varying accuracy, we check in the following that our high accuracy calibration is actually mandatory to obtain satisfying prediction results from real data. This comparison is presented in Section 6.2.

4.2.2 Dataset Generation

To generate the cloth simulation dataset, we implement the exact path of the motorised stage used for the physical data captures, which is analytically accessible, leading to a high temporal synchronisation between the real data and the simulations. Furthermore, the calibrated camera parameters of the real data captures are used to render a similar view for our simulations. This results in simulations that can be considered physically valid to train a discriminatory model.

In order to ascertain the usefulness of our validation experiment, described in the previous section, we generate 3 datasets with varying simulator parameters. For all datasets we use a high resolution mesh. However, we degrade the accuracy of the generated data by relaxing the solver tolerance and the simulation timestep. The solver tolerance is relaxed by keeping the residual value constant at $1e^{-15}$ and by lowering down the number of solver iterations (which results in a increase of the actual residual at each timestep). The parameters used for each dataset are described in Table 1. Out of the 3 datasets, High-accurate-solver-small-timestep dataset corresponds to calibrated settings obtained from the validation experiment.

In each dataset the drop and drag motion is simulated for the 10 material classes measured by Wang *et al.* [39]. For each material, 16 friction coefficients evenly distributed between 0.0 and 1.5 are explored, which represents a

TABLE 1
Parameter Specifications of 3 Datasets, Generated for Varying Levels of Accuracy

Dataset	Solver Iterations	Timestep
Low-accurate-solver-big-timestep	2000	1 ms
Low-accurate-solver-small-timestep	10000	0.5 ms
High-accurate-solver-small-timestep	50000	0.5 ms

reasonable range for fabrics according to values tabulated in [16] and is in agreement with our reference measurements. The resulting simulated 3D sequences contain 300 frames each. To generate a dataset of 2D videos, each simulated sequence is rendered using 8 different texture maps and from 3 different viewpoints and varying lighting conditions, using the free Blender software. The addition of texture variation increases generalisation of learning as shown in Section 6.3. One of the rendered viewpoints is calibrated based on the real data to replicate the experimental settings in our simulated data. Two additional viewpoints, are selected to increase visual variation. Furthermore, to match the experimental setup and reduce variability due to environmental factors, we render a dark background and a substrate floor in all of our renderings. The viewpoints for simulated data, alongside real data, are depicted in Fig. 4. Further illustrations of the dataset are provided in the supplemental material, available online.

5 PREDICTION MODEL ARCHITECTURE

We now specify our pipeline to estimate friction of cloth from a 2D input video. As the input to our model is a sequence of images, we use a Long-term Recurrent Convolutional Network (LRCN) model similar to those used for action classification [15]. We are inspired by a previous model that recovers a material class label of cloth from an input image sequence [42]. A variant of this model, called *baseline model* in the following, is explained in Section 5.1.

While the baseline model is effective at predicting a material class, it performs significantly worse when trained to predict friction coefficients (see Section 6.3 for details). We therefore propose a novel pipeline for this task that conditions friction on material classes, as outlined in Section 5.2.

5.1 Baseline Model

In the baseline model by Yang *et al.* [42], convolution and pooling layers are used to extract image features from each frame. Their architecture of this feature extraction block is a modified version of AlexNet [20]. We replace this architecture by a simplified version of VGG16 [34] with ReLU activations as shown in Table 2. The weights are shared between all feature extraction blocks. Let

$$f_i = CNN_{VGG}(I_i), \qquad (1)$$

denote the image features extracted from frame I_i, where CNN_{VGG} is the simplified VGG16 of Table 2.

A sequence of learned image features $f_1, f_2, \ldots, f_k$ is then passed to long short term memory (LSTM) layers, which extract temporal information. The output of these layers is finally passed to fully connected layers to learn a function from the extracted spatial and temporal features to the data labels. This can be written as
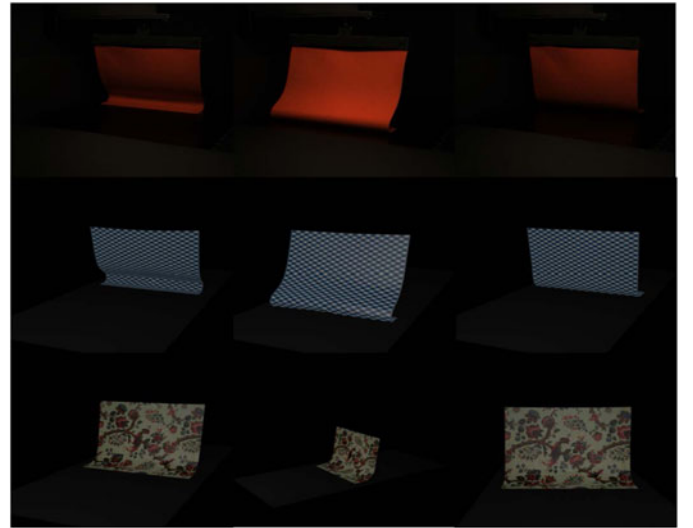

Fig. 4. Dataset Examples: First and second row show corresponding frames from real and synthetic data respectively. Third row shows 3 viewpoints rendered in the simulated dataset.

TABLE 2
Architecture Details for a Feature Extractor Block

Name	Description
Input	Image 224, 224, 3
conv1a	Conv 3 x 3, 64, ReLU
conv1b	Conv 3 x 3, 64, ReLU
maxpool1	3 x 3, stride 2 x 2
conv2a	Conv 3 x 3, 128, ReLU
conv2b	Conv 3 x 3, 128, ReLU
maxpool2	3 x 3, stride 2 x 2
conv3a	Conv 3 x 3, 256, ReLU
maxpool3	3 x 3, stride 2 x 2
conv4a	Conv 3 x 3, 512, ReLU

$$\hat{h} = FC(LSTM(f_1, f_2, \ldots, f_k)), \qquad (2)$$

where $\hat{h}$ is the final likelihood computed for each material class label, $LSTM$ is a set of two LSTM layers and FC is a set of two fully connected layers. The architecture is depicted in Fig. 5. This architecture is trained with a standard categorical cross-entropy classification loss, and for prediction, the class label with the highest likelihood is reported.

5.2 Conditional Friction Model

The baseline model is significantly worse at predicting friction than at predicting material class. The reason is that different materials combined with different friction values can yield visually similar features. For a fixed material, however, different friction values typically lead to visually distinctive behaviours. Motivated by this observation, we estimate the friction coefficient using a model that is conditioned on material classes.

Model. The model is shown in Fig. 5. In addition to a sequence of 2D video frames $I_i, i = 1, \ldots, k$, the model takes as input a material class label. The material class label is represented by a one-hot vector m which is passed to a fully connected layer with softmax activation. For this architecture, the video frames are processed using the same
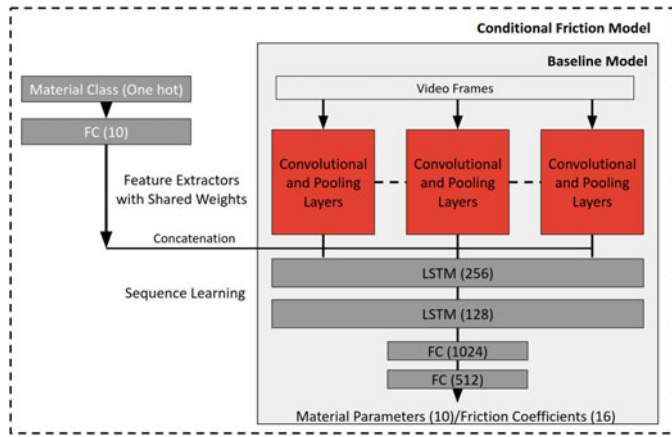
Fig. 5. Proposed architecture to estimate friction conditioned on material parameters. The coloured inlay shows the baseline model, which is augmented with material class information to form the conditional friction model.

convolution and pooling layers as for the baseline model. The material information is then cloned for each input frame and concatenated with the feature vectors of each frame, before being passed to the LSTM layers and on to the fully connected layers. That is, the vector $\hat{y}$ containing the likelihood for each friction class label is computed as

$$\hat{y} = FC(LSTM(FC_s(m) \frown f_1, \ldots, FC_s(m) \frown f_k)), \quad (3)$$

where FC_s denotes a fully connected layer with softmax activation and $\frown$ is the concatenation operator.

Training. The training loss can be written as

$$\mathcal{L} = -log P(y|I_1, I_2, \ldots, I_k, m), \quad (4)$$

where y is the friction label provided for the training examples. This loss is implemented as the categorical cross-entropy loss function. In both models, we use dropout layers between fully connected layers for regularisation. We train the baseline model to predict the material class. We use true material class labels to train the conditional friction model.

Data Representation. A data point for training or testing consists of 30 frames sampled at regular intervals from a video sequence and corresponding material class and friction coefficient as label. We select an input size of 30 frames as the maximum number of frames from a single video sequence that we can fit during the training cycle on a single GPU, without running into memory constraints. We use a split of $80, 10, 10$ percent for training, cross-validation and testing, respectively, which results in training on 92160 images from 3072 video sequences and testing on 11520 images from 384 video sequences for each dataset.

Model Initialisation. We observed experimentally that model initialisation is important for training convergence. Intuitively, this might be due to the fact that certain viewpoints provide better discriminatory information than others. To obtain training convergence in practice, we therefore train our models progressively by adding one rendered viewpoint from the dataset in each training cycle to our training data. After training simultaneously on all viewpoints, we finely tune our model on the viewpoint which is calibrated based on our experimental setup.

Prediction. At test time, we provide the true material class label to the conditional friction model in order to observe the effect of lower to higher accuracy friction information present in our datasets. For real captured data, we provide the model with closest matching material class label that is present in our synthetic dataset. For a class label that is not present in our synthetic training datasets i.e., silk, we provide the conditional model with a material class label predicted by our baseline model.

6 RESULTS AND DISCUSSION

In this section we present and analyse the results of testing our model on real data, before discussing limitations. We also perform ablation studies on synthetic data, in order to perform a comparison between our conditional friction estimator and baseline models, and to evaluate the model's generalisation capability.

6.1 Implementation and Evaluation Details

Our implementation uses Keras [10] and Tensorflow [1]. While we experimented with various optimisers, we empirically observed Adadelta [45] to converge faster. We use a learning rate of 1.0 and a decay factor of 0.95. The training time of our model and the baseline is between 4 to 6 hours on all datasets with a single NVIDIA TitanX GPU and the training converges in circa 30 epochs.

Evaluation Protocol. As our goal is to build a friction measurement protocol for cloth, we evaluate our friction prediction by considering the absolute difference between the predicted value and the reference value. As the reference measurement r for real data is only known up to a measurement error e, we consider any value within the interval $r \pm e$ as having no error, and report the absolute difference of our prediction to this interval. This provides an optimistic estimate of the error as the error is calculated from the range and not from an absolute value. For material parameter estimation we report the top-1 and top-2 accuracy.

6.2 Results on Real Test Data

We present results on real data captured through our experimental setup presented in Section 4.1. Our datasets contains 5 videos for each material / substrate pair, leading to 5 test datapoints. We take the median prediction value out of the five test datapoints to be the final prediction for any material-substrate pair. We evaluate our test data using models trained on 3 different datasets generated using the ARGUS simulator as mentioned in Section 4.2.2.

Impact of the Simulator's Accuracy. As can be observed in Fig. 6, the model trained on the dataset High-accurate-small-timestep outperforms the models trained on the two other datasets Low-accurate-big-timestep and Low-accurate-small-timestep (error of <0.1 on 65.07 percent of test data compared to 25.3 and 34.23 percent, respectively). Recall that the former dataset is generated using our calibrated parameters settings based on the validation experiment described in Section 4.2.1.

Furthermore, it can be observed that while the <0.1 error is not significantly different between models trained on datasets Low-accurate-big-timestep and Low-accurate-small-timestep, the model trained on the latter makes significantly less catastrophic errors, thus achieving an error of <0.2 on 52.3 percent of test data as compared to 33.3 percent of the model trained on the former.
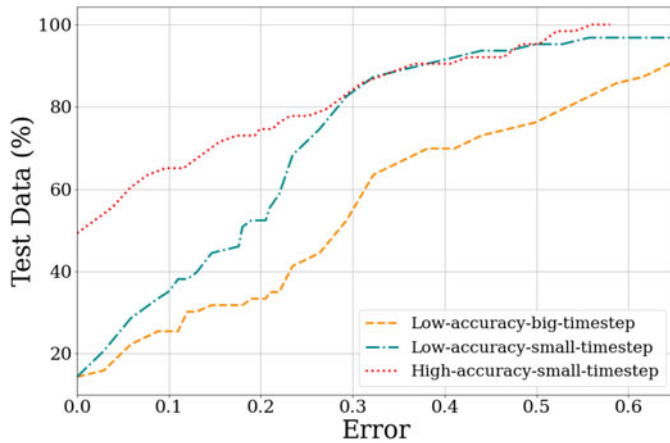
Fig. 6. Cumulative error plots for all datasets.

Fig. 7. Cumulative error plots for dataset High-accurate-small-timestep (red, black, light blue, green) and simulated test data (dark blue).

Overall, this study shows that degrading the accuracy of the simulator causes a downgrade in prediction of our learning-based method, hence it confirms the need for a realistic enough simulator to generate input data. This justifies the need for validating and calibrating the ARGUS simulator properly before learning from it.

Impact of the Reflectance.

Fig. 7 shows a detailed breakdown of results for the model trained on the dataset High-accurate-small-timestep. It is noteworthy that the model accuracy reached on real experiments when tested on all substrates (red curve, error of <0.1 on "only" 65.07 percent of data) is significantly lower than the model accuracy obtained on synthetic test data (dark blue curve, see details in next subsection). This decay in model accuracy can be mostly attributed to the presence of reflecting substrates in the real data, a factor that is not modelled in our training data. The light blue curve shows that the model accuracy is significantly worse on substrates which are reflective. Removing these substrates from the evaluation yields an error of <0.1 for 86.1 percent of data, as shown in the black curve. This shows that the model has indeed learnt the friction behaviour of cloth which is transferable from simulated to real data, barring the confounding factors. One way to mitigate this degeneration on reflective substrates is to vary the reflectance of the substrate floor during rendering of our simulated training data. After adding reflective substrate renderings, we obtain an error of <0.1 for 93.6 percent of real test data on all substrates, as shown in the green curve. Note that adding the variation to the training data improves the results overall, yielding better results than training and testing without reflective materials (black curve).

6.3 Results on Simulated Data

We perform ablation studies on simulated data in order to assess our trained model's ability to generalise for our trained model. In the following experiments, we use a model trained on our calibrated simulated dataset (High-accurate-small-timestep).

6.3.1 Utility of the Conditional Friction Model

To compare the performance of our conditional model with the baseline model, we train our model on synthetic data while leaving out 10 percent of our dataset for testing.
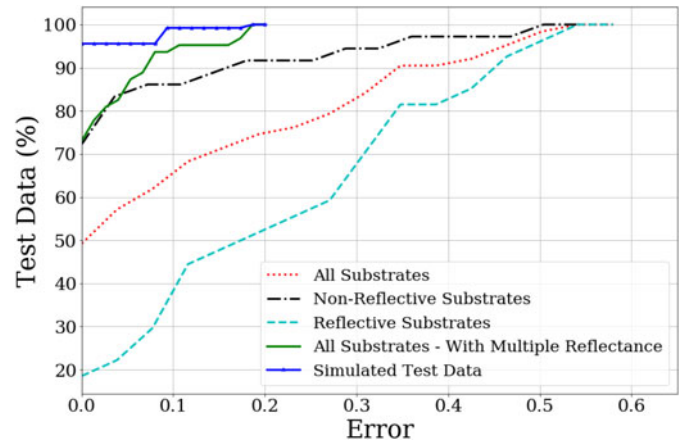
As shown in Table 3, baseline model performs vastly better at predicting material parameters than at predicting friction coefficients. The conditional estimation model performs better at predicting friction coefficients than the baseline model which indicates that the material and friction parameters are not decoupled in the global behaviour of the cloth, and that adding material information as an input parameter reduces the search space. Furthermore, the error distribution for test data indicates that our model learns coherently.

6.3.2 Generalisation to Unseen Textures

We render our training data with different textures to make the model agnostic to appearance variations. To test this generalisation ability, we render our simulations with a texture that has not been seen by the model during training. We test our model on 160 sequences, and our model achieves an error of <0.1 on 87.2 percent, and an error of <0.2 on 95 percent of the samples. This shows that the model's predictions do not degrade dramatically if an unseen texture is encountered, implying that the model is capturing more nuanced spatio-temporal phenomena to classify friction behaviour.

6.3.3 Generalisation to Unseen Viewpoints

We demonstrate the model's generalisation ability to unseen viewpoints. Starting with the camera position from one of our simulated viewpoints, we rotate the camera origin, in both azimuthal and transverse directions by ±5 degree intervals on either side to generate 6 viewpoints which are 5, 10 and 15 degrees apart on either side. Afterwards we randomly select 20 material, friction and texture combinations for each viewpoint and render them as our test data. We report the accuracy for each viewpoint in Table 4. While the accuracy degenerates with unseen viewpoints, the decay happens progressively based on the difference in viewpoint from the original one. The asymmetry of results in different directions of the azimuthal movement can be attributed to the fact that in one direction less of the cloth contacting the substrate is visible as compared with the other one. Furthermore, during horizontal movement, frontal viewpoints display better accuracy as opposed to side viewpoints.

TABLE 3
Results on Simulated Test Data

Task	Material Estimation		Friction Estimation	
	Acc	Top-2 Acc.	Err <0.1	Err <0.2
Baseline	99.5 %	100.0 %	78.6 %	88.5 %
Conditional	-	-	98.3 %	100 %

TABLE 4
Results on Unseen Viewpoints

Horizontal	-15°	-10°	-5°	5°	10°	15°
Error < 0.1	65%	75%	80%	75%	65%	50%
Error < 0.2	70%	85%	95%	80%	80%	65%
Azimuthal	**-15°**	**-10°**	**-5°**	**5°**	**10°**	**15°**
Error < 0.1	60%	80%	90%	80%	55%	40%
Error < 0.2	80%	80%	100%	85%	70%	60%

6.3.4 Generalisation to Unseen Rendering Conditions

As our training data contains a variation of rendering conditions, we demonstrate the model's generalisation ability by arbitrarily changing the lighting conditions and reflectance of the substrate. We train our model without adding additional reflectance variance mentioned in Section 6.2. We test our model on 160 sequences, and our model achieves an error of <0.1 on 72.7 percent, and an error of <0.2 on 84.1 percent of the samples. This degradation is in line with the results observed on real captured data. The rendering condition variability is thus necessary to increase the generalisation of the model.

7 EXTENDING TO CLOTH ON CLOTH FRICTION

We extend our proposed non-invasive parameter estimation model to cloth on cloth motion data captured by a physical experimental setup. Estimating cloth on cloth friction has been a relatively more difficult task due to the woven nature of the materials. In the following sections, we describe the modalities for generating our experimental data, extending our synthetic dataset and the results from applying our method in cloth on cloth friction regime.

7.1 Cloth on Cloth Friction Test Data

Using the same setup as detailed in Section 4.1 we generate a new test dataset, using our 9 cloth samples from the previous experimental dataset. We now use each cloth both as a substrate and as the suspended cloth for the 'drop and drag' motion, thus having 81 cloth-cloth pairs in total. As a substrate, the cloth is stretched out on to the base of the experiment, such that the suspended cloth can drag on it. The motion is further explained in Section 4.1. In total we obtain 5 videos of 300 frames each of 81 cloth-cloth pairs. An extract from this dataset is shown in Fig. 8. Note that our setup is not symmetric, as the roles played by the suspended cloth and the cloth substrate are not equivalent due to the stretching of the cloth substrate. This asymmetry will be revealed by our friction measurements below.

In addition to our main setup, we record a baseline measurement of friction of each cloth-cloth pair based on the
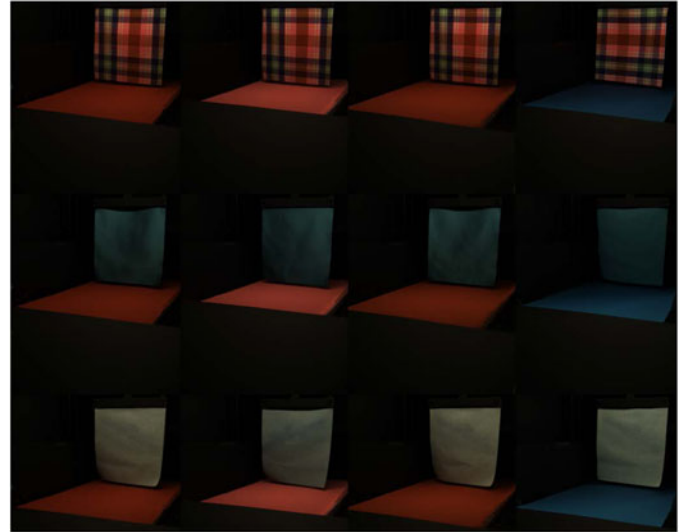

Fig. 8. The figure shows an excerpt from our new Cloth to Cloth friction evaluation dataset. This dataset includes a cloth stretched out as a substrate and another cloth sliding over it using a motorised clamp.

inclined plane method described in Section 4.1, using a similar stretching for the cloth serving as a substrate. As our substrate is a non-smooth surface, some cloth pairs display a thread interlocking behaviour such that a stable measurement of a baseline value is not possible. Thread interlocking is a phenomenon that lies outside the scope of our friction model. Therefore, we omit these pairs from our evaluation, considering only 5 cloth substrates, and as a result, only 45 cloth-cloth pairs out of the 81 initial ones. The distribution of measured baseline values for this reduced dataset is provided in Fig. 9

Out of the measured baseline values, certain cloth pairs display a friction measurement higher than 1.5. Such values, shown as orange bars in Fig. 9, clearly lie outside the purview of our synthetic training dataset. We thus expect our prediction method to fail on such data. Finally, as anticipated before, it is noticeable on the distribution of friction coefficients that symmetry is not preserved between a cloth pair in terms of friction values. This is understandable given the nature of the setup, the cloth used as substrate being stretched out while the cloth at the top sliding across it.

7.2 Extending the Synthetic Dataset

As our synthetic dataset does not include coloured substrates, in order to improve our model's capacity to generalise to cloth on cloth friction test data we extend our synthetic dataset as follows. We use the dataset High-accuracy-small-timestep and out of the simulation sequence of each friction coefficient (0.0 to 1.5) we take each material and render it with a random colour chosen for both cloth and substrate out of the colours picked from our real cloth dataset. Images rendered from simulation sequences of the same friction coefficient are presented in Fig. 10.

7.3 Evaluation and Discussion

In order to estimate friction on the cloth on cloth data, we use a model pre-trained on the synthetic dataset High-accuracy-small-timestep. We further fine tune this model by training it on additional synthetic renderings mentioned in
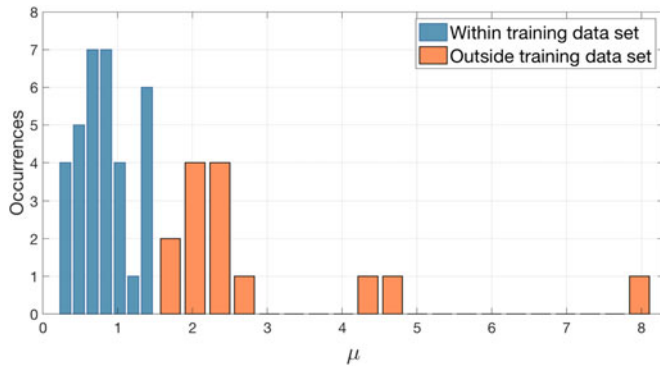
Fig. 9. Distribution of baseline friction values for cloth to cloth dataset. The orange bars are outside the range of our synthetic training data.
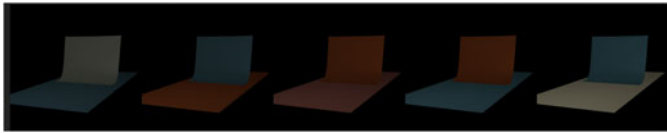

Fig. 10. Extension of our synthetic dataset, with renderings more similar to the new cloth on cloth dataset.
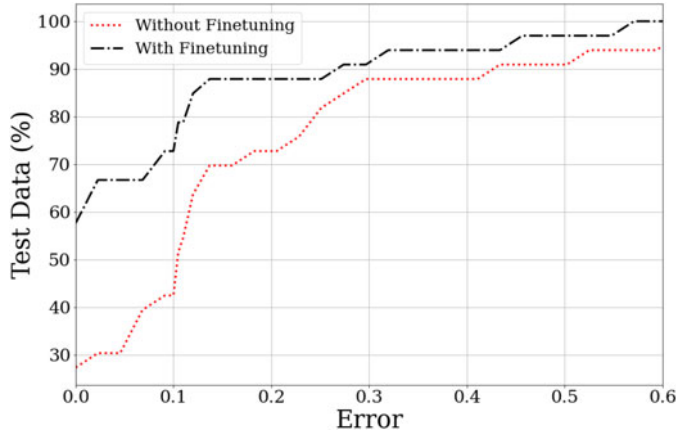

Fig. 11. Prediction errors on Cloth to Cloth friction test dataset.

the previous section. We present prediction error plot from both before and after fine-tuning our model in Fig. 11

In Fig. 11, we show results on test data that is within the scope of our training dataset. Before fine-tuning the model predicts an error of <0.2 on 72.7 percent of the data and with fine-tuning the model improves to an error of <0.2 on 87.8

percent of the data and an error of <0.1 on 72.7 percent of test data. This is comparable with the results achieved for our previous experimental test dataset discussed in Section 6.2. If we include friction values that are beyond the scope of our training data i.e., >1.5, we achieve an error of <0.2 on 66.6 percent and <0.1 on 53.4 percent of test data using the model fine tuned on the extended synthetic dataset. Given the fact that estimating cloth on cloth friction is a relatively more difficult task, our method provides reasonably accurate results on the test dataset while being trained purely on synthetic data.

8 APPLICATION TO CLOTH SIMULATION

Our measurement method can be applied successfully to the simulation of cloth with realistic materials. Instead of testing and trying various simulation parameters independently to obtain a desired effect, as often done in Computer Graphics, here the user can directly pick real materials for both clothes and external bodies, and simply inject the corresponding measured parameters in the simulator. Resulting simulations then capture cloth dynamics as well as cloth-cloth and cloth-body interactions in a consistent fashion, since all the parameters only stem from the chosen materials. To illustrate such a benefit, we perform two kinds simulations, shown in Figs. 12 and 13.

First, we simulate two square pieces of cloth of different material such that each one is dropped on a rotating sphere made of a particular substrate, and dragged on the floor. One scenario was performed by choosing Sparkle Sweat (M02) for cloth and Aluminium for the sphere, and another one by choosing Denim (M03) combined with Ceramic. These scenarios take as friction coefficients 0.3 and 0.6, respectively, which exactly correspond to our measurements (see Table S6 in supplemental material, available online). As shown in Fig. 12, the wrinkle patterns in both experiments differ significantly given the choice of material and friction coefficients.

Second, we demonstrate a realistic interaction between real life cloth materials on an animated character. We simulate two sequences of cloth on cloth friction using our test materials Sparkle Sweat (M02) and Tango Red (M06) with the ARGUS simulator. In each sequence a shawl is dropped on a clothed animated character.

As shown in Fig. 13, the behaviour of cloth to cloth contact is dictated by the specified friction coefficient. The figure shows corresponding frames from two simulated sequences. The materials corresponding to M02 and M06 have been used for both garments in the top and bottom rows respectively.
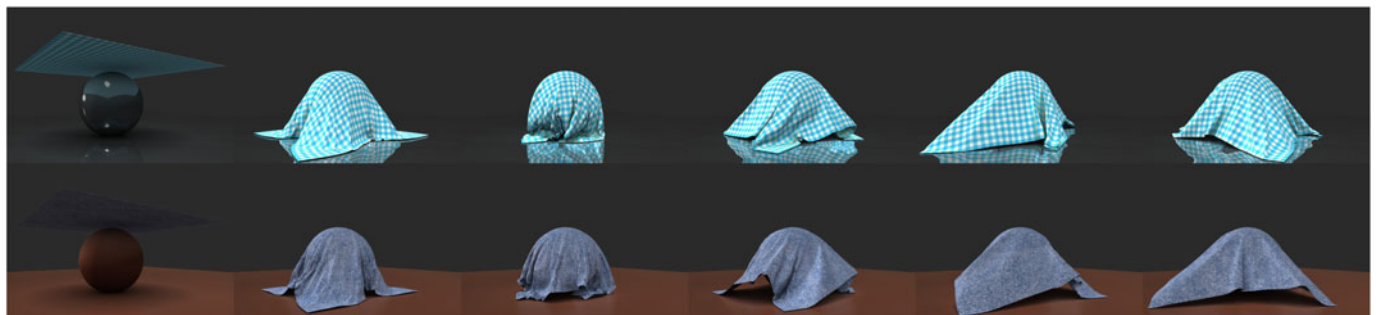

Fig. 12. Corresponding frames showing the cloth coming into contact with a rotating sphere substrate. Top Row: M02 (Polyester 95 percent, Spandex 5 percent) with $\mu = 0.3$, substrate: aluminium, Bottom Row M03 (Cotton 100 percent) with $\mu = 0.6$, substrate: ceramic.
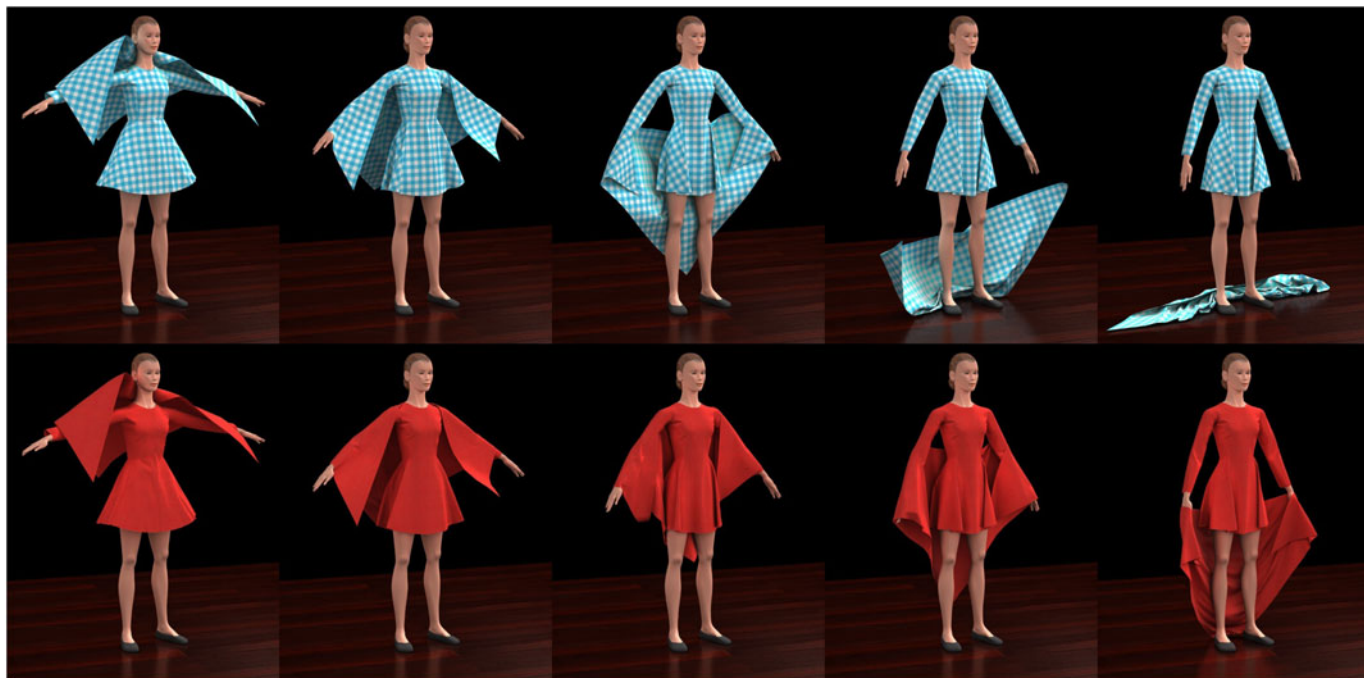
Fig. 13. Corresponding frames of a simulated character displaying cloth on cloth friction. Top Row: M02 (Polyester 95 percent, Spandex 5 percent) with $\mu = 0.3$, Bottom Row M06 (100 percent Polyester) with $\mu = 0.7$.

The friction coefficients used are the ones predicted by our measurement protocol on the test data, i.e., 0.3 and 0.7 (see Table S9 in supplemental material, available online).

9 CONCLUSION

We propose the first protocol for vision-based measurement of dry friction in cloth coming into contact with a substrate and with another cloth. This protocol relied on a CNN-based model trained on data generated by a physics based simulator which has been verified using a physical experiment. In this extended paper, we conclusively show the effect of performing such verification on the simulator before using it to generate training data. We also demonstrate the measurement of cloth on cloth friction using the same protocol and release a new dataset for the community. These contributions open interesting future directions. First, our method paves the way towards estimation of friction in-the-wild by progressively relaxing the video-acquisition protocol, with interesting applications for non-invasive physics measurements, finer-grain capture of real surfaces, and physically accurate re-simulations of pre-observed surfaces. Second, the successful use of calibrated synthetic simulator-based training could be transposed to other inverse parameter estimation problems.

REFERENCES

[1] M. Abadi *et al.*, "TensorFlow: Large-scale machine learning on heterogeneous systems," Software, 2015. [Online]. Available: tensorflow.org

[2] G. Amontons, "De la résistance causée dans les machines," *Mémoires De l'Académie Royale*, vol. 4, pp. 257–282, 1699.

[3] K. S. Bhat, C. D. Twigg, J. K. Hodgins, P. K. Khosla, Z. Popović, and S. M. Seitz, "Estimating cloth simulation parameters from video," in *Proc. ACM SIGGRAPH/Eurographics Symp. Comput. Animation*, 2003, pp. 37–51.

[4] K. L. Bouman, B. Xiao, P. Battaglia, and W. T. Freeman, "Estimating the material properties of fabric from video," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2013, pp. 1984–1991.

[5] M. Brando, K. Hashimoto, and A. Takanishi, "Friction from vision: A study of algorithmic and human performance with consequences for robot perception and teleoperation," in *Proc. IEEE-RAS 16th Int. Conf. Humanoid Robots.*, 2016, pp. 428–435.

[6] R. Bridson, R. Fedkiw, and J. Anderson, "Robust treatment of collisions, contact and friction for cloth animation," *ACM Trans. Graph.*, vol. 21, no. 3, pp. 594–603, 2002.

[7] R. Candelier, A. Prevost, and G. Debrégeas, "The role of exploratory conditions in bio-inspired tactile sensing of single topogical features," *Sensors*, vol. 11, no. 8, pp. 7934–7953, 2011.

[8] Z. Chen, R. Feng, and H. Wang, "Modeling friction and air effects between cloth and deformable bodies," *ACM Trans. Graph.*, vol. 32, no. 4, pp. 1–4, 2013.

[9] N. R. Chevalier, "Hair-on-hair static friction coefficient can be determined by tying a knot," *Colloids Surfaces B: Biointerfaces*, vol. 159, pp. 924–928, 2017.

[10] F. Chollet *et al.*, "Keras," 2015. [Online]. Available: https://keras.io

[11] D. Clyde, J. Teran, and R. Tamstorf, "Modeling and data-driven parameter estimation for woven fabrics," in *Proc. ACM SIGGRAPH/Eurographics Symp. Comput. Animation*, 2017, Art. no. 17.

[12] R. Courtel and L. Tichvinsky, "A brief history of friction," *Naval Eng. J.*, vol. 76, no. 3, pp. 451–460, 1964.

[13] P. R. Dahl, "A solid friction model," Aerospace Corp El Segundo Ca, El Segundo, CA, USA, Tech. Rep. TOR-0158 (3107–18), 1968.

[14] A. Davis *et al.*, "Visual vibrometry: Estimating material properties from small motions in video," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 4, pp. 732–745, Apr. 2017.

[15] J. Donahue *et al.*, "Long-term recurrent convolutional networks for visual recognition and description," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 2625–2634.

[16] E. C. Dreby, "A friction meter for determining the coefficient of kinetic friction of fabrics," *J. Res. Nat. Bur. Standards*, vol. 31, no. 4, 1943, Art. no. 237.

[17] K. Erleben, M. Macklin, S. Andrews, and P. G. Kry, "The matchstick model for anisotropic friction cones," *Comput. Graph. Forum*, vol. 39, no. 1, pp. 450–461, 2020.

[18] M. W. Gondal *et al.*, "On the transfer of inductive bias from simulation to the real world: A new disentanglement dataset," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 15714–15725.

[19] S. James *et al.*, "Sim-to-real via sim-to-sim: Data-efficient robotic grasping via randomized-to-canonical adaptation networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 12627–12637.

[20] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2012, pp. 1097–1105.

[21] S. Kunitomo, S. Nakamura, and S. Morishima, "Optimization of cloth simulation parameters by considering static and dynamic features," in *Proc. ACM SIGGRAPH Posters*, 2010, pp. 15:1–15:1.

[22] J. Li *et al.*, "An implicit frictional contact solver for adaptive cloth simulation," *ACM Trans. Graph.*, vol. 37, no. 4, 2018, Art. no. 52.

[23] J. Liang, M. C. Lin, and V. Koltun, "Differentiable cloth simulation for inverse problems," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 771–780.

[24] X. Liu, Z Yue, Z. Cai, D. G. Chetwynd, and S. Smith, "Quantifying touch-feel perception: Tribological aspects," *Meas. Sci. Technol.*, vol. 19, 2008, Art. no. 084007.

[25] N. Mao, Y. Wang, and J. Qu, "Smoothness and roughness: Characteristics of fabric-to-fabric self-friction properties," in *Proc. 90th Textile Inst. World Conf.*, 2016, pp. 125–135.

[26] E. Miguel *et al.*, "Data-driven estimation of cloth simulation models," *Comput. Graph. Forum*, vol. 31, pp. 519–528, 2012.

[27] E. Miguel *et al.*, "Modeling and estimation of internal friction in cloth," *Trans. Graph.*, vol. 32, no. 6, pp. 212:1–212:10, 2013.

[28] R. R. Moorthy and P. Kandhavadivu, "Surface friction characteristics of woven fabrics with nonconventional fibers and their blends," *J. Textile Apparel Technol. Manage.*, vol. 9, no. 3, pp. 1–14, 2015. [Online]. Available: https://ojs.cnr.ncsu.edu/index.php/JTATM/article/view/7307/4087

[29] R. Narain, A. Samii, and J. F. O'Brien, "Adaptive anisotropic remeshing for cloth simulation," *ACM Trans. Graph.*, vol. 31, no. 6, pp. 152:1–152:10, Nov. 2012.

[30] S. Pabst, B. Thomaszewski, and W. Straßer, "Anisotropic friction for deformable surfaces and solids," in *Proc. ACM SIGGRAPH/ Eurographics Symp. Comput. Animation*, 2009, pp. 149–154.

[31] A. H. Rasheed, V. Romero, F. Bertails-Descoubes, S. Wuhrer, J.-S. Franco, and A Lazarus, "Learning to measure the static friction coefficient in cloth contact," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 9909–9918.

[32] T. G. Sano, T. Yamaguchi, and H. Wada, "Slip morphology of elastic strips on frictional rigid substrates," *Physical Rev. Lett.*, vol. 118, 2017, Art. no. 178001.

[33] A. Selle, J. Su, G. Irving, and R. Fedkiw, "Robust high-resolution cloth using parallelism, history-based collisions, and accurate friction," *IEEE Trans. Vis. Comput. Graphics*, vol. 15, no. 2, pp. 339–350, Mar./Apr. 2009.

[34] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. 3rd Int. Conf. Learn. Representations*, 2015. [Online]. Available: http://arxiv.org/abs/1409.1556

[35] C. Syllebranque and S. Boivin, "Estimation of mechanical parameters of deformable solids from videos," *Vis. Comput.*, vol. 24, no. 11, pp. 963–972, 2008.

[36] G. H. Thorndike and L. Varley, "Measurement of the coefficient of friction between samples of the same cloth," *J. Textile Inst. Proc.*, vol. 52, no. 6, pp. P255–P271, 1961.

[37] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2017, pp. 23–30.

[38] B. Wang, L. Wu, K. Yin, U. M. Ascher, L. Liu, and H. Huang, "Deformation capture and modeling of soft objects," *ACM Trans. Graph.*, vol. 34, no. 4, 2015, Art. no. 94.

[39] H. Wang, J. F. O'Brien, and R. Ramamoorthi, "Data-driven elastic models for cloth: Modeling and measurement," *ACM Trans. Graph.*, vol. 30, 2011, Art. no. 71.

[40] J. Wu, J. J. Lim, H. Zhang, J. B. Tenenbaum, and W. T. Freeman, "Physics 101: Learning physical object properties from unlabeled videos," in *Proc. Brit. Mach. Vis. Conf.*, 2016, no. 12, pp. 39.1–39.12, Art. no. 39, doi: 10.5244/C.30.39.

[41] J. Wu, I. Yildirim, J. J. Lim, B. Freeman, and J. Tenenbaum, "Galileo: Perceiving physical object properties by integrating a physics engine with deep learning," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2015, pp. 127–135.

[42] S. Yang, J. Liang, and M. C. Lin, "Learning-based cloth material recovery from video," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 4383–4393.

[43] S. Yang and M. C. Lin, "MaterialCloning: Acquiring elasticity parameters from images for medical applications," *IEEE Trans. Vis. Comput. Graphics*, vol. 22, no. 9, pp. 2122–2135, Sep. 2016.

[44] W. Yuan, S. Wang, S. Dong, and E. Adelson, "Connecting look and feel: Associating the visual and tactile properties of physical materials," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 4494–4502.

[45] M. D. Zeiler, "Adadelta: An adaptive learning rate method," *CoRR*, vol. abs/1212.5701, 2012.

[46] H. Zhang, K. Dana, and K. Nishino, "Friction from reflectance: Deep reflectance codes for predicting physical surface properties from one-shot in-field reflectance," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 808–824.

Abdullah Haroon Rasheed received the master's degree in informatics with specialization in data science from INP Grenoble, Grenoble, France. He is currently working toward the PhD degree with INRIA Grenoble, Montbonnot-Saint-Martin, France. He is working simultaneously with the ELAN and Morpheo teams at INRIA on topics related to utilizing machine learning for solving inverse measurement and design problems in mechanics of slender and elastic objects.

Victor Romero received the PhD degree in physics from the Université Pierre et Marie Curie, Paris, France and the Universidad de Santiago, Santiago, Chile. He worked in contact and friction mechanics with Laboratoire Jean Perin, UPMC, Paris, and in Non-equilibrium matter at ESPCI, France. Currently, he works as a research engineer with ELAN Team, INRIA Grenoble Rhône-Alpes, France, in collaboration with Laboratoire Jean Le Rond D'Alembert, France. His main interests include the domains of Tribology and the elasticity of slender elastic structures.

Florence Bertails-Descoubes received the PhD degree in computer science from INP Grenoble, Grenoble, France, in 2006, and after a post-doc with the University of British Columbia. She is a senior researcher with Inria, and the head of the ELAN Team. She has joined Inria as a permanent researcher in 2007. Her research interests deal with the direct and inverse numerical modeling of complex mechanical phenomena, in particular slender elastic objects (rods, plates, shells) and frictional contact. Target applications of her work encompass computer graphics, virtual prototyping, and computational physics.

Stefanie Wuhrer received the PhD degree in computer science from Carleton University, Ottawa, Canada, in 2009, and after working as research associate with the National Research Council of Canada, and junior research group leader with Saarland University in Germany, she joined INRIA in 2015. She is a research scientist with INRIA Grenoble Rhône-Alpes in France, conducting her research in the Morpheo Team. Her research interests include 3D shape and 4D motion analysis, human body and face modeling, and geometry processing.

Jean-Sebastien Franco received the PhD degree from the Institut National Polytechnique de Grenoble, Grenoble, France, in 2005 with the INRIA MOVI Team. He is an associate professor of computer science with the Ensimag (School of Computer Science and Applied Mathematics, INP Grenoble University), and a researcher with the INRIA Grenoble Rhône-Alpes, France, with the Morpheo Team since 2011. His expertise and interest is in the field of computer vision, dynamic 3D modeling and 4D spatio-temporal modeling from multiple views, 3D interaction.

Arnaud Lazarus received the PhD degree in mechanical engineering from École Polytechnique, Palaiseau, France, in 2008. After two post-doctoral years in Paris, he joined the Massachusetts Institute of Technology as an associate researcher from 2010 to 2013. Since 2013, he is an assistant professor with Sorbonne Université, Paris, France, doing his research with Institut Jean le Rond d'Alembert. His current interests include the stability of dynamical systems and the mechanical behavior of slender elastic structures.

▷ **For more information on this or any other computing topic, please visit our Digital Library at** www.computer.org/csdl.