



HAL
open science

Inverse Dynamic Modeling of Cloth - Deep Learning using Physics based Simulations

Abdullah-Haroon Rasheed

► **To cite this version:**

Abdullah-Haroon Rasheed. Inverse Dynamic Modeling of Cloth - Deep Learning using Physics based Simulations. Computer Vision and Pattern Recognition [cs.CV]. Inria; Université Grenoble Alpes, 2021. English. NNT: . tel-03501532v1

HAL Id: tel-03501532

<https://inria.hal.science/tel-03501532v1>

Submitted on 25 Jan 2022 (v1), last revised 16 May 2022 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

THÈSE

Pour obtenir le grade de

DOCTEUR DE L'UNIVERSITE GRENOBLE ALPES

Spécialité: **Mathématiques et Informatique**

Arrêté ministériel: 25 mai 2016

Présentée par

Abdullah Haroon RASHEED

Thèse dirigée par **Florence BERTAILS-DESCOUBES**, DR, INRIA
et codirigée par **Jean-Sébastien FRANCO**, MCF, UGA et
Stefanie WUHRER, CR, INRIA

préparée au sein du Laboratoire **Institut National de Recherche
en Informatique et en Automatique**
dans l'**École Doctorale Mathématiques, Sciences et
technologies de l'information, Informatique**

**Modélisation dynamique inverse de tissus
- Apprentissage profond à l'aide de
simulations basées sur la physique**

**Inverse Dynamic Modeling of Cloth - Deep
Learning using Physics based Simulations**

Thèse soutenue publiquement le **9 décembre 2021**,
devant le jury composé de:

Mme Ming LIN

Professeuse des Universités, University of Maryland, Rapportrice

M. Ken MUSETH

Directeur de Recherche, NVIDIA, Rapporteur

M. Antoine CHATEAUMINOIS

Directeur de Recherche, CNRS - ESPCI Paris, Examineur

M. Franck HÉTROUY-WHEELER

Professeur des Universités, Université de Strasbourg, Président

Mme Florence BERTAILS-DESCOUBES

Directrice de recherche, INRIA, Directrice de thèse

M. Jean-Sébastien FRANCO

Maître de Conférences, Université Grenoble Alpes, Co-Directeur de
thèse

Mme Stefanie WUHRER

Chargée de Recherche, INRIA, Co-Directrice de thèse



Abstract

Inverse problems arise in various physical domains and solving them from real-world visual observations poses a significant challenge due to the high dimensional nature of the data. Furthermore gathering enough observations that a data driven model can accurately capture the complete distribution of a physical phenomenon is often intractable. In this work we use deep learning to solve inverse problems by applying two basic principles. Deep learning models can be trained using synthetic data generated from physics based simulations. And the employed simulator itself needs to be verified for physical accuracy thus allowing the model to learn the exact physical phenomenon that is desired.

To validate the simulator, we introduce rich and compact physical protocols, originally proposed in soft matter physics literature to measure physical parameters. These protocols can be easily replicated in a simulator to test the physical correctness of the model, and the validity of the simulator.

We solve the inverse measurement problem of estimating contact friction in soft-bodies which otherwise requires a specialized physics bench and entails tedious acquisition protocols. 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 [YLL17a; Dav+17; Bou+13], but this has never been applied to the friction estimation problem which results in even more subtle visual variations. Since acquiring a large dataset for this problem is impractical, we generate it using a frictional contact simulator. As the simulator has been calibrated and verified using controlled experiments, the results are not only visually plausible, but physically-correct enough to match observations made at the macroscopic scale. 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. We also acquire an extensive dataset of real world experiments for testing. Both the training and test datasets have been made freely available to the community.

We also utilize the same protocol for solving the inverse measurement problem of estimating the deformed curvature of a suspended Kirchhoff rod. In order to do such estimation on physical rods, we utilize a deep learning model to visually predict a curvature field from a suspended rod. As creating a dataset from physical rods (even if synthetically constructed), that faithfully covers a representative manifold of deformed curvatures is intractable, we rely

on generating such a dataset from a verified simulator. Our work shows a promising way forward for utilizing deep learning models as part of an inversion measurement pipeline.

Résumé

Des problèmes inverses surviennent dans divers domaines physiques et les résoudre à partir d'observations visuelles du monde réel pose un défi important en raison de la nature hautement dimensionnelle des données. De plus, rassembler suffisamment d'observations pour qu'un modèle basé sur les données puisse capturer avec précision la distribution complète d'un phénomène physique est souvent insoluble. Dans ce travail, nous utilisons l'apprentissage profond pour résoudre des problèmes inverses en appliquant deux principes de base. Les modèles d'apprentissage profond peuvent être entraînés à l'aide de données synthétiques générées à partir de simulations basées sur la physique. Et la précision physique du simulateur employé, lui-même, doit être vérifiée, permettant ainsi au modèle d'apprendre le phénomène physique exact souhaité.

Afin de valider le simulateur, nous introduisons des protocoles physiques riches et compacts, proposés à l'origine dans la littérature de physique de la matière molle pour mesurer des paramètres physiques. Ces protocoles peuvent être facilement répliqués dans un simulateur pour tester l'exactitude physique du modèle et la validité du simulateur.

Nous résolvons le problème de mesure inverse de l'estimation du frottement de contact dans les corps mous qui nécessite sinon un banc de physique spécialisé et un protocole d'acquisition fastidieux. Cela rend la perspective d'une technique de mesure purement non invasive basée sur la vidéo particulièrement attrayante. Des travaux antérieurs ont montré qu'une telle estimation basée sur la vidéo est réalisable pour les paramètres de matériaux en utilisant l'apprentissage profond [YLL17a; Dav+17; Bou+13], mais cela n'a jamais été appliqué au problème d'estimation de la friction qui entraîne des variations visuelles encore plus subtiles. Étant donné qu'il n'est pas pratique d'acquérir un grand ensemble de données pour ce problème, nous le générons à l'aide d'un simulateur de contact frictionnel. Comme le simulateur a été calibré et vérifié à l'aide d'expériences contrôlées, les résultats sont non seulement visuellement plausibles, mais suffisamment corrects physiquement pour correspondre aux observations faites à l'échelle macroscopique. Nous proposons à notre connaissance le premier réseau de mesure non invasif et un jeu de données d'entraînement synthétique adjacent pour estimer le frottement du tissu au contact, à la fois pour les contacts tissu-corps dur et tissu-tissu. Nous acquérons également un vaste ensemble de données d'expériences du monde réel pour les tests. Les ensembles de données de formation et de test ont été mis gratuitement à la disposition de la communauté.

Nous utilisons également le même protocole pour résoudre le problème de mesure inverse de l'estimation de la courbure déformée d'une tige de Kirchhoff suspendue. Afin de faire une telle estimation sur des tiges physiques, nous utilisons un modèle d'apprentissage profond pour prédire visuellement un champ de courbure à partir d'une tige suspendue. Comme la création d'un ensemble de données à partir de tiges physiques (même si elles sont synthétiquement construites), qui couvre fidèlement une variété représentative de courbures déformées est insoluble, nous comptons sur la génération d'un tel ensemble de données à partir d'un simulateur vérifié. Notre travail montre une voie prometteuse pour l'utilisation de modèles d'apprentissage profond dans le cadre d'un pipeline de mesure d'inversion.

Acknowledgement

This work would not have been possible without the support of my supervisors, Florence Bertails-Descoubes, Jean-Sébastien Franco and Stefanie Wuhler.

As the Persian polymath, Omar Khayyam, writes in an apropos quatrain.

این گنبدِ لاجوردی و زرین طشت
بسیار بگشت است و دگر خواهد گشت
یکچند ز اقتضای دورانِ قضا
مانیز چو دیگران رسیدیم و گذشت

Contents

1	Introduction	1
1.1	Challenges in Using Neural Networks	1
1.2	Our Approach	2
1.3	Physical Verification of the Simulator	2
1.4	Visual Measurement of Friction in Cloth	4
1.5	Curvature Estimation in Rods	5
1.6	Contributions	6
2	Related Work	9
2.1	Validation of Computer Graphics Simulators	9
2.1.1	Qualitative vs. Quantitative Validation Protocols	10
2.1.2	Fitting Simulator Parameters to Real Data	10
2.2	Estimating Friction and Material Properties of Cloth	11
2.2.1	Dry Friction Models	11
2.2.2	Friction Measurement	12
2.2.3	Material Parameter Estimation	12
2.3	Curvature Estimation for Suspended Kirchhof Rods	13
2.3.1	Curvature Estimation of Discrete Curves	13
2.3.2	Curvature Estimation Using Neural Networks	14
3	Physical Verification of the Simulator	15
3.1	Problem Definition	15
3.2	The ARGUS Simulator	16
3.3	Verification of the Friction Model	16
3.3.1	Amontons-Coulomb’s Law for Friction	17
3.3.2	The Stick-Slip Test	17
3.3.3	ARGUS Verification Results	21
3.4	Verification of the Bending Model	24
3.4.1	The Cantilever test	24
3.4.2	ARCSIM Verification Results	27
4	Friction Estimation in Cloth	31

4.1	Problem Definition	31
4.2	Experimental Test Data	33
4.2.1	Cloth materials	34
4.2.2	Substrates Selection	34
4.2.3	Experimental setup	36
4.2.4	Video acquisition	37
4.2.5	Reference friction measurements	37
4.3	Dataset Generation from Verified Simulator	39
4.4	Model Description	41
4.4.1	Baseline Model	41
4.4.2	Conditional Friction Model	43
4.5	Results Analysis	45
4.5.1	Results on Simulated Data	45
4.5.2	Results on Real Test Data	47
4.6	Application to Cloth Simulation	52
4.7	Limitations and Future Work	53
5	Curvature Estimation in Rods	55
5.1	Problem Definition	55
5.2	Background and Modalities	56
5.3	Experimental Data Capture	57
5.3.1	Images to 3D Centerline Pipeline	57
5.4	Synthetic Dataset Generation from Verified Simulator	59
5.4.1	Non-Dimensional Data Generation	60
5.5	Model Description	61
5.5.1	Baseline Model	62
5.6	Results and Analysis	62
5.6.1	Future Outlook	62
6	Conclusion and Future Work	67
	Bibliography	69

Introduction

Inverse problems appear in various domains of science. They generally involve the computation of some internal parameter of a mathematical model given external observations or the satisfaction of certain bounding constraints. In this work we concern ourselves with inverse measurement problems that appear in visual measurement of slender deformable structures, such as cloth and rods/filaments that may be subject to contact and friction. Visual measurement systems, while challenging, provide a non-invasive measurement protocol to determine certain properties of an object without disturbing its structural integrity. Furthermore such measurements can be used to simulate real objects with fidelity.

However, such measurement problems are essentially ill-posed due to the high dimensional nature of visual data. Owing to their flexible representation of high dimensional and fundamentally ambiguous functions, *deep neural networks* emerge as a promising solution to visual measurement problems. Such models perform exceedingly well on computer vision tasks due to their automatic feature extraction and efficient optimization techniques. Therefore, in this work we explore deep learning as the primary tool for solving our inverse measurement problems.

1.1 Challenges in Using Neural Networks

The use of neural networks presents the following two primary challenges:

- **Data Acquisition:** Neural network models require significant amounts of data for supervised training. It is specifically important for the dataset to be sampled over the complete distribution of the phenomenon being captured. Acquiring such a dataset in the real world can be quite challenging and in some cases an intractable task. In order to overcome this shortcoming and with recent development of accurate simulators, training on simulated data has become increasingly popular in various domains such as robotics [Tob+17; Goo+19; Che+19; ZQW20] and high energy physics [Car+19; CBL20]. However, reliance on simulation for training can lead to the following challenge of accuracy.

- **Accurate Capture of a Physical Quantity:** Tautologically, when generalizing to real world data for inference, training on simulated data can only be as accurate as the simulator itself. When measuring a physical quantity, as is primarily the case in our work, mere visual plausibility of simulated data does not remain sufficient and accurately inferring physical properties from visual data can prove quite challenging [Wu+16]. It is thus necessary that the simulator accurately models the physical quantity that the neural network has to learn to measure.

1.2 Our Approach

In order to cope with the challenges stated above we adopt a novel strategy in our work for creating visual measurement protocols. This involves the following:

- Select physical experimental protocols that define a master curve which delineates a threshold effect and isolates the physical phenomenon being studied thus defining the correct measurement for a given physical behaviour.
- Validate the simulator using the physical experimental protocol that verifies the physical model and implementation modalities used in the simulator. This involves replicating the physical experiment in the simulator and calibrating the free parameters until the correct physical behaviour is achieved.
- Make acquisitions of experimental test data in tightly bounded conditions, using an easily replicable experimental bench.
- Use the simulator to generate an extensive training dataset replicating the conditions of the experimental test data, while ensuring that the dataset is representative of the complete distribution of the parameters to be measured.
- Do inference using the trained model on experimental test data.

1.3 Physical Verification of the Simulator

In this work, which is detailed in Chapter 3 we establish the first building block of our approach, which is to introduce physical experiment protocols that can be used to verify the simulator behaviour. More specifically we apply the **Stick-Slip** [SYW17] and **Cantilever** [Duc06; Far17] tests to verify the static friction and elastic bending model of the ARGUS simulator [Li+18b], along with others.




Fig. 1.1: Depiction of the **Stick-Slip** experiment implemented, both physically and in simulation [Rom+21]

Although this work was first initiated in the context of calibrating the simulator to produce physically valid data for training a neural network model for inverse measurement [Ras+20]; later on it became part of an elaborate effort in collaboration with colleagues from the Elan team at Inria and from the experimental physics community [Rom+21] to introduce and enhance general protocols which have been known in the soft matter physics community, for quantitative verification in the context of Computer Graphics (CG) simulators.

CG simulators have traditionally been used to animate, for instance, hair, fur and virtual clothing [BW98; BWK03; McA+09; DBB11a; Kau+14; Dav20] mainly in content production for the entertainment industry. Quantitative evaluation of simulator correctness has thus not been the main thrust of the community. But as the entertainment industry moves towards increasingly complex scenarios of mixing simulated and physical reality it becomes more and more necessary to achieve a certain level of quantitative realism [Mus20]. Furthermore due to their open source availability, CG simulators have also recently found wide applicability in physical object design and prototyping space cloth material identification and prototyping [Bar+16; YLL17b; Wan18], architectural design [Pan+19; Lac+19; Gav+20], fabrication of soft robots [CED19; Van+20], and the exploration of new functional and meta materials [Sch+18; Mar+19; Gus+20]. Thus, our work plugs a broader concern of

quantitatively evaluating CG simulators for increasing fidelity to realism. Figure 1.2 shows the comparison between a simulated and real experiment.

1.4 Visual Measurement of Friction in Cloth

This work, which is elaborated in Chapter 4, aims at estimating physical parameters from the rich dynamic effects that can be observed in cloth motion. Specifically, measuring a friction coefficient accurately in cloth is a longstanding problem, mainly studied in textile engineering [TV61; MK15]. There are various applications for such a visual protocol ranging 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, i.e. its stiffness and mass [Bha+03; YLL17a; Dav+17]. Friction at contact, on the other hand, is a much less studied problem due to the difficulty of modelling, measuring, and simulating dry frictional contact accurately.

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) [Don+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 [Li+18b] for capturing dry frictional contact in cloth, and apply 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, when the motion is strictly replicated in both simulation and reality and sufficient variety is added to simulation rendering conditions.

In this work we also show that calibrating the simulator properly before generating the training data is a mandatory step for achieving good estimates on real data. We show results of friction estimation for cloth-hard body and cloth-cloth scenarios and finally we show examples of applications of our measurement protocol by creating simulations of cloth materials and substrates from our real captured test data. A scene from cloth-cloth scenario is shown in Figure 1.2.



Fig. 1.2: Scene from re-simulation application, where a correct estimation of cloth-cloth friction is necessary to replicate the behaviour of selected cloth materials.

1.5 Curvature Estimation in Rods

This work, which is presented in Chapter 5, is an elaboration of an inverse design problem introduced by Bertails-Descoubes *et al.* [Ber+18]. They prove that for a suspended isotropic rod there exists a unique natural shape corresponding to a deformed equilibrium shape. The computation of this natural shape requires the deformed curvature of the suspended rod to be known. For physical rods, this computation of this curvature presents another inverse measurement problem, which we propose to resolve utilizing a visual protocol. These structures depicted by rods can be used to model and simulate various phenomena in nature including but not limited to plant growth [GT98], and human hair [Ber+06] *etc.*

We thus utilize a verified simulator SUPER-HELIX [Ber+06] to generate a dataset of suspended rods that match the physical experiment. Figure 1.3 shows a simulated rod overlaid on a physical one, both having the same material properties. After pre-processing and acquisition of a 3D centerline for each data point, we train a neural network model to depict the curvature of this 3D centerline. This network can then be utilized to make predictions on experimental data. We present preliminary results from this work which show a promising direction to build a more generic model for estimating curvature from 3D centerlines.



Fig. 1.3: Depiction of a simulated suspended rod over a physical physical one showing very good agreement. The simulation has been made using the exact known material parameters as that of the synthetically constructed physical rod.

1.6 Contributions

The joint contributions from the work presented in this document are as follows:

- To generate physically correct training data, we validate the ARGUS simulator physically and calibrate its accuracy against constrained, measurable real-world physics experiments.
- 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.
- We demonstrate the usefulness of physically validating the simulator before generating 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 of the friction coefficient using 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% of real test data capturing friction between a cloth and hard substrate and an error of < 0.1 from the baseline range on 81.1% of the test dataset for cloth on cloth experiments.
- We present preliminary results on visual measurement of deformed 3D curvature field for suspended Kirchhoff rods, where a baseline neural network model is trained on simulated data and tested on real-world experimental captures.

Related Work

Neural networks are increasingly used to solve inverse problems, where the high dimensional nature of the observed data makes it ill-posed to estimate model parameters that satisfy the observations. However, the networks are significantly dependant on the data being used to train them and acquisition of such data. In situations where acquisition of such training data is prohibitive in real world scenarios, simulated data is used and the sim-to-real generalization is achieved via some form of cross domain adaptation [Jam+19; Tob+17; Gon+19] or architecture innovations unique to the problem [FY20; FY19]. In our work since the real world test data is acquired using strictly controlled physical settings, the main generalization concern is the correct physical behaviour of the simulator. Hence in the following sections, we first look at state of the art works for simulator validation. Second, we present the related works for measurement of friction in cloth. And finally, the last section discusses the related works for curvature estimation in rods.

2.1 Validation of Computer Graphics Simulators

Since computer graphics (CG) simulators have primarily found their applications in the entertainment industry, validation of simulators has mainly been conducted qualitatively. Various well known simulators have been released by the community for animating hair, fur, virtual clothing, snow *etc.* [BW98; BWK03; McA+09; DBB11a; Kau+14; Dav20; Li+18b; Gis+20]. As long as a simulator fits a particular aesthetic narrative, it is considered useful. However, recently, with CG simulators being applied in real world scenarios such as material and structure design, robotics, bio-mechanical applications *etc.*, the need for quantitative validation has become more apparent. Indeed, for the application of physics based simulators discussed in this work, for generating realistic training data, it is necessary to validate the simulators to ensure they produce results commiserate with physical experiments. In the following subsections we detail the state of the art for simulator validation.

2.1.1 Qualitative vs. Quantitative Validation Protocols

As discussed earlier, qualitative validation has been the norm in computer graphics research. These have included tests such as cloth on rotating sphere [BFA02]; where a cloth is dropped on a rotating sphere and its wrinkling pattern is observed for visual plausibility, the cloth funnel [Har+08]; in which a cloth is passed through a funnel to determine convergence and robustness for collision detection and self contact, the reef knot [Har+09]; a knot tying simulation to determine robustness and performance on contact detection and management in a scenario with multiple simultaneous contacts, or rod plectonemes [ST07]; where a twisted filament forms a plectonemic structures for verification of rod models. These tests became standard in the community and were referred to by subsequent benchmarking studies [Ber+08; Li+18a; LKJ20]. However, there are a few studies, conducted mostly in collaboration with physicists, which perform quantitative validation. These include works on rods [Ber+08], viscous threads [Ber+10], and rigid body impacts [Smi+12].

2.1.2 Fitting Simulator Parameters to Real Data

In the fitting approach researchers have devised models whose parameters can be roughly estimated by fitting the geometric output of the simulator to the geometric output of data generated by a physical experiment. This strategy has been widely used for the fast simulation of soft tissues [Bic+09], the characterization of garments properties [WRO11; Mig+12a], or the design of inflatable structures [Sko+12; Sko+14] and materials with desired properties [Bic+10]. Once fitted to real data, the simulator output can be considered close to the real world data. While this approach is useful for specific applications, it does not truly validate a simulator and is rather limited as compared to our approach, which involves carefully evaluating the physical validity of the simulator through a testing protocol that can provide strong evidence of the accuracy of the simulator against a general physical law.

In our work we use experiments from the physics literature [SYW17; Duc06; Far17] in order to validate our data generating simulators. Each of these protocols propose a scaling law which compactly represents several regimes of a complex physical experiment, hence describing a wide parameter range. As my contribution to this collaborative work primarily involves the validation of ARGUS, we will mainly discuss the validation protocols used for the same in Chapter 3.

2.2 Estimating Friction and Material Properties of Cloth

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.

2.2.1 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 [Amo99] 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 [BFA02; Sel+09], some recent works have explored anisotropic variations of the Coulomb model when simulating interacting rigid bodies and cloth [PTS09; CFW13; Erl+20]. Interestingly, Chen et al. [CFW13] 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.

2.2.2 Friction Measurement

Estimating friction and material parameters jointly from visual data has recently become a topic of study in computer vision. Miguel *et al.* [Mig+13] propose a reparameterization of Dahl’s friction model [Dah68] 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.* [Wu+15] 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.* [ZDN16] analyse visible reflections, while Yuan *et al.* [Yua+17] 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 [Dre43; TV61; MK15; MWQ16]. Some studies in other fields connect the perception of friction (*i.e.* feeling) to actual measurements [CPD11; Liu+08] and visual features [BHT16]. 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 [Che17], 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 [SYW17]. 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 [CT64; TV61].

2.2.3 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 [SB08; YL16; Wan+15].

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 [Mig+12b; CTT17; WOR11]. To avoid capturing deformations under controlled conditions, some works have explored extracting material parameters from casually captured videos of cloth [Bha+03; KNM10]. Bouman *et al.* [Bou+13] capture stiffness and density of fabrics from video using handcrafted discriminatory features. Davis *et al.* [Dav+17] estimate material parameters based on videos that show small vibrating motions, and apply this method to fabrics. Recently, Liang *et al.* [LLK19] propose a computationally efficient way of estimating material parameters on synthetic data with a differentiable cloth simulator.

The closest methods to ours (see Chapter 4 Section 4.4), identify cloth material parameters from videos [YLL17a; Dav+17; LLK19]. Yang *et al.* [YLL17a] train a neural network on simulated cloth deformations and use the resulting architecture to infer bending and stiffness parameters. As discussed earlier, we follow the same approach of training our network with simulated data however, unlike previous methods, we carefully validate and finely tune our simulator for physical realism, through an experimental verification as discussed in the previous section.

2.3 Curvature Estimation for Suspended Kirchhof Rods

Kirchhoff rod equations [Kir] describe the theory of inextensible elastic rods (three dimensional structure where two dimensions are very small as compared to the third one). These equations have been studied extensively in literature for their wide ranging applications which include hair simulation, simulation of plant growth [GT98], DNA supercoiling [Neu04], simulation of hair [Ber+06] *etc.*

More recently Bertails-Descoubes *et al.* [Ber+18] present the formulation for solving the inverse static design of a Kirchhoff rod suspended under gravity. Their work analytically determines a unique natural configuration for a suspended rod given its deformed curvature and material frame. Our work presented in this thesis attempts to extend this inverse design protocol to experimentally captured rods, by filling in the missing piece of curvature estimation using a 3D centerline of the rod.

2.3.1 Curvature Estimation of Discrete Curves

The curvature estimation task can be summarized as follows. Given as input a curve in Euclidean space (typically of dimension 2 or 3) represented by a discrete set of points sampled along the curve, the goal is to estimate the curvature field at each point. This

problem is challenging as curvature is a derived value, and its estimation is therefore made complicated by numerical stability issues. Various algorithms have been presented in literature for curvature estimation of discrete curves. Nguyen *et al.* [ND07] present an algorithm for curvature estimation of discrete noisy curves in 2D. Similarly Fiorio *et al.* [RFM10] propose adaptive convolutional masks for computing curvature, however also remain restricted to 2D representations. Lewiner *et al.* [Lew+05] propose a method very relevant to our problem which relies on parameteric curve fitting using weighted least squares fitting and local arc length approximation.

2.3.2 Curvature Estimation Using Neural Networks

The main literature on curvature estimation using neural networks deals with estimation of principle curvatures on 2 or 3 dimensional surfaces. Littmann *et al.* [LR95] use neural networks for curvature estimation on surfaces present in grayscale images. Ben-Shabat *et al.* [BG20] use deep learning for weighted least squares surface fitting on point clouds, to extract normals and principal curvatures. And finally, [Pat+19; Qi+19; LG21] use a multilayer perceptron for estimation of curvature for the Volume of Fluid method (which is an Eulerian method used in fluid dynamics for tracking and locating the interface between two interacting fluid surfaces) by generating synthetic datasets for supervised training.

In contrast to the stated methods, we apply deep learning for estimating curvature of a spatial curve (the estimated 3D centerline of a suspended rod) and generate our training data from a verified simulator in order to achieve generalization on real rods. To the best of our knowledge, this is a novel approach to curvature estimation on 3D curves.

Physical Verification of the Simulator

3.1 Problem Definition

Several simulators of slender elastic structures, possibly under frictional contact, have been produced by the computer graphics community. These simulators were initially utilized in the entertainment industry in order to animate characters with hair, fur and cloth etc. [BW98; BWK03; McA+09; DBB11b; Dav20]. Thus the main thrust of development has been to increase visual plausibility and reduce the speed of computation.

However in our work we tackle the idea of using the simulator to generate datasets that can be utilized for training a deep learning based model capable of doing visual measurement of physical phenomena, such as static friction. Such a model can only be as realistic as the data it is trained on. Hence visual plausibility of a simulator is not enough if one is to generate a dataset that represents reality. Thus it is necessary to verify the simulator being used against a benchmark physical experiment. While graphics simulators have been previously used to learn measurement models that are applied to real data [YLL17a; YL16], to our knowledge, verification of a graphics simulator before generating training data has not been previously undertaken. Another point to note is that while the mechanical engineering community does have protocols for verification of its simulators, in the specific case of cloth simulation this community does not possess simulators able to cope with large plate deformation and dynamics, including frictional contact.

In our work we introduce rich and compact physical protocols, originally designed to measure physical parameters in soft matter physics, to quantitatively verify and calibrate the simulator. These protocols can be easily replicated in the simulator to test the physical correctness of the model, and the validity of the simulator.

One must note that the idea of verification of simulators is not just useful for generating simulated training datasets. Indeed, with the growing demand of mixing real and virtual scenes in movies, which requires an enhanced level of predictability, the entertainment industry seamlessly benefits from such a quantitative evaluation. Furthermore, computer graphics simulators are increasingly used in virtual prototyping for building real world

objects, where the demand on simulator predictability/realism is paramount. Recent developments in this field have given rise to applications such as cloth material identification and prototyping [Bar+16; YLL17a; Wan18], architectural design [Pan+19; Lac+19; Gav+20], fabrication of soft robots [CED19; Van+20], and the exploration of new functional and meta materials [Sch+18; Mar+19; Gus+20]. These applications can benefit from using physical protocols for verification of simulator accuracy, such as the ones introduced in our work.

The work presented in this chapter was first introduced in [Ras+20] and later elaborated in [Rom+21] in collaboration with colleagues from the Elan team at Inria and from the experimental physics community. My contribution to this work includes the replication of the physical protocols in the ARGUS simulator and calibration of the simulator to achieve the desired results.

3.2 The ARGUS Simulator

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. In our work we use the ARGUS implementation [Li+18b] to simulate cloth deformation subject to frictional contact. ARGUS is a combination of ARCSIM [NSO12], a cloth simulator that provides dynamic adaptive remeshing; and SO-BOGUS solver [DBB11b], a generic implicit solver for frictional contact. In SO-BOGUS the contact forces are treated as unknowns and are tied to the velocities through complementarity constraints that are resolved implicitly. This solver has been successfully used to simulate large assemblies of hair [Kau+14], cloth [Li+18a] and granular material [DB16].

A free implementation of ARGUS is available on [github](#).

3.3 Verification of the Friction Model

ARGUS simulates Amontons-Coulomb's dry frictional contact which is the behaviour we aim to visually measure in our work. As a quick summary the physical phenomena is described in the following section.

3.3.1 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 \mathbf{R} denote the reaction force on the surface in contact with a piece of cloth. We can divide \mathbf{R} into two components: the force component normal to the surface and the one tangential to the surface, denoted by \mathbf{P} and \mathbf{Q} , respectively. The force \mathbf{P} keeps the two contacting surfaces from interpenetrating, and \mathbf{Q} opposes relative displacements between the two surfaces in contact; more precisely, the two surfaces stick if $\mathbf{Q} \leq \mu\mathbf{P}$ and they slip if $\mathbf{Q} = \mu\mathbf{P}$ (the case $\mathbf{Q} > \mu\mathbf{P}$ is not admissible).

In order to verify the friction solver used in ARGUS (SO-BOGUS) we utilize the **Stick-Slip** test as described in the following section.

3.3.2 The Stick-Slip Test

This protocol for thin elastic strips was originally introduced by Sano *et al.* [SYW17] and replicated experimentally with enhanced details by my colleague Victor Romero [Rom+21]. In this latter work, our group of collaborators also recomputed accurately the limit between the different regimes of this experiment. The following analysis relies upon this recent study. In this test, a clamped elastic strip, of length L is pushed quasi-statically against a solid substrate, as shown in Figure 3.1. The strip exerts normal (\mathbf{P}) and tangential (\mathbf{Q}) forces on the substrate at the point of contact. As the vertical displacement (Δ_y) increases, so do the forces.

The strip is geometrically constrained, hence depending on the value of the friction coefficient (μ) and the vertical strain ($\epsilon_y = \Delta_y/L$), the experiment results in three well-defined regimes. These regimes are stick, slip, and extended contact (orange, turquoise, and gray regions respectively in Figure 3.2). In the stick phase, contact is localised at the lower end of the elastic strip and the contact point does not move. In the slip phase, the strip slips along the vertical plane as friction cannot hold the strip in place. Finally, in case of high enough friction coefficients, the extended contact regime is observed where a larger portion of the elastic strip comes into contact with the substrate without any slippage.

These three regions are separated by (i) a *Master curve* (computed below), (ii) the line $\epsilon_y = \epsilon_{y,c} \simeq 0.33$, and (iii) the line $\mu = \mu_c \simeq 0.36$. These three border curves meet at the triple point $(\epsilon_{y,c}, \mu_c)$.




Fig. 3.1: Stick-Slip experiment schematic for $\epsilon_y = 0.2$. The superposition of the numerical planar elastica solution (blue line), for the same value $\epsilon_y = 0.2$, illustrates the very good agreement between experiments and numerics

This phase diagram of Figure 3.2 does not depend on any mechanical or geometrical property of the system. Therefore a friction solver can be evaluated regardless of the material model being used and the results would be independent of the accuracy of the material model.

Master Curve: Threshold Between Stick and Slip Regimes

To compute the **Stick-Slip** master curve, which partitions the stick and slip regions, we compute the equilibrium of a clamped planar elastica as the parameter ϵ_y is varied. The equilibrium solution gives the forces $\mathbf{P} = -P \mathbf{e}_y$ and $\mathbf{Q} = Q \mathbf{e}_x$ and friction is accounted for with the Amontons-Coulomb law. More precisely, once contact is made, the strip buckles as soon as the vertical force P exceeds $20.2 EI/L^2$. At buckling the horizontal force is $Q = 0$ and friction is not required for the equilibrium to hold. As ϵ_y increases the ratio Q/P keeps on increasing, and slip occurs as soon as Q/P reaches the Coulomb friction coefficient μ .

The threshold curve is then reached when the clamped-pinned equilibrium is such that $Q/P = \mu$. We therefore compute the ratio Q/P at equilibrium, as function of ϵ_y and plot it in the phase diagram, see the black curve in Figure 3.2.

Using the dimensionless arc length $\bar{s} = s/L$, equilibrium is found as the solution of the following boundary value problem:

$$\frac{d^2\theta}{d\bar{s}^2} = \frac{PL^2}{EI} \sin \theta + \frac{QL^2}{EI} \cos \theta \quad \text{with } \theta(0) = \pi/2, \quad \frac{d\theta}{d\bar{s}}(1) = 0 \quad (3.1a)$$

$$\frac{d\bar{x}}{d\bar{s}} = \sin \theta \quad \text{with } \bar{x}(0) = 0, \quad \bar{x}(1) = 0 \quad (3.1b)$$

$$\frac{d\bar{y}}{d\bar{s}} = \cos \theta \quad \text{with } \bar{y}(0) = 0, \quad \bar{y}(1) = 1 - \epsilon_y^* \quad (3.1c)$$

where we note ϵ_y^* the value of ϵ_y on the transition curve.

Note that in Eq. (3.1a), \mathbf{P} and \mathbf{Q} only appear divided by the scaling force EI/L^2 and are therefore going to be proportional to it. Hence, the ratio $\mu = Q/P$ will be independent of the mechanical (EI) and geometrical (L) properties of the elastic strip.

One can analytically solve Eq. (3.1) up to a second order in θ by writing $\sin \theta \simeq \theta$ and $\cos \theta \simeq 1 - \theta^2$. With this weakly nonlinear approximation, the solution for the master curve reads

$$\frac{Q}{P} \simeq 0.445 \sqrt{\epsilon_y} \quad (3.2)$$

This solution (dashed white line in Figure 3.2) is not accurate enough to serve in the present test. A detailed, semi-analytical, solution of this problem can be found, see *e.g.* [Mik06], but involves elliptic functions and a root-finding step. Here, without any loss of accuracy, we

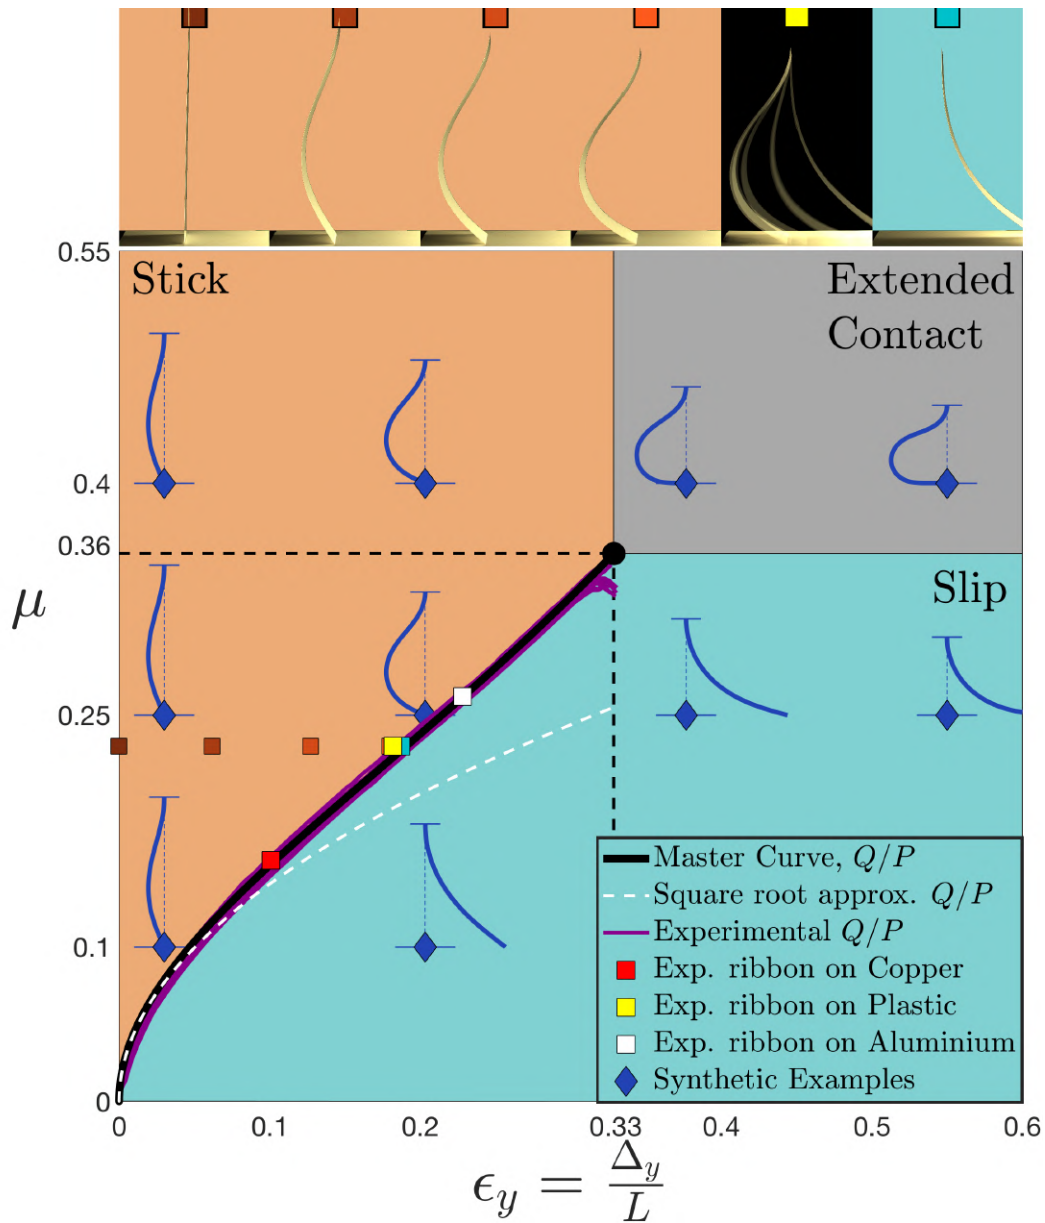


Fig. 3.2: Master curve and experimental validation for the **Stick-Slip** test. Phase diagram where the master curve (in black) separates the sticking regime (in orange) from the slipping regime (in turquoise). The extended contact region (in gray) is not used in our validation protocol. The (numerical) black and (experimental) purple curves are plots of the ratio Q/P as a function of ϵ_y .

simply solve system (3.1) numerically and plot Q/P as a function of ϵ_y , see the black curve in Figure 3.2. The triple point $(\epsilon_{y,c}, \mu_c) \simeq (0.33, 0.36)$ is reached when the equilibrium solution is such that $\theta(s = L) = \pi$.

3.3.3 ARGUS Verification Results

I now describe my contribution in this work, consisting in validating the ARGUS simulator against the **Stick-Slip** protocol. The **Stick-Slip** test is a 2D test which is valid when the effect of gravity can be considered negligible. The remaining two degrees of freedom in the test thus include the friction coefficient μ and the bending force EI/L^2 . The test should be insensitive to the latter parameter, so it can be taken equal to 1. We succeeded in matching the threshold defined in the **Stick-Slip** test with the output of ARGUS, but this requires 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. As shown in Figure 3.3, although it generates fair results, the original adaptive version of ARGUS does not properly match the master curve, presumably due to small impacts during adaptive refinement.

After disabling adaptivity in ARGUS code, we want to mimic a quasi-static experiment by increasing damping forces in order to remove spurious dynamical vibrations of the cloth. And finally, we intend to find the right level of accuracy. We do this by sufficiently decreasing the timestep of the simulation as well as the solver tolerance, to a point where further refinement would lead to indistinguishable results at our observation scale. We set the solver tolerance to 10^{-13} and a maximum number of iterations of 50000. In our comparisons, we use a timestep of 0.5 ms, a mesh resolution of 149 vertices for a 10 cm long ribbon and a damping value set to 0.002. We repeat the simulation using 2 materials having different Young's moduli. This is done in order to ensure that the numerical experiment is insensitive to the material properties of the ribbon as long as gravity plays a negligible role.




Fig. 3.3: ARGUS Results on **Stick-Slip** test with Adaptive Remeshing




Fig. 3.4: ARGUS Results on **Stick-Slip** test without Adaptive Remeshing

These results, as can be seen in Figure 3.4, 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, as explained in Chapter 4, 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 prediction results presented in Chapter 4 Sec. 4.5.2 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 the solver, hence an unnecessarily large training cost. By generating training datasets with varying accuracy, we check that our high accuracy calibration is actually mandatory to obtain satisfying prediction results from real data. This comparison is presented in Chapter 4 Sec. 4.5.2.

3.4 Verification of the Bending Model

As explained earlier, ARGUS is a combination of the adaptive cloth simulator ARCSIM and the SO-BOGUS friction contact solver. We have validated the frictional contact part (SO-BOGUS) on the Stick-Slip protocol. Our next goal is to validate the bending elasticity part (ARCSIM). To this aim, we leverage the Cantilever protocol, which is detailed below. For our tests we used version 0.2.1 of ARCSIM, as recommended by its authors over version 0.3.1 which is useful for shells and fracture and not needed in our case.

3.4.1 The Cantilever test

The **Cantilever** test is a 2D test commonly used by the soft-matter community for inferring mechanical parameters of rods and ribbons [Duc06; Far17]. The main advantage of this test is that, from purely geometrical considerations and measurements, it is possible to estimate, with good accuracy, the elastic behaviour of the material. The soft cantilever consists of a slender, naturally straight rod or ribbon, clamped horizontally, and deformed by the action of gravity, see Figure 3.5, left. The parameters involved in the problem are the bending rigidity (EI), the material density (ρ), the rod's length (L), the area A of the cross-section of the rod, and the acceleration due to gravity (g). Here E is the Young elastic modulus and I the second moment of area of the cross-section of the rod. One can construct a unique length

scale for this problem, namely the gravito-bending length $L_{gb} = \sqrt[3]{EI/(\rho Ag)}$, which compares the resistance to bending with the gravitational force. Thus, for an elastic object deformed by gravity with length L , its equilibrium shape is effectively determined solely by the dimensionless gravito-bending parameter $\Gamma = (L/L_{gb})^3$. Note that in the case of ribbons, as is the case with ARCSIM, when the cross-section is a rectangle with large aspect ratio, the Young modulus E is replaced with $E^* = E/(1 - \nu^2)$ and Γ by $\Gamma^* = (1 - \nu^2)\Gamma$, where ν is the Poisson ratio of the material [Shi92].

The parameter Γ gives an immediate notion of the expected behaviour for the deformation of an elastic object. Large values of Γ lead to large deformations, and conversely low values yield small deformations. Thus by utilizing this test one can quantitatively compare simulations with experimental and analytical results for a broad range of mechanical and geometrical parameters. For such comparison, here we chose the aspect ratio of the final shape, H/W (see Figure 3.5), which is an accurate indicator of the quality of the simulation and the mapping between Γ and H/W is unique (strictly monotonic).

Master Curve: Solution of the Planar Elastica

The final shape for rods and ribbons is governed by the planar elastica equations (*i.e.* the Kirchhoff equations for 2D rods, see *e.g.* [LL59]). The equilibrium for the internal force \mathbf{F} and moment \mathbf{M} for an infinitesimal variation of the arc length s read $d\mathbf{F}/ds + \mathbf{f}_{\text{ext}} = \mathbf{0}$ and $d\mathbf{M}/ds + \mathbf{t} \times \mathbf{F} = \mathbf{0}$, with \mathbf{t} the unit tangent to the rod. Free-end boundary conditions reads $\mathbf{F}(s = L) = \mathbf{0} = \mathbf{M}(s = L)$. The constitutive relation between the bending moment and the curvature $d\theta/ds$ reads $M = EI d\theta/ds$, with θ the angle between the horizontal and the tangent, see Figure 3.5, left. When gravity is the only external force, $\mathbf{f}_{\text{ext}} = \rho A g \mathbf{e}_y$, yielding $\mathbf{F} = \rho A g (L - s) \mathbf{e}_y$ and $d\mathbf{M}/ds + \rho A g (L - s) \cos \theta \mathbf{e}_z = \mathbf{0}$. Introducing dimensionless variables $\bar{s} = s/L$, $\bar{x} = x/L$, and $\bar{y} = y/L$, the boundary value problem to be solved finally reads

$$\frac{d^2\theta}{d\bar{s}^2} + \Gamma(1 - \bar{s}) \cos \theta = 0 \text{ with } \theta(0) = 0 \text{ and } \frac{d\theta}{d\bar{s}}(1) = 0 \quad (3.3a)$$

$$\frac{d\bar{x}}{d\bar{s}} = \cos \theta \text{ with } \bar{x}(0) = 0 \quad (3.3b)$$

$$\frac{d\bar{y}}{d\bar{s}} = \sin \theta \text{ with } \bar{y}(0) = 0 \quad (3.3c)$$

where the aforementioned parameter Γ arises naturally. For Γ values smaller than ~ 200 , this boundary value problem is easily solved with simple shooting techniques [AMR95]. However as Γ gets larger the problem becomes stiff, and more powerful continuation methods are required. We use the collocation-based continuation package AUTO07p [DKK91]

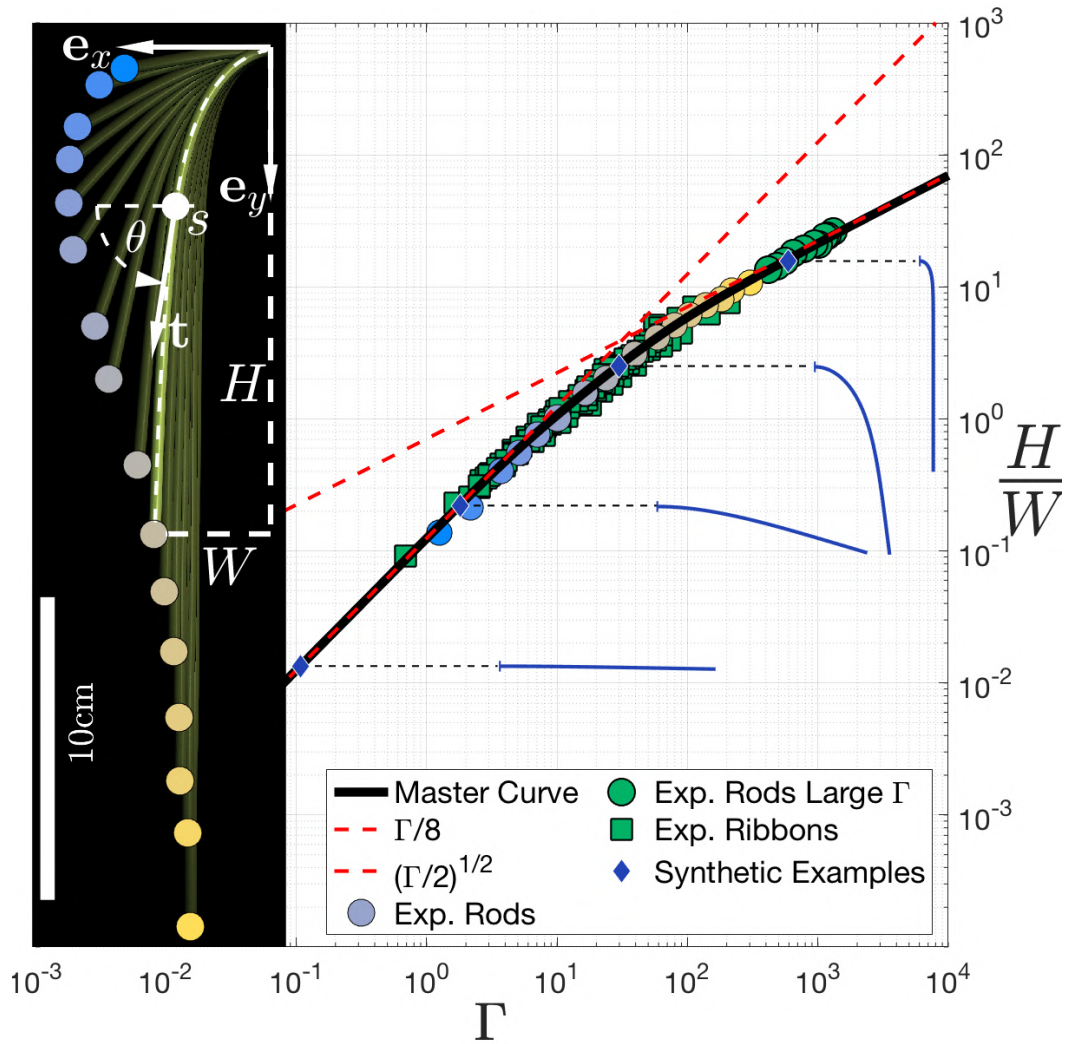


Fig. 3.5: Master curve and experimental validation for the **Cantilever test**. Left: snapshots of experimental rods under gravity, with 16 different Γ values. Right: our computed master curve (in black) representing the aspect ratio of the rod's shape as a function of Γ , together with experimental results from [Rom+21]

to compute the solution of (3.3) over 8 decades of Γ values, and we plot the aspect ratio $H/W = y(L)/x(L)$ as a black solid line in Figure 3.5, right. Furthermore, one can find analytical expressions for the aspect ratio for small and large Γ values, these are $H/W = \Gamma/8$ for $\Gamma \ll 1$ [Ger04], and $H/W = \sqrt{\Gamma/2}$ for $\Gamma \gg 1$, plotted as red dashed lines in Figure 3.5.

3.4.2 ARCSIM Verification Results

Like for the Stick-Slip test, my simulation results are compared against the recent experimental data from Victor Romero, and the master curve accurately derived in our collective work [Rom+21]. We use two different mesh resolutions, Res0 (120 nodes) and Res+ (240 nodes) for evaluating ARCSIM. Furthermore we test ARCSIM with adaptive remeshing turned off. While the fixed remeshing and increased resolution produce slightly better results, in all cases ARCSIM fails to pass the **Cantilever** test, see Figures 3.6a, 3.6b, 3.6c, 3.6d. One can observe a shift from the master curve as well as scattering of results, which intensifies for higher values of Γ .




Fig. 3.6: ARCSIM results on **Cantilever** experiment. Figure taken from our collective work [Rom+21]


Modifying the Bending Energy

In this section we discuss the impact of changing the bending energy model in ARCSIM, work done in collaboration with my colleague Mickaël Ly. ARCSIM uses a popular formulation of bending energy for plate simulators proposed by [BMF03]. This formulation was simultaneously proposed along with [Gri+03]. Although the latter derives it from a discretization of the square of the difference of the mean curvatures and the former from heuristics, both yield a similar (up to a scaling coefficient) hinge energy based on the dihedral angle between faces at their shared edges. The two formulations only differ from each other by a multiplicative factor and by introducing a slight modification of the bending coefficient (from $\frac{1}{4}$ to 3), one can readily switch to [Gri+03]’s bending model in ARCSIM.


However, after making this modification results lie slightly below the curve and still remain excessively scattered, as illustrated in Figures 3.7a, 3.7b. Testing the code in its non-adaptive version and using an increased mesh resolution unfortunately does not help, as shown in Figures 3.7c, 3.7d.

Implications for Visual Measurement Results


To obtain physically accurate material parameter settings in the simulator, we leverage the work of Wang *et al.* [WOR11] 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: 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.* [YLL17a] for recovering cloth material parameters. Despite ARCSIM failing to clear the **Cantilever** test, data generated from the model, gives passable results on material parameter estimation (as will be shown in Chapter 4). The reason for this is that unlike the prediction of friction coefficient, which is a physical quantity, in case of material estimation the model does not predict physical material parameters such as (Young’s Moduli, volumetric density etc.), but discrete material classes as defined by Wang *et al.* [WOR11]. These material classes have been derived by fitting real world cloth materials to simulator data and thus inferring the appropriate simulator parameters that produce similar visual appearance.




(a) ARCSIM Adaptive Remeshing Res0 with Grinspun *et al.* [Gri+03] Bending Energy



(b) ARCSIM Adaptive Remeshing Res+ with Grinspun *et al.* [Gri+03] Bending Energy



(c) ARCSIM Fixed Resolution Res0 with Grinspun *et al.* [Gri+03] Bending Energy



(d) ARCSIM Fixed Resolution Res+ with Grinspun *et al.* [Gri+03] Bending Energy

Fig. 3.7: ARCSIM results on **Cantilever** experiment with Grinspun *et al.* [Gri+03] Bending Energy. Figure taken from our collective work [Rom+21]

Friction Estimation in Cloth

4.1 Problem Definition

In the spirit of inferring physical properties of objects from visual data, we present the problem of estimating physical parameters from the rich dynamic effects that can be observed in cloth motion. As described in Chapter 2, most of the existing methods seeking to solve this problem have focused on estimating cloth material parameters, that is, its stiffness and mass [Bha+03; YLL17a; Dav+17]. Friction at contact is a much less studied problem due to the difficulty of modelling, measuring, and simulating dry frictional contact accurately. Such friction can manifest between a piece of cloth and a rigid substrate or between two pieces of cloth. However, friction has a high impact on the overall cloth dynamics, as illustrated in Figure 4.1 where cloth samples contacting a horizontal plane were simulated with the same material parameters, but with different friction coefficients at contact.

An accurate measurement of friction is usually achieved at the expense of tedious and generally invasive protocols with specialised mechanical sensors and hardware. However, the visual variability induced by friction shown in Figure 4.1 shows that a purely vision-based protocol is plausible for discriminating the different motions and folding patterns observed under such 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 [Che17] or stiff ribbons [SYW17]. 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 (Figure 4.1, centre and bottom rows).

The main impediments in achieving a complete vision based protocol for friction measurement are two fold. First that a data driven, vision based approach requires a large representative dataset. This requirement presents a challenge as acquiring such a dataset from real life would require acquisition of cloth materials with exact and varying friction properties that can represent the whole range of observable cloth. Thus making real life




Fig. 4.1: 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 $\mu = 0.6$.

acquisition prohibitive. In scenarios where capturing significant amounts of real data is intractable, many works have explored the idea of training on simulated data with the aim to generalise inference on real data. These include approaches for cross domain transfer learning [Jam+19; Tob+17] and interpretable low dimensional representation learning [Gon+19]. In our work we create such a dataset using the ARGUS simulator (expanded on in Section 4.3). This is where the second challenge arises; the simulated dataset needs to accurately reflect the real world by capturing the desired physical phenomena precisely. Our approach, as detailed in Chapter 3, can be considered orthogonal to these aforementioned efforts as we present a protocol to calibrate simulated data generation by experimental verification, before learning a model from this data. We perform calibration of our used simulator ARGUS against a physical experiment proposed by Sano *et al.* [SYW17] before we utilize it for dataset generation.

Thus in this chapter, 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. This method relies on the generation of accurate synthetic data which can be used to learn a model that can generalize its inference to data captured in the real world. We discuss the acquisition of experimental test data, synthetic dataset generation from a verified simulator, and a deep learning based model to learn material parameter classifier and friction estimator from video sequences. This work has been published [Ras+20] in collaboration with colleagues from the experimental physics community. My main contributions are the conceptualization of the cloth motion sequence, synthetic data set generation on a verified simulator and design, training and testing of the deep learning model.

4.2 Experimental Test Data

Our first contribution comprises of generating a dataset of closely matching captured videos and simulations with, for each video, corresponding material classes and friction coefficients. For the real-world data capture 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 4.2.3, 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 Figure 4.8. In addition to material substrates we also generate data for cloth on cloth friction, in which a cloth material acts as a substrate. Estimating cloth on cloth friction has been a relatively more difficult task due to the woven nature of the materials. 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.

4.2.1 Cloth materials

We use 9 materials in bright colours, out of which 8 are close in composition and density to material classes defined in [WOR11], and one (silk) is not covered by these material classes. Details on the material features can be found in Table 4.1. The equivalent material classes from [WOR11] are also shown in the table. From each material we laser cut 20x20 cm pieces with 1x1cm holders at two corners of one edge for clamping purposes.

4.2.2 Substrates Selection

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 choose 7 rigid 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. Experiments with such substrates are shown in Figure 4.2.




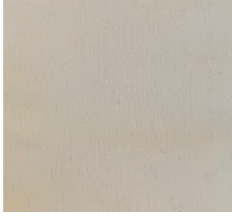


Fig. 4.2: Experimental images with various hard substrates for the same cloth material. From left to right: Aluminium, Aluminium-PET, Ceramic, Rough Glass, Smooth Glass, Polyester, and Stainless.

We also present our work in more complex regimes of cloth-on-cloth friction, that are not covered by Amontons-Coulomb's law, as shown in Figure .4.3. 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 reference measurement using the inclined plane setup (explained in Subsection 4.2.5) is performed immediately before running the drop and drag experiment to capture the same state of the textile substrate. Note

		
M01	M02	M03
Polyester 50% Rayon 50%	Polyester 95% Spandex 5%	Cotton 100%
0.187(Kg/m ³)	0.187(Kg/m ³)	0.75(Kg/m ³)
Camel Ponte Roma in [WOR11]	Navy Sparkle Sweat in [WOR11]	11oz Black Denim in [WOR11]
		
M04	M05	M06
Cotton 65% Polyester 33% Spandex 2%	Polyester 100%	Polyester 100%
0.207(Kg/m ³)	0.198(Kg/m ³)	0.083(Kg/m ³)
Royal Target in [WOR11]	Pink Ribbon Brown in [WOR11]	Tango Red Jet Set in [WOR11]
		
M07	M08	M09
Polyester 100%	Silk 100%	Cotton 95% Spandex 5%
0.150(Kg/m ³)	0.065(Kg/m ³)	0.268(Kg/m ³)
White Dots on Black in [WOR11]	Not included in [WOR11]	Ivory Rib Knit in [WOR11]

Tab. 4.1: Experimental Cloth Materials

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 in the following sections.




Fig. 4.3: Experimental images with different textile substrates for the same cloth material

4.2.3 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 movement is followed by a led lamp to ensure consistent illumination. 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 mm/s^2 until it reaches 10 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 mm/s^2 until it reaches 10 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. The motion is captured from a well characterised view point. Two annotated views of the experimental setup for producing our real test data-set are presented in Figure 4.4.




Fig. 4.4: Experimental setup for real data-set creation

4.2.4 Video acquisition

With a calibrated camera, we record the cloth motion from the viewpoint shown in Figure 4.8, 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.

4.2.5 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 [CT64; TV61]. 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 friction coefficient is then computed based on the slope of the plane at the point where the object slips.




Fig. 4.5: Sketch of the inclined plane setup for frictional characterisation

As depicted in Figure 4.5, an object (with mass m , gravitational acceleration g) on an inclined plane, with angle of inclination α , is in equilibrium as long as the tangential part of its weight, $mg \sin(\alpha)$, is equal to the frictional force \vec{Q} and the normal part, $mg \cos(\alpha)$, is equal to the normal \vec{P} . As the angle α increases so does $mg \sin(\alpha)$. The mass m remains in equilibrium up to the limit where the frictional force reaches its maximum possible value $\vec{Q} = \mu \vec{P}$. This gives the relation for the friction coefficient $\mu = Q/P = \sin(\alpha)/\cos(\alpha) = \tan(\alpha)$.

For each substrate we test the angle at which a piece of cloth slides on it. From the angle we compute its tangent and we use this value as a friction coefficient. We test each piece of cloth on various parts of a substrate and do not find a noticeable difference in the slippage angle.

Since our measurement is the angle of maximum inclination before sliding, and we move the inclined plane with a discrete interval of inclination angle $\delta\alpha$, the propagation of the error to the measured parameter $\mu = \tan(\alpha)$ is given by $\delta\mu \sim \pm \sec(\alpha)^2 \delta\alpha$.

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. Figure 4.6 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.




Fig. 4.6: Distribution of baseline friction values (μ) for all material-substrate pairs.

Peculiarities for Cloth-Cloth Test Data

We record a baseline measurement of friction of each cloth-cloth pair based on the same inclined plane method, while stretching the cloth at the bottom to serve as a substrate. As in this case, 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 Figure 4.7

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.




Fig. 4.7: Distribution of baseline friction values (μ) for all cloth-cloth pairs.

4.3 Dataset Generation from Verified Simulator

For training and testing, we simulate a physically accurate dataset that closely resembles the videos captured using the experimental protocol. We use the ARGUS implementation [Li+18b] 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.

The simulator is rigorously tested and calibrated against a physical experiment as discussed in Chapter 3 Section 3.3).

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 experiments described in Chapter 3, 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	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

Tab. 4.2: Parameter specifications of 3 datasets, generated for varying levels of accuracy.

dataset are described in Table 4.2. Out of the 3 datasets, High-accurate-solver-small-timestep dataset corresponds to calibrated settings obtained from the validation experiment.




Fig. 4.8: 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.

In each dataset the drop and drag motion is simulated for the 10 material classes measured by Wang *et al.* [WOR11]. For each material, 16 friction coefficients evenly distributed between

0.0 and 1.5 are explored, which represents a reasonable range for fabrics according to values tabulated in [Dre43] that 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 4.5.1). 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 Figure 4.8.

4.4 Model Description

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 [Don+15]. We are inspired by a previous model that recovers a material class label of cloth from an input image sequence [YLL17a]. A variant of this model, called *baseline model* in the following, is explained in Sec. 4.4.1.

While the baseline model is effective at predicting a material class, it performs significantly worse when trained to predict friction coefficients (see Section 4.5.1 for details). We therefore propose a novel pipeline for this task that conditions friction on material classes, as outlined in Sec. 4.4.2.

4.4.1 Baseline Model

In the baseline model by Yang *et al.* [YLL17a], 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 [KSH12]. We replace this architecture by a simplified version of VGG16 [SZ14] with ReLU activations as shown in Table 5.1. The weights are shared between all feature extraction blocks. Let

$$f_i = CNN_{VGG}(I_i) \quad (4.1)$$

denote the image features extracted from frame I_i , where CNN_{VGG} is the simplified VGG16 of Table 5.1.




Fig. 4.9: 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.

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

Tab. 4.3: Architecture details for a feature extractor block.

A sequence of learned image features f_1, f_2, \dots, 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

$$\hat{h} = FC(LSTM(f_1, f_2, \dots, f_k)), \quad (4.2)$$

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 Figure 4.9. This architecture is trained with a standard categorical cross-entropy classification loss, and for prediction, the class label with the highest likelihood is reported.

4.4.2 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 Figure 4.9. In addition to a sequence of 2D video frames $I_i, i = 1, \dots, k$, 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 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

$$\hat{y} = FC(LSTM(FC_s(m) \frown f_1, \dots, FC_s(m) \frown f_k)), \quad (4.3)$$

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

$$\mathcal{L} = -\log P(y|I_1, I_2, \dots, I_k, m), \quad (4.4)$$

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.

Implementation Details

This algorithm is implemented uses Keras [Cho+15] and Tensorflow [M A+15]. While we experimented with various optimisers, we empirically observed Adadelta [Zei12] 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.

4.5 Results Analysis

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.

4.5.1 Results on Simulated Data

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

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% of our dataset for testing.

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 %

Tab. 4.4: Results on Simulated Test Data

The 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.

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%, and an error of < 0.2 on 95% of the samples. This shows that the model's predictions do not degrade dramatically if an unseen texture is encountered, implying that the model captures more nuanced spatio-temporal phenomena to classify friction behaviour.

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 4.5. 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.

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%

Tab. 4.5: Results on unseen viewpoints

Generalisation to Unseen Rendering Conditions

We test 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 the following Section 4.5.2. We test our model on 160 sequences, and our model achieves an error of < 0.1 on 72.7%, and an error of < 0.2 on 84.1% 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.

4.5.2 Results on Real Test Data

We present results on real data captured through our experimental setup presented in Section 4.2. The following subsections present the evaluation protocol and the results acquired on the cloth-substrate and the cloth-cloth test data sets, along with their analysis.

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 consider the top-1 and top-2 accuracy.

Results on Cloth-Substrate Test Data

Our dataset 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.

Impact of the Simulator’s Accuracy

As can be observed in Figure 4.10, 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% of test data compared

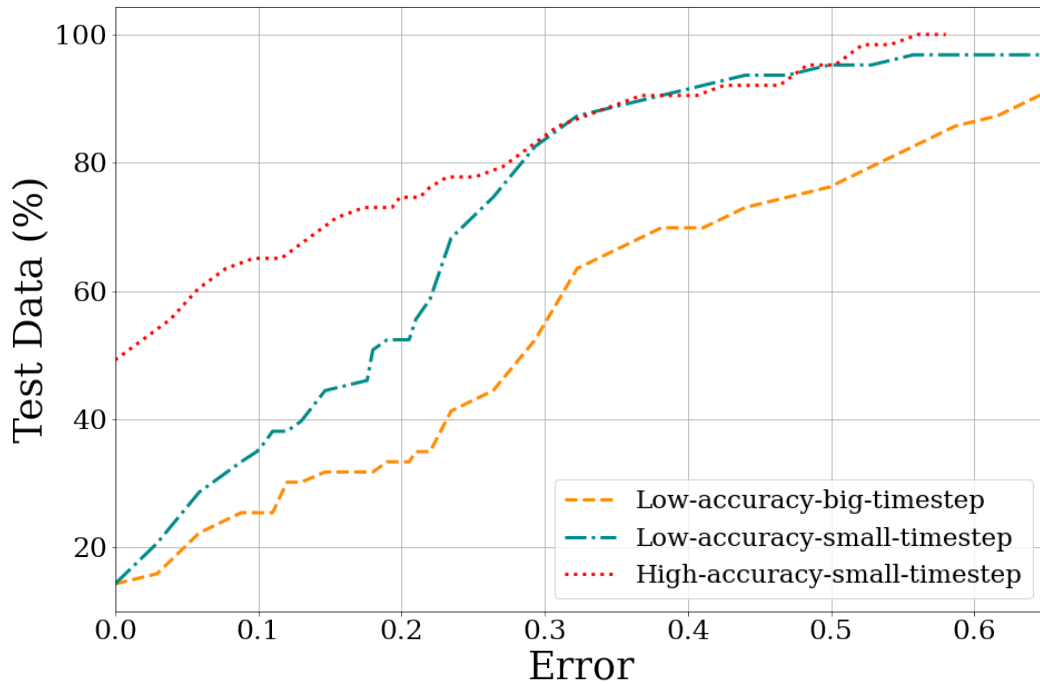


Fig. 4.10: Cumulative error plots for all datasets

to 25.3% and 34.23%, respectively). Recall that the former dataset is generated using our calibrated parameters settings based on the validation experiment described in Chapter 3 Section 3.3.

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% of test data as compared to 33.3% of the model trained on the former.

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

Figure 4.11 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% of data) is significantly lower than the model accuracy obtained on synthetic test data (dark




Fig. 4.11: Cumulative error plots for dataset High-accurate-small-timestep (red, black, light blue, green) and simulated test data (dark blue).

blue curve). 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% 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% 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).

Results on Cloth-Cloth Test Data

In order to extend our non-invasive parameter estimation model to more complex regimes we collect a cloth on cloth test dataset as explained in Section 4.2. As with cloth-substrate test data, our cloth-cloth dataset contains 5 videos for each cloth eligible 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 cloth pair.

Outliers with respect to Training Dataset

Out of the measured baseline values, certain cloth pairs display a friction measurement higher than 1.5. Such values, shown as orange bars in Figure 4.12, clearly lie outside the purview of our synthetic training dataset. It would be thus reasonable to expect our prediction method to fail on such data.




Fig. 4.12: Distribution of baseline friction values for cloth to cloth dataset. The orange bars are outside the range of our synthetic training data.

Evaluation With and Without Fine Tuning

In order to estimate friction in the cloth on cloth data, we use a model pre-trained on the synthetic dataset High-accuracy-small-timestep. To further improve the generalization of our model, we further fine tune it by training it on additional synthetic renderings. In order to generate these renderings 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 Figure 4.13




Fig. 4.13: Extension of our synthetic dataset, with renderings more similar to the new cloth on cloth dataset




Fig. 4.14: Prediction errors on Cloth to Cloth friction test dataset

We present prediction error plot from both before and after fine-tuning our model in Figure 4.14

In Figure 4.14, 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% of the data and with fine-tuning the model improves to an error of < 0.2 on 87.8% of the data and an error of < 0.1 on 72.7% of test data. This is comparable with the results achieved for our previous experimental test dataset discussed in Sec. 4.5.2. 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% and < 0.1 on 53.4% 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.




Fig. 4.15: Corresponding frames showing the cloth coming into contact with a rotating sphere substrate. Top Row: M02 (Polyester 95 %, Spandex 5 %) with $\mu = 0.3$, substrate: aluminium, Bottom Row M03 (Cotton 100 %) with $\mu = 0.6$, substrate: ceramic

4.6 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 Figure 4.15 and 4.16.

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. As shown in Figure 4.15, the wrinkle patterns in both experiments differ significantly given the choice of material and friction coefficients.

Secondly, 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 Figure 4.16, 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. The friction coefficients used are the ones predicted by our measurement protocol on the test data, *i.e.* 0.3 and 0.7.




Fig. 4.16: Corresponding frames of a simulated character displaying cloth on cloth friction. Top Row: M02 (Polyester 95 %, Spandex 5 %) with $\mu = 0.3$, Bottom Row M06 (100 % Polyester) with $\mu = 0.7$

4.7 Limitations and Future Work

Despite the physical plausibility of the simulated training data, and the rather simple setup for the physical video acquisition setup, the presented protocol is restricted to the alignment between the physical experiment and its replication in the simulated data. However the method opens up an interesting direction for estimation of friction in-the-wild. The restrictions of a strictly bounded physical setup for video acquisition can be progressively relaxed in future work, leading to a more generally usable solution.

Curvature Estimation in Rods

5.1 Problem Definition

In the previous Chapter we dealt with the so called *inverse measurement problem* where the parameters of a physical model (friction coefficient and material class) are estimated treating data corresponding to the model. In this Chapter we introduce our work on an *inverse design problem* where material parameters are known and one is looking for the intrinsic shape of the object. In particular we are concerned with the problem of finding the natural shape of a rod (also referred as fiber or filament) while the equilibrium state of the fiber is known under the deformation due to gravity. A solution to this problem has been presented by Bertails-Descoubes *et al.* [Ber+18]. In this paper, they prove that for a suspended isotropic rod with known material parameters (mass, stiffness), there exists a unique natural shape corresponding to a deformed equilibrium shape. And that this natural shape can be computed by solving three linear initial value problems (IVP) in series.

However the IVP assumes that the deformed curvature of the rod¹ in its equilibrium state (κ) is known. This assumption turns out to be prohibitive in the case of real world rods. Thus in order to apply the method to a physical object, one requires a visual measurement technique that can estimate the deformed curvature of a suspended rod. Failing to recover the curvature field with sufficient accuracy using standard techniques (finite differences, spline fitting, etc.) [Rom+19], we here again resorted to deep learning models to visually estimate a curvature field from a suspended rod. As creating a dataset from physical rods (even if synthetically constructed), that faithfully covers a representative manifold of deformed curvatures is intractable, we rely on generating such a dataset from a verified simulator. In this work we aim to train a deep learning based curvature estimator on synthetic data that can be generalized for curvature estimation on both manufactured and naturally existing physical rods. This curvature estimation can then be used to solve the inverse problem as described in [Ber+18].

This work is currently ongoing in collaboration with Victor Romero, Arnaud Lazarus and my thesis supervisors. My contribution to it involves design and exploitation of a deep learning model to recover the deformed curvature of suspended rods, while ensuring that the model trained on synthetic data can generalize well to real world fabricated and naturally occurring

¹In reality a material rod possesses two curvatures and a twist. Here curvature is being used as a simplification.

rods. The experimental data capture and synthetic dataset generation were performed by Victor.

5.2 Background and Modalities

We follow the convention described in Bertails-Descoubes *et al.* [Ber+18] and consider an inextensible and unshearable rod of length L , represented by a centerline $\Gamma(s)$ and a material frame $\mathcal{R}(s)$, which are both parameterized by arc length $s[0, L]$. Thus $\Gamma(s)$ is the 3D position of the centerline at location s . Furthermore the rod is clamped at $s = 0$, with clamped position $\Gamma(0) = \Gamma_0$ and orientation $\mathcal{R}(0) = \mathcal{R}_0$, while the end of the rod at $s = L$ is free.

In this paper the authors are interested in the inverse design for Kirchhoff rods, which can be stated more precisely as follows. If the equilibrium geometry Γ is known for the deformed configuration of the rod, one needs to compute the natural strain vector field (or natural curvature) $\bar{\kappa}$.

The authors show the following property. For any framed curve $\{\Gamma_0; \mathcal{R}_0; \kappa(s)\}$ Kirchhoff's equilibrium equations are satisfied if and only if $\bar{\kappa}(s)$ is solution to the explicit linear ODE of first order,

$$\mathbb{K}_3(\kappa - \bar{\kappa})' + [\kappa]_{\times} \mathbb{K}_3(\kappa - \bar{\kappa}) = [\mathbf{e}_x]_{\times} \mathcal{R}^T \mathbf{T}(s) \quad (5.1)$$

where $\mathbb{K}_3 = \text{diag}_3(\mathcal{K}_0, \mathcal{K}_1, \mathcal{K}_2)$ is a diagonal 3×3 matrix collecting the twisting and bending stiffness and $\mathbf{T}(s)$ represents the internal forces acting on the rod; in the case when only gravity is applied, this term boils down to a function of the rod weight distribution, which can be computed easily.

For our work we assume that for a given experimental rod, \mathbb{K}_3 , $\mathbf{T}(s)$ and the material frame \mathcal{R} are known. Thus from Eq. 5.1, one can compute the natural curvature $\bar{\kappa}(s)$ given one can estimate the deformed curvature κ .

However estimating the deformed curvature κ is not trivial. In practice, slight errors in estimation can lead to a drastically different natural curvature. Furthermore, reconstruction of centerline of a deformed experimental rod, using image analysis techniques, already introduces a certain reconstruction noise, which exacerbates the problem.

For these reasons we use machine learning to formulate a visual solution to estimating the deformed curvature of a noisy centerline obtained from a suspended Kirchhoff rod, which




Fig. 5.1: Experimental setup of Victor Romero, showing a manufactured suspended rod being captured.

takes into account the noise from the estimation process while yielding the correct curvature estimation.

5.3 Experimental Data Capture

In order to capture experimental test data, rods of Vinyl Polysiloxane were fabricated such that the lengths and the material properties of the rods are completely known. The material used has an approximate Young's modulus on the order of 1 MPa. The rod is suspended on a stage with dark background and images are taken using 4 viewpoints for 2D to 3D projection. The setup is shown in Figure 5.1. The cameras are strategically placed, such that occlusions are minimized and there are no aberrations in the 3D reconstruction. The 4 viewpoints are illustrated in Figure 5.2.

5.3.1 Images to 3D Centerline Pipeline

In order to capture the 3D centerline from the experimental images, edge detection and region of interest definition were performed to separate foreground from the background.

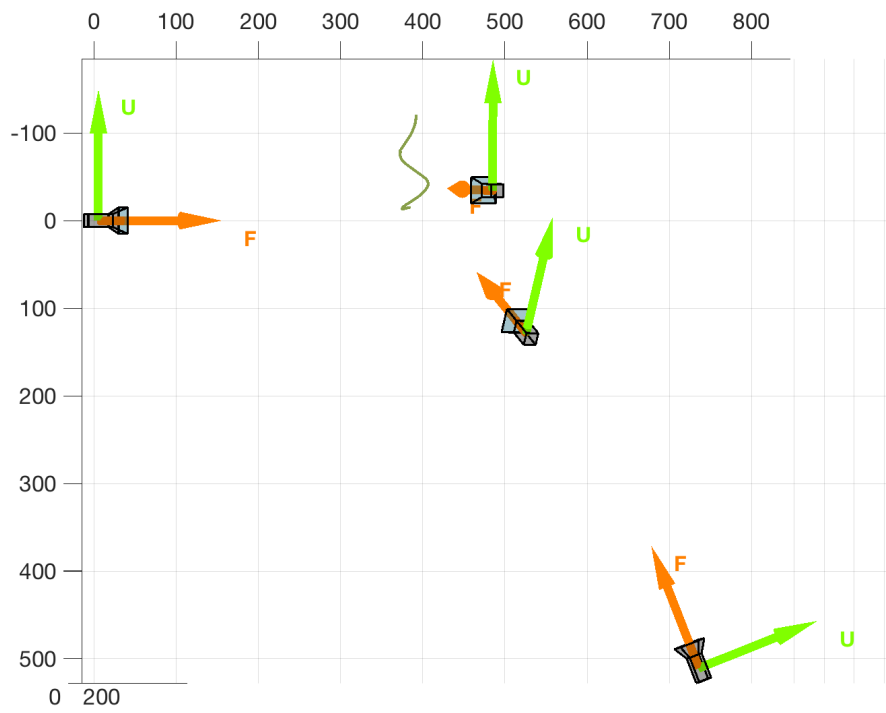


Fig. 5.2: Camera positions shown for capture of the experimental rod. Image courtesy of Victor Romero.

The generated silhouettes are shown in Figure 5.3. This was done using standard image analysis functions in MATLAB Image Analysis toolbox.



Fig. 5.3: Silhouettes captured from experimental rods. Image courtesy of Victor Romero.

From the regions of interest a visual hull was defined in space and by defining the points that are inside the silhouettes an unordered point cloud of the rod was reconstructed. In order to define the centerline in space the centroid of a moving cylindrical selection was used through the point cloud as depicted in Figure 5.4.




Fig. 5.4: Capturing the centerline from unordered point clouds. Image courtesy of Victor Romero.

In ongoing work we are attempting to use a similar process to capture naturally occurring rods, such as curly human hair and plant tendrils.

5.4 Synthetic Dataset Generation from Verified Simulator

In order to generate the dataset we use the verified simulator SUPER-HELIX [Ber+06]. The verification of the simulator was done by Victor Romero and performed using the so-called

Bend-Twist experiment. The experiment, is first described by [Mil+14] and replicated with additional data by Victor Romero in our work [Rom+21]. It involves an elastic curly rod that is clamped vertically at one end and left sagging under the action of gravity, with its other end being free. Depending on the material and physical parameters of the rod, the suspended equilibrium shape can be planar (with no twist), or deform into a 3D curly shape. The details of the experiment can be read in Romero *et al.* [Rom+21].

5.4.1 Non-Dimensional Data Generation

In order to explore the space of curved rods with a minimum number of parameters, we exploit the use of non-dimensional parameters such that various curvatures are represented in the dataset with a scaling factor. For suspended rods with natural bend and twist one needs 3 different parameters to completely characterize the deformation. These parameters are γ , ϕ_0 and ϕ_1 , where γ measures the weight against the resistance to bending and ϕ_0 and ϕ_1 determine the number of turns the rod makes and depend on the natural curvature of the rod:

$$\gamma = \frac{\pi r^2 \rho g L^3}{EI} \quad (5.2a)$$

$$\phi_0 = \kappa_0^0 * L \quad (5.2b)$$

$$\phi_1 = \kappa_1^0 * L \quad (5.2c)$$

where r is the radius of the cross-section, ρ is volumetric density and EI is the bending stiffness of the rod. With these parameters, Eq. 5.1, takes the non dimensional form:

$$\bar{K}_3 \frac{d}{d\bar{s}} (\bar{\kappa} - (\phi_0, 0, \phi_1)) + [\kappa]_{\times} \bar{K}_3 (\bar{\kappa} - (\phi_0, 0, \phi_1)) = \gamma (\bar{s} - 1) [e_x]_{\times} Re_z \quad (5.3)$$

We sample 100 values of γ between the range 100 – 1300, 10 values of ϕ_0 between the range 0.2 – 12 and 10 values of ϕ_1 between the range 6 – 20. The ranges of these parameters were determined experimentally to satisfy rendering constraints, such that the rod remains in frame for all 4 views. Furthermore, there also exist certain data points within these ranges, where the rod does not remain in frame. Such data points were manually sifted and discarded. Using these non dimensional parameters we create a dataset of 6215 data points. Out of which 5593 are randomly selected for training and 662 for testing. For each data point, we have the corresponding deformed curvature as its label. Each mesh is rendered using Blender, with the same calibrated viewpoints as the ones described for the experimental data. A sample of the dataset is shown in Figure 5.5.




Fig. 5.5: Samples from the synthetic dataset. Each data point is represented by 4 views and rendered with a randomly selected colour.

5.5 Model Description

We now specify our deep learning model to estimate the curvature of the suspended rod. The input to our model is a centerline of the rod and the output is the 3D deformed curvature of the rod, both represented by vectors of dimension 150×3 . Training on centerline data, makes the learning process agnostic to visual condition of the rod and extraneous conditions such as lighting and camera parameters. The model is trained only on the synthetic data described in the previous section.

Name	Description
Input	Vector 150, 3
conv1a	Conv1D 3, 64, ReLU
conv1b	Conv1D 3, 64, ReLU
maxpool1	3, stride 1
conv2a	Conv1D 3, 128, ReLU
conv2b	Conv1D 3, 128, ReLU
maxpool2	2, stride 1
conv2a	Conv1D 3, 128, ReLU
conv2b	Conv1D 3, 128, ReLU
maxpool2	2, stride 1
fc1	512, ReLU
fc2	450, None
Reshape	Vector 150, 3

Tab. 5.1: Architecture for baseline deep learning model.

5.5.1 Baseline Model

Our baseline learning model uses 1D convolution layers and max pooling layers as feature extractors and fully connected layers to learn the regression between centerline and deformed curvature. The model is described in detail in Table 5.1.

We use ReLU activation between layers. We train our model using root mean square loss and use the Adadelta optimizer. Total training time is approximately 1 hour on an NVIDIA Titan X GPU.

5.6 Results and Analysis

Figure 5.6 and 5.7 show the results of deformed curvature prediction using the baseline model on synthetic test data. This is data not seen at the training time by the model. As we have the true deformed curvature available to us, we can make a direct comparison with the predicted curvature values. In this figure one can observe that the curvature is predicted near perfectly by the model.

In the case of experimental data, the deformed curvature predictions are used to solve Eq. 5.1 as described in Section 5.2 to acquire the natural shape. The natural shape and the known material frame are used to simulate the deformed shape of the rod using the SUPER-HELIX [Ber+06] simulator. The deformed centerline thus obtained is compared with the one acquired experimentally (model input). The results from this process are shown in Figures 5.8 and 5.9.

As can be seen from the figures, the results are not accurate. While the baseline model captures the global shape in many instances (number of curls, overall length etc), the local features do not match precisely. This can be attributed to the lack of local feature representation in our baseline model.

5.6.1 Future Outlook

In future work we aim to mitigate the discrepancies in experimental results for curvature prediction by pooling local and global features learned separately in our model. This can be done by explicitly training a model to estimate curvature on local patches. Furthermore we aim to generalize our prediction results to capture naturally occurring rods, such as human hair.




Fig. 5.6: Curvature predictions for synthetic test data that the model has not seen. The left panel of each sub-figure shows the comparison of deformed centerline acquired after the inversion process following curvature prediction (in black) with the true centerline (in red). The right panel shows the deformed curvature prediction (in black) and the true deformed curvature (in red).

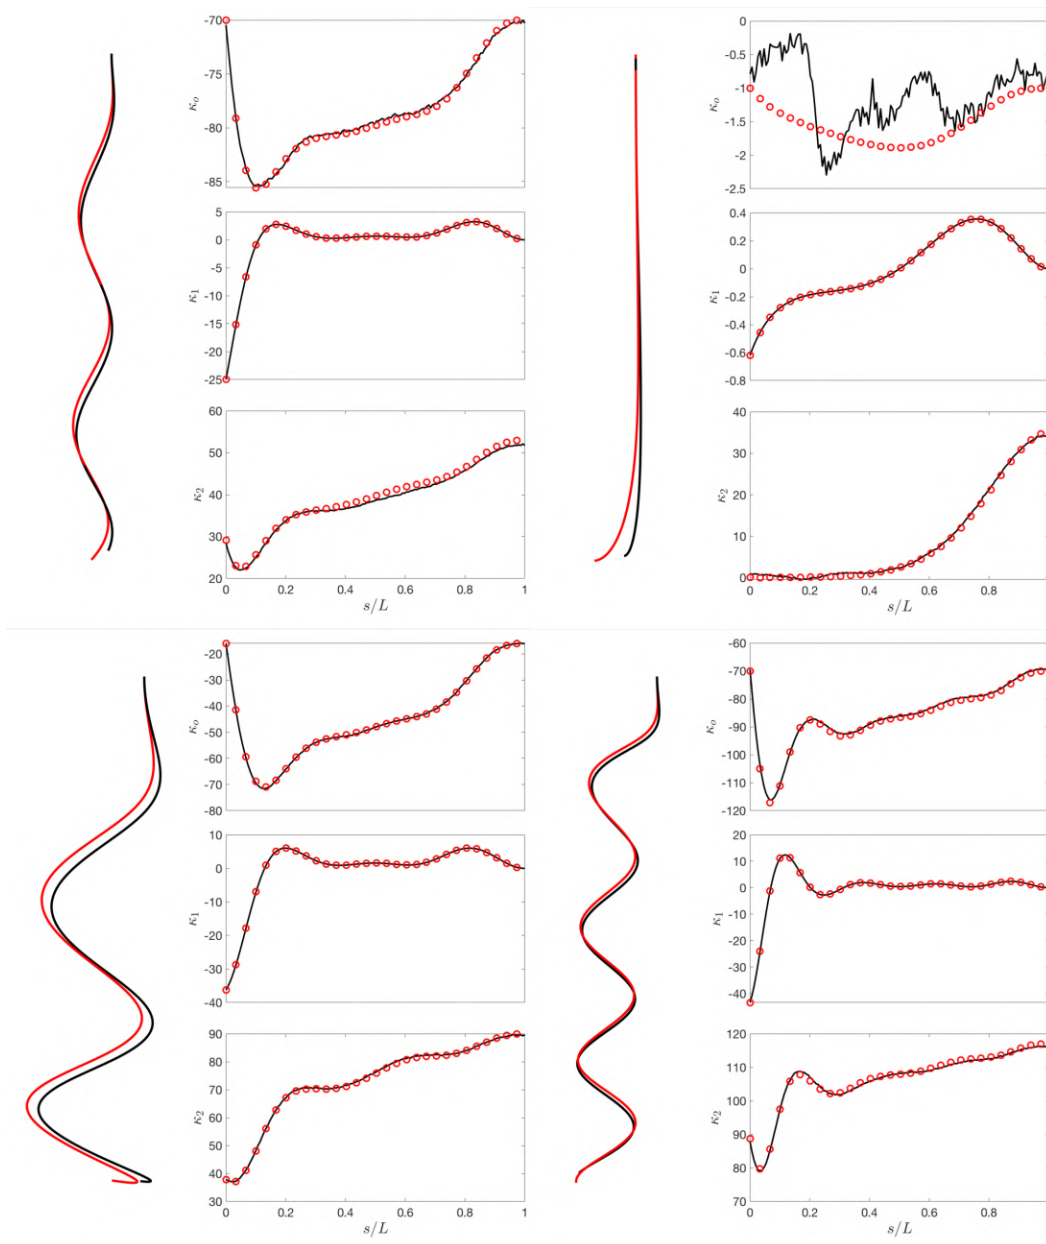


Fig. 5.7: Curvature predictions for synthetic test data that the model has not seen. The left panel of each sub-figure shows the comparison of deformed centerline acquired after the inversion process following curvature prediction (in black) with the true centerline (in red). The right panel shows the deformed curvature prediction (in black) and the true deformed curvature (in red).




Fig. 5.8: Curvature predictions for experimental test data. The left panel of each sub-figure shows the comparison of deformed centerline acquired after the inversion process following curvature prediction (in black) with the one acquired experimentally (in red). The right panel shows the deformed curvature predictions.

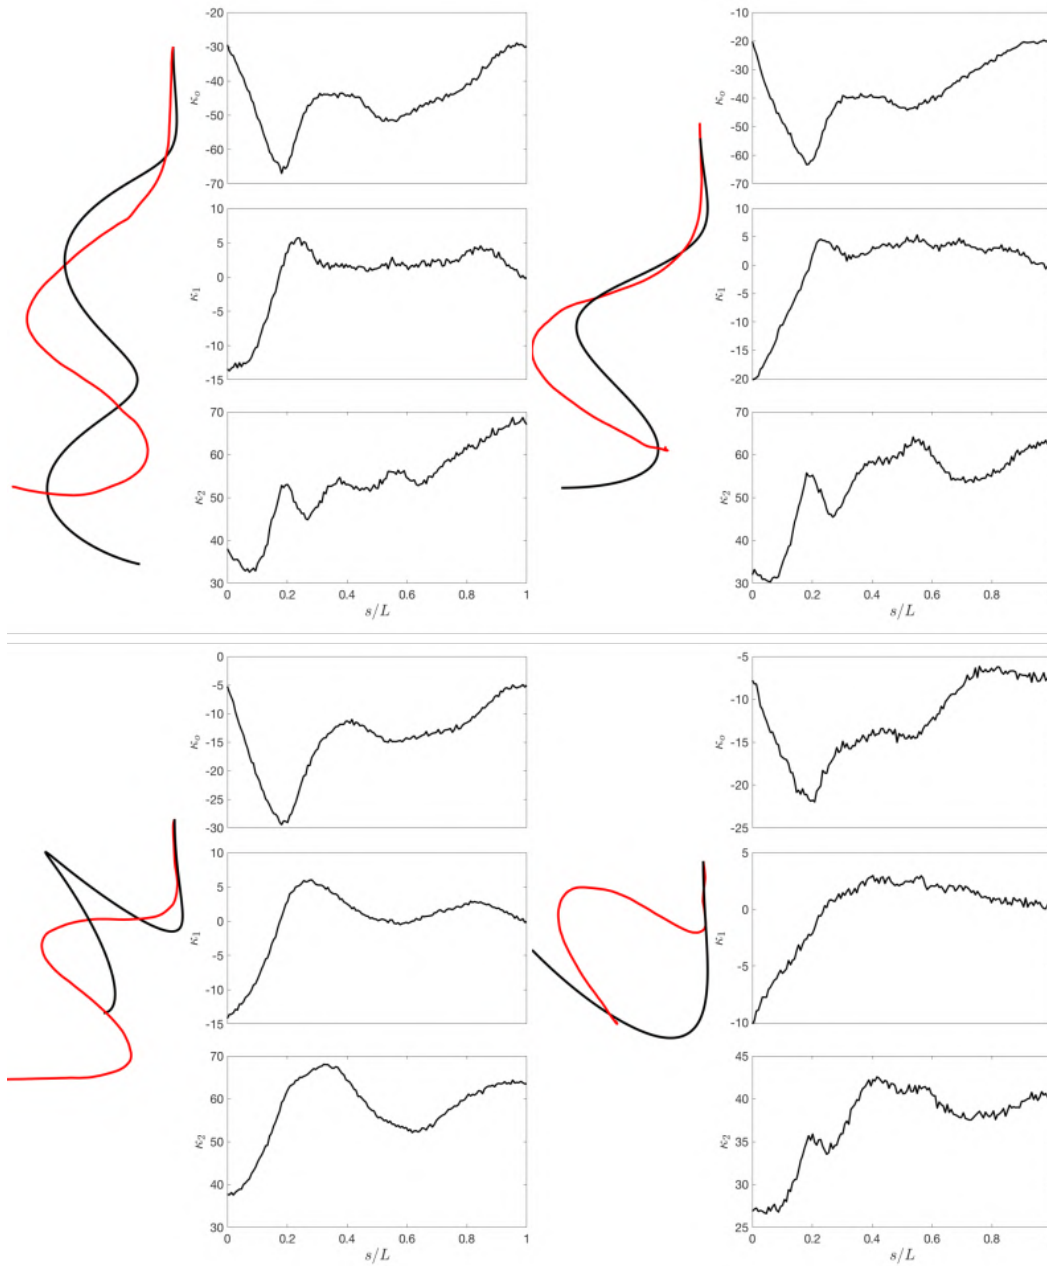


Fig. 5.9: Curvature predictions for experimental test data. The left panel of each sub-figure shows the comparison of deformed centerline acquired after the inversion process following curvature prediction (in black) with the one acquired experimentally (in red). The right panel shows the deformed curvature predictions.

Conclusion and Future Work

We propose a new approach to solving inverse measurement problems in soft body physics using deep neural networks. However rather than focusing on domain adaptation techniques or unique problem specific architecture innovations, our pipeline includes generating training data from a physically verified and calibrated simulators. Thus we end up training the neural network models on data that accurately represents the physical phenomenon being measured.

In this regard we present a new framework to quantitatively validate the physical realism of numerical simulators for plates and frictional contact. We introduce physical protocols from soft matter physics literature, which can be used to robustly compare any simulator against a clear master curve. This study is also aimed to help the Computer Graphics community to better validate and distribute their simulators, which can then be used for downstream applications that demand physical realism.

In order to demonstrate the effectiveness of our approach, 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 relies on a neural network model trained on data generated by a physics based simulator which has been verified using the physical protocols we presented. We conclusively show the effect of performing such verification on the simulator before using it to generate training data. These contributions open interesting future directions. 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.

The successful use of calibrated synthetic simulator-based training could also be transposed to other inverse parameter estimation problems. We demonstrate this by applying the same protocol to 3D curvature field estimation of suspended Kirchhoff rods. We show the utility of a neural network model to learn from synthetic data and make better than random predictions on experimental data. While the work is in its preliminary stage, a baseline model already shows promising results. In future work the model can be enhanced to capture local as well as global features of the rod to make it more robust and invariant to disparities between synthetic and real data. Furthermore the real data acquisition protocol

can be progressively relaxed in order to use the model for predictions on less restrained real world scenarios.

Bibliography

- [Amo99] G. Amontons. “De la résistance causée dans les machines”. In: *Mémoires de l’Académie royale* 4 (1699), pp. 257–282 (cit. on p. 11).
- [AMR95] U. Ascher, R. Mattheij, and R. Russell. *Numerical Solution of Boundary Value Problems for Ordinary Differential Equations*. Classics in Applied Mathematics. SIAM, 1995 (cit. on p. 25).
- [BW98] D. Baraff and A. Witkin. “Large Steps in Cloth Simulation”. In: *Computer Graphics Proceedings*. 1998, pp. 43–54 (cit. on pp. 3, 9, 15).
- [BWK03] D. Baraff, A. Witkin, and M. Kass. “Untangling Cloth”. In: *ACM Trans. Graph.* 22.3 (2003), pp. 862–870 (cit. on pp. 3, 9, 15).
- [Bar+16] A. Bartle, A. Sheffer, V. Kim, et al. “Physics-driven Pattern Adjustment for Direct 3D Garment Editing”. In: *ACM Trans. Graph.* 35.4 (July 2016), 50:1–50:11 (cit. on pp. 3, 16).
- [BG20] Yizhak Ben-Shabat and Stephen Gould. “DeepFit: 3D Surface Fitting via Neural Network Weighted Least Squares”. In: *ArXiv abs/2003.10826* (2020) (cit. on p. 14).
- [Ber+10] M. Bergou, B. Audoly, E. Vouga, M. Wardetzky, and E. Grinspun. “Discrete Viscous Threads”. In: *ACM Trans. Graph. (Proc. ACM SIGGRAPH’10)* 29.4 (2010) (cit. on p. 10).
- [Ber+08] M. Bergou, M. Wardetzky, S. Robinson, B. Audoly, and E. Grinspun. “Discrete elastic rods”. In: *ACM Trans. Graph. (Proc. ACM SIGGRAPH’08)* 27.3 (2008), pp. 1–12 (cit. on p. 10).
- [Ber+06] Florence Bertails, Basile Audoly, Marie-Paule Cani, et al. “Super-Helices for Predicting the Dynamics of Natural Hair”. In: *ACM Trans. Graph.* 25.3 (July 2006), pp. 1180–1187 (cit. on pp. 5, 13, 59, 62).
- [Ber+18] Florence Bertails-Descoubes, Alexandre Derouet-Jourdan, Victor Romero, and Arnaud Lazarus. “Inverse design of an isotropic suspended Kirchhoff rod: theoretical and numerical results on the uniqueness of the natural shape”. In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 474.2212 (2018), p. 20170837 (cit. on pp. 5, 13, 55, 56).
- [Bha+03] K.S. Bhat, C.D. Twigg, J.K. Hodgins, et al. “Estimating cloth simulation parameters from video”. In: *Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation*. Eurographics Association. 2003, pp. 37–51 (cit. on pp. 4, 13, 31).
- [Bic+09] B. Bickel, M. Bächer, M. Otaduy, et al. “Capture and Modeling of Non-Linear Heterogeneous Soft Tissue”. In: *ACM Trans. Graph.* 28.3 (July 2009) (cit. on p. 10).

- [Bic+10] B. Bickel, M. Bächer, M. Otaduy, et al. “Design and Fabrication of Materials with Desired Deformation Behavior”. In: *ACM Trans. Graph.* 29.4 (July 2010) (cit. on p. 10).
- [Bou+13] K.L. Bouman, B. Xiao, P. Battaglia, and W.T. Freeman. “Estimating the Material Properties of Fabric from Video”. In: *2013 IEEE International Conference on Computer Vision*. 2013, pp. 1984–1991 (cit. on pp. iii, v, 13).
- [BHT16] 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: *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*. 2016, pp. 428–435 (cit. on p. 12).
- [BMF03] R. Bridson, S. Marino, and R. Fedkiw. “Simulation of Clothing with Folds and Wrinkles”. In: *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. SCA '03. 2003, pp. 28–36 (cit. on p. 28).
- [BFA02] Robert Bridson, Ronald Fedkiw, and John Anderson. “Robust Treatment of Collisions, Contact and Friction for Cloth Animation”. In: *ACM Trans. Graph.* 21.3 (2002), pp. 594–603 (cit. on pp. 10, 11).
- [CPD11] R. Candelier, A. Prevost, and G. Debrégeas. “The Role of Exploratory Conditions in Bio-Inspired Tactile Sensing of Single Topological Features”. In: *Sensors* 11.8 (2011), pp. 7934–7953 (cit. on p. 12).
- [Car+19] Giuseppe Carleo, Ignacio Cirac, Kyle Cranmer, et al. “Machine learning and the physical sciences”. In: *Rev. Mod. Phys.* 91 (4 Dec. 2019), p. 045002 (cit. on p. 1).
- [Che+19] Yevgen Chebotar, Ankur Handa, Viktor Makoviychuk, et al. “Closing the Sim-to-Real Loop: Adapting Simulation Randomization with Real World Experience”. In: *2019 International Conference on Robotics and Automation (ICRA)*. 2019, pp. 8973–8979 (cit. on p. 1).
- [CFW13] Zhili Chen, Renguo Feng, and Huamin Wang. “Modeling Friction and Air Effects between Cloth and Deformable Bodies”. In: *ACM Trans. Graph.* 32.4 (2013) (cit. on p. 11).
- [Che17] N.R. Chevalier. “Hair-on-hair static friction coefficient can be determined by tying a knot”. In: *Colloids and Surfaces B: Biointerfaces* 159 (2017), pp. 924–928 (cit. on pp. 12, 31).
- [Cho+15] F. Chollet et al. *Keras*. <https://keras.io>. 2015 (cit. on p. 44).
- [CTT17] D. Clyde, J. Teran, and R. Tamstorf. “Modeling and data-driven parameter estimation for woven fabrics”. In: *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. ACM. 2017, p. 17 (cit. on p. 13).
- [CED19] E. Coevoet, A. Escande, and C. Duriez. “Soft robots locomotion and manipulation control using FEM simulation and quadratic programming”. In: *RoboSoft 2019 - IEEE International Conference on Soft Robotics*. Seoul, South Korea, Apr. 2019 (cit. on pp. 3, 16).
- [CT64] R. Courtel and L. Tichvinsky. “A brief history of friction”. In: *Naval Engineers Journal* 76.3 (1964), pp. 451–460 (cit. on pp. 12, 37).

- [CBL20] Kyle Cranmer, Johann Brehmer, and Gilles Louppe. “The frontier of simulation-based inference”. In: *Proceedings of the National Academy of Sciences* 117.48 (2020), pp. 30055–30062. eprint: <https://www.pnas.org/content/117/48/30055.full.pdf> (cit. on p. 1).
- [Dah68] Phil R Dahl. *A solid friction model*. Tech. rep. Aerospace Corp El Segundo Ca, 1968 (cit. on p. 12).
- [Dav20] G. Daviet. “Simple and Scalable Frictional Contacts for Thin Nodal Objects”. In: *ACM Trans. Graph.* 39.4 (July 2020) (cit. on pp. 3, 9, 15).
- [DBB11a] G. Daviet, F. Bertails-Descoubes, and L. Boissieux. “A Hybrid Iterative Solver for Robustly Capturing Coulomb Friction in Hair Dynamics”. In: *ACM Trans. Graph.* 30.6 (Dec. 2011), pp. 1–12 (cit. on pp. 3, 9).
- [DBB11b] G. Daviet, F. Bertails-Descoubes, and L. Boissieux. “A Hybrid Iterative Solver for Robustly Capturing Coulomb Friction in Hair Dynamics”. In: *ACM Trans. Graph.* 30.6 (2011), 139:1–139:12 (cit. on pp. 15, 16).
- [DB16] Gilles Daviet and Florence Bertails-Descoubes. “Nonsmooth simulation of dense granular flows with pressure-dependent yield stress”. In: *Journal of Non-newtonian Fluid Mechanics* 234 (2016), pp. 15–35 (cit. on p. 16).
- [Dav+17] A. Davis, K.L. Bouman, J.G. Chen, et al. “Visual Vibrometry: Estimating Material Properties from Small Motions in Video”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39.4 (2017), pp. 732–745 (cit. on pp. iii, v, 4, 13, 31).
- [DKK91] E. Doedel, H. Keller, and J.-P. Kernevez. “Numerical Analysis and Control of Bifurcation Problems (I) Bifurcation in Finite Dimensions”. In: *International Journal of Bifurcation and Chaos* 1.3 (1991), pp. 493–520 (cit. on p. 25).
- [Don+15] J. Donahue, L. Anne Hendricks, S. Guadarrama, et al. “Long-term recurrent convolutional networks for visual recognition and description”. In: *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*. 2015, pp. 2625–2634 (cit. on pp. 4, 41).
- [Dre43] E.C. Dreby. “A friction meter for determining the coefficient of kinetic friction of fabrics”. In: *Journal of Research of the National Bureau of Standards* 31.4 (1943), p. 237 (cit. on pp. 12, 41).
- [Duc06] V. Duclaux. “Pulmonary occlusions, eyelid entropion and aneurysm : a physical insight in physiology”. PhD thesis. Université de Provence - Aix-Marseille I, 2006 (cit. on pp. 2, 10, 24).
- [Erl+20] K. Erleben, M. Macklin, S. Andrews, and P. G. Kry. “The Matchstick Model for Anisotropic Friction Cones”. In: *Computer Graphics Forum* 39.1 (2020), pp. 450–461 (cit. on p. 11).
- [FY20] Yuwei Fan and Lexing Ying. “Solving electrical impedance tomography with deep learning”. In: *Journal of Computational Physics* 404 (2020), p. 109119 (cit. on p. 9).
- [FY19] Yuwei Fan and Lexing Ying. “Solving Inverse Wave Scattering with Deep Learning”. In: *ArXiv abs/1911.13202* (2019) (cit. on p. 9).

- [Far17] A. Fargette. “Soft Interfaces: from elastocapillary snap-through to droplet dynamics on elastomers”. PhD thesis. Université Pierre et Marie Curie - Paris VI, 2017 (cit. on pp. 2, 10, 24).
- [Gav+20] K. Gavriil, R. Guseinov, J. Pérez, et al. “Computational Design of Cold Bent Glass Façades”. In: *ACM Trans. Graph. (SIGGRAPH Asia 2020)* 39.6 (Dec. 2020) (cit. on pp. 3, 16).
- [Ger04] J. M. Gere. *Mechanics of Materials*. 6th. Thomson-Brooks/Cole, 2004 (cit. on p. 26).
- [Gis+20] Christoph Gissler, Andreas Henne, Stefan Band, Andreas Peer, and Matthias Teschner. “An Implicit Compressible SPH Solver for Snow Simulation”. In: *ACM Trans. Graph.* 39.4 (July 2020) (cit. on p. 9).
- [Gon+19] M.W. Gondal, M. Wüthrich, D. Miladinovic, et al. *On the Transfer of Inductive Bias from Simulation to the Real World: a New Disentanglement Dataset*. Tech. rep. arXiv, 2019 (cit. on pp. 9, 33).
- [Goo+19] Christopher T. Goodin, Suvash Sharma, Matthew Doude, et al. “Training of Neural Networks with Automated Labeling of Simulated Sensor Data”. In: *SAE Technical Paper Series* (2019) (cit. on p. 1).
- [GT98] Alain Goriely and Michael Tabor. “Spontaneous Helix Hand Reversal and Tendril Perversion in Climbing Plants”. In: *Phys. Rev. Lett.* 80 (7 Feb. 1998), pp. 1564–1567 (cit. on pp. 5, 13).
- [Gri+03] E. Grinspun, A. Hirani, M. Desbrun, and P. Schröder. “Discrete Shells”. In: *ACM SIGGRAPH - EG Symposium on Computer Animation (SCA'03)*. ACM-EG SCA. Aug. 2003, pp. 62–67 (cit. on pp. 28, 29).
- [Gus+20] R. Guseinov, C. McMahan, J. Pérez, C. Daraio, and B. Bickel. “Programming temporal morphing of self-actuated shells”. In: *Nature Communications* 11 (Jan. 2020) (cit. on pp. 3, 16).
- [Har+09] D. Harmon, E. Vouga, B. Smith, R. Tamstorf, and E. Grinspun. “Asynchronous Contact Mechanics”. In: *ACM Trans. Graph.* 28.3 (July 2009), 87:1–87:12 (cit. on p. 10).
- [Har+08] D. Harmon, E. Vouga, R. Tamstorf, and E. Grinspun. “Robust Treatment of Simultaneous Collisions”. In: *ACM Trans. Graph.* 27.3 (Aug. 2008), 23:1–23:4 (cit. on p. 10).
- [Jam+19] S. James, P. Wohlhart, M. Kalakrishnan, et al. “Sim-to-real via sim-to-sim: Data-efficient robotic grasping via randomized-to-canonical adaptation networks”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2019, pp. 12627–12637 (cit. on pp. 9, 33).
- [Kau+14] D. Kaufman, R. Tamstorf, B. Smith, J.-M. Aubry, and E. Grinspun. “Adaptive Nonlinearity for Collisions in Complex Rod Assemblies”. In: *ACM Trans. Graph.* 33.4 (July 2014), 123:1–123:12 (cit. on pp. 3, 9, 16).
- [Kir] Gustav R. Kirchhoff. “Ueber das Gleichgewicht und die Bewegung eines unendlich dünnen elastischen Stabes.” In: *Journal für die reine und angewandte Mathematik (Crelles Journal)* 1859 (), pp. 285–313 (cit. on p. 13).

- [KSH12] A. Krizhevsky, I. Sutskever, and G.E. Hinton. “Imagenet classification with deep convolutional neural networks”. In: *Advances in Neural Information Processing Systems*. 2012, pp. 1097–1105 (cit. on p. 41).
- [KNM10] S. Kunitomo, S. Nakamura, and S. Morishima. “Optimization of Cloth Simulation Parameters by Considering Static and Dynamic Features”. In: *ACM SIGGRAPH 2010 Posters*. SIGGRAPH ’10. Los Angeles, California: ACM, 2010, 15:1–15:1 (cit. on p. 13).
- [Lac+19] F. Laccone, L. Malomo, J. Pérez, et al. “FlexMaps Pavilion: a twisted arc made of mesostructured flat flexible panels”. In: *FORM and FORCE, IASS Symposium 2019, Structural Membranes 2019*. CIMNE, Oct. 2019, pp. 498–504 (cit. on pp. 3, 16).
- [LL59] L. D. Landau and E. M. Lifshitz. *Theory of elasticity*. English. Pergamon London, 1959 (cit. on p. 25).
- [LG21] Luis Ángel Larios-Cárdenas and Frédéric Gibou. “A Deep Learning Approach for the Computation of Curvature in the Level-Set Method”. In: *SIAM J. Sci. Comput.* 43 (2021), A1754–A1779 (cit. on p. 14).
- [Lew+05] Thomas Lewiner, João D. Gomes, Hélio Lopes, and Marcos Craizer. “Curvature and torsion estimators based on parametric curve fitting”. In: *Computers Graphics* 29.5 (2005), pp. 641–655 (cit. on p. 14).
- [Li+18a] J. Li, G. Daviet, R. Narain, et al. “An Implicit Frictional Contact Solver for Adaptive Cloth Simulation”. In: *ACM Trans. Graph.* 37.4 (Aug. 2018) (cit. on pp. 10, 16).
- [Li+18b] J. Li, G. Daviet, R. Narain, et al. “An implicit frictional contact solver for adaptive cloth simulation”. In: *ACM Transactions on Graphics (TOG)* 37.4 (2018), p. 52 (cit. on pp. 2, 4, 9, 16, 39).
- [LKJ20] M. Li, D. Kaufman, and C. Jiang. *Codimensional Incremental Potential Contact*. 2020 (cit. on p. 10).
- [LLK19] J. Liang, M.C. Lin, and V. Koltun. “Differentiable Cloth Simulation for Inverse Problems”. In: *Conference on Neural Information Processing Systems*. 2019 (cit. on p. 13).
- [LR95] Enno Littmann and Helge Ritter. “Curvature Estimation with a DCA neural network”. In: *Mustererkennung 1995*. Ed. by Gerhard Sagerer, Stefan Posch, and Franz Kummert. Berlin, Heidelberg: Springer Berlin Heidelberg, 1995, pp. 610–617 (cit. on p. 14).
- [Liu+08] X. Liu, Z. Yue, Z. Cai, D. G. Chetwynd, and S. Smith. “Quantifying touch-feel perception: Tribological aspects”. In: *Measurement Science & Technology - MEAS SCI TECHNOL* 19 (2008) (cit. on p. 12).
- [M A+15] M. Abadi, A. Agarwal, P. Barham, et al. *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. Software available from tensorflow.org. 2015 (cit. on p. 44).
- [MWQ16] N. Mao, Y. Wang, and J. Qu. “Smoothness and roughness: characteristics of fabric-to-fabric self-friction properties”. In: *The Proceedings of 90th Textile Institute World Conference*. The Textile Institute. 2016 (cit. on p. 12).
- [Mar+19] J. Martínez, M. Skouras, C. Schumacher, et al. “Star-Shaped Metrics for Mechanical Metamaterial Design”. In: *ACM Trans. Graph.* 38.4 (July 2019), Article No. 82 :1–13 (cit. on pp. 3, 16).

- [McA+09] A. McAdams, A. Selle, K. Ward, E. Sifakis, and J. Teran. “Detail preserving continuum simulation of straight hair”. In: *ACM Trans. Graph. (Proc. ACM SIGGRAPH’09)* 28.3 (2009), pp. 1–6 (cit. on pp. 3, 9, 15).
- [Mig+12a] E. Miguel, D. Bradley, B. Thomaszewski, et al. “Data-Driven Estimation of Cloth Simulation Models”. In: *Computer Graphics Forum* 31.2 (May 2012) (cit. on p. 10).
- [Mig+12b] E. Miguel, D. Bradley, B. Thomaszewski, et al. “Data-driven estimation of cloth simulation models”. In: *Computer Graphics Forum*. Vol. 31. Wiley Online Library. 2012, pp. 519–528 (cit. on p. 13).
- [Mig+13] E. Miguel, R. Tamstorf, D. Bradley, et al. “Modeling and Estimation of Internal Friction in Cloth”. In: *Transactions on Graphics* 32.6 (2013), 212:1–212:10 (cit. on p. 12).
- [Mik06] Y Mikata. “Complete solution of elastica for a clamped-hinged beam, and its applications to a carbon nanotube”. In: *Acta mechanica* 190.1-4 (Oct. 2006), pp. 133–150 (cit. on p. 19).
- [Mil+14] J. Miller, A. Lazarus, B. Audoly, and P. Reis. “Shapes of a Suspended Curly Hair”. In: *Physical Review Letters* 112.6 (2014) (cit. on p. 60).
- [MK15] R.R. Moorthy and P. Kandhavadi. “Surface friction characteristics of woven fabrics with nonconventional fibers and their blends”. In: *Journal of Textile and Apparel, Technology and Management* 9.3 (2015) (cit. on pp. 4, 12).
- [Mus20] K. Museth. *Physics simulations: Is it Hollywood magic or rocket science*. Oct. 2020 (cit. on p. 3).
- [NSO12] R. Narain, A. Samii, and J.F. O’Brien. “Adaptive Anisotropic Remeshing for Cloth Simulation”. In: *ACM Trans. Graph.* 31.6 (Nov. 2012), 152:1–152:10 (cit. on p. 16).
- [Neu04] Sébastien Neukirch. “Extracting DNA Twist Rigidity from Experimental Supercoiling Data”. In: *Phys. Rev. Lett.* 93 (19 Nov. 2004), p. 198107 (cit. on p. 13).
- [ND07] Thanh Phuong Nguyen and Isabelle Debled-Renneson. “Curvature Estimation in Noisy Curves”. In: *Computer Analysis of Images and Patterns*. Ed. by Walter G. Kropatsch, Martin Kampel, and Allan Hanbury. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 474–481 (cit. on p. 14).
- [PTS09] Simon Pabst, Bernhard Thomaszewski, and Wolfgang Straßer. “Anisotropic Friction for Deformable Surfaces and Solids”. In: *Proceedings of the 2009 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. New Orleans, Louisiana, 2009, pp. 149–154 (cit. on p. 11).
- [Pan+19] J. Panetta, M. Konaković-Luković, F. Isvoranu, E. Bouleau, and M. Pauly. “X-Shells: A New Class of Deployable Beam Structures”. In: *ACM Trans. Graph.* 38.4 (July 2019) (cit. on pp. 3, 16).
- [Pat+19] H. V. Patel, A. Panda, J. A. M. Kuipers, and Eajf Frank Peters. “Computing interface curvature from volume fractions: A machine learning approach”. In: *Computers & Fluids* (2019) (cit. on p. 14).

- [Qi+19] Yinghe Qi, Jiakai Lu, Ruben Scardovelli, Stéphane Zaleski, and Grétar Tryggvason. “Computing curvature for volume of fluid methods using machine learning”. In: *Journal of Computational Physics* 377 (2019), pp. 155–161 (cit. on p. 14).
- [Ras+20] A.H. Rasheed, V. Romero, F. Bertails-Descoubes, et al. “Learning to Measure the Static Friction Coefficient in Cloth Contact”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2020 (cit. on pp. 3, 16, 33).
- [RFM10] Frédéric Rieux, Christophe Fiorio, and Christian Mercat. “Curvature Estimation for Discrete Curves Based on Auto-adaptive Masks of Convolution”. In: May 2010 (cit. on p. 14).
- [Rom+19] Victor Romero, Florence Bertails-Descoubes, Alexandre Derouet-Jourdan, and Arnaud Lazarus. “Inverse design of a suspended Kirchhoff rod: From theory to practice”. In: *APS 2019 - American Physical Society March Meeting*. Boston, United States, Mar. 2019 (cit. on p. 55).
- [Rom+21] Victor Romero, Mickaël Ly, Abdullah Haroon Rasheed, et al. “Physical validation of simulators in Computer Graphics: A new framework dedicated to slender elastic structures and frictional contact”. In: *ACM Transactions on Graphics* (Aug. 2021), pp. 1–18 (cit. on pp. 3, 16, 17, 26, 27, 29, 60).
- [SYW17] T.G. Sano, T. Yamaguchi, and H. Wada. “Slip Morphology of Elastic Strips on Frictional Rigid Substrates”. In: *Physical Review Letters* 118 (2017) (cit. on pp. 2, 10, 12, 17, 31, 33).
- [Sch+18] C. Schumacher, S. Marschner, M. Gross, and B. Thomaszewski. “Mechanical Characterization of Structured Sheet Materials”. In: *ACM Trans. Graph.* 37.4 (July 2018) (cit. on pp. 3, 16).
- [Sel+09] Andrew Selle, Jonathan Su, Geoffrey Irving, and Ronald Fedkiw. “Robust High-Resolution Cloth Using Parallelism, History-Based Collisions, and Accurate Friction”. In: *IEEE Transactions on Visualization and Computer Graphics* 15.2 (2009), pp. 339–350 (cit. on p. 11).
- [Shi92] R. T. Shield. “Bending of a beam or wide strip”. In: *Quarterly Journal of Mechanics and Applied Mathematics* 45.4 (1992), pp. 567–573 (cit. on p. 25).
- [SZ14] K. Simonyan and A. Zisserman. “Very deep convolutional networks for large-scale image recognition”. In: *arXiv preprint arXiv:1409.1556* (2014) (cit. on p. 41).
- [Sko+12] M. Skouras, B. Thomaszewski, B. Bickel, and M. Gross. “Computational Design of Rubber Balloons”. In: *Comput. Graphics Forum (Proc. Eurographics)* (2012) (cit. on p. 10).
- [Sko+14] M. Skouras, B. Thomaszewski, P. Kaufmann, et al. “Designing Inflatable Structures”. In: *ACM Trans. Graph.* 33.4 (July 2014) (cit. on p. 10).
- [Smi+12] B. Smith, D. Kaufman, E. Vouga, R. Tamstorf, and E. Grinspun. “Reflections on Simultaneous Impact”. In: *ACM Trans. Graph.* 31.4 (July 2012) (cit. on p. 10).
- [ST07] J. Spillmann and M. Teschner. “CoRdE: Cosserat rod elements for the dynamic simulation of one-dimensional elastic objects”. In: *ACM SIGGRAPH - EG Symposium on Computer Animation (SCA’07)*. ACM-EG SCA. 2007, pp. 63–72 (cit. on p. 10).

- [SB08] C. Syllebranque and S. Boivin. “Estimation of mechanical parameters of deformable solids from videos”. In: *The Visual Computer* 24.11 (2008), pp. 963–972 (cit. on p. 12).
- [TV61] G.H. Thorndike and L. Varley. “Measurement of the coefficient of friction between samples of the same cloth”. In: *Journal of the Textile Institute Proceedings* 52.6 (1961), P255–P271 (cit. on pp. 4, 12, 37).
- [Tob+17] J. Tobin, R. Fong, A. Ray, et al. “Domain randomization for transferring deep neural networks from simulation to the real world”. In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2017, pp. 23–30 (cit. on pp. 1, 9, 33).
- [Van+20] F. Vanneste, O. Goury, J. Martinez, et al. “Anisotropic soft robots based on 3D printed meso-structured materials: design, modeling by homogenization and simulation”. In: *IEEE Robotics and Automation Letters* 5.2 (Jan. 2020), pp. 2380–2386 (cit. on pp. 3, 16).
- [Wan+15] B. Wang, L. Wu, K. Yin, et al. “Deformation capture and modeling of soft objects”. In: *ACM Transactions on Graphics (TOG)* 34.4 (2015), pp. 94–1 (cit. on p. 12).
- [Wan18] H. Wang. “Rule-Free Sewing Pattern Adjustment with Precision and Efficiency”. In: *ACM Trans. Graph.* 37.4 (July 2018) (cit. on pp. 3, 16).
- [WOR11] H. Wang, J.F. O’Brien, and R. Ramamoorthi. “Data-driven elastic models for cloth: modeling and measurement”. In: *ACM Transactions on Graphics (TOG)*. Vol. 30. ACM. 2011, p. 71 (cit. on pp. 13, 28, 34, 35, 40).
- [WRO11] H. Wang, R. Ramamoorthi, and J. O’Brien. “Data-driven elastic models for cloth: modeling and measurement”. In: *ACM Trans. Graph.* 30.4 (Aug. 2011), 71:1–71:12 (cit. on p. 10).
- [Wu+16] J. Wu, J.J. Lim, H. Zhang, J.B. Tenenbaum, and W.T. Freeman. “Physics 101: Learning physical object properties from unlabeled videos”. In: *British Machine Vision Conference*. 2016 (cit. on p. 2).
- [Wu+15] 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: *Advances in Neural Information Processing Systems*. 2015, pp. 127–135 (cit. on p. 12).
- [YLL17a] S. Yang, J. Liang, and M.C. Lin. “Learning-based cloth material recovery from video”. In: *Proceedings of the IEEE International Conference on Computer Vision*. 2017, pp. 4383–4393 (cit. on pp. iii, v, 4, 13, 15, 16, 28, 31, 41).
- [YLL17b] S. Yang, J. Liang, and M.C. Lin. “Learning-based cloth material recovery from video”. In: *Proceedings of the IEEE International Conference on Computer Vision*. 2017, pp. 4383–4393 (cit. on p. 3).
- [YL16] S. Yang and M.C. Lin. “MaterialCloning: Acquiring elasticity parameters from images for medical applications”. In: *IEEE Transactions on Visualization and Computer Graphics* 22.9 (2016), pp. 2122–2135 (cit. on pp. 12, 15).
- [Yua+17] W. Yuan, S. Wang, S. Dong, and E. Adelson. “Connecting Look and Feel: Associating the visual and tactile properties of physical materials”. In: *Conference on Computer Vision and Pattern Recognition*. 2017 (cit. on p. 12).

- [Zei12] M.D. Zeiler. “ADADELTA: an adaptive learning rate method”. In: *arXiv preprint arXiv:1212.5701* (2012) (cit. on p. 44).
- [ZDN16] 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: *European Conference on Computer Vision*. 2016 (cit. on p. 12).
- [ZQW20] Wenshuai Zhao, Jorge Peña Queralta, and Tomi Westerlund. “Sim-to-Real Transfer in Deep Reinforcement Learning for Robotics: a Survey”. In: *2020 IEEE Symposium Series on Computational Intelligence (SSCI)*. 2020, pp. 737–744 (cit. on p. 1).

List of Figures

1.1	Depiction of the Stick-Slip experiment implemented, both physically and in simulation [Rom+21]	3
1.2	Scene from re-simulation application, where a correct estimation of cloth-cloth friction is necessary to replicate the behaviour of selected cloth materials.	5
1.3	Depiction of a simulated suspended rod over a physical physical one showing very good agreement. The simulation has been made using the exact known material parameters as that of the synthetically constructed physical rod.	6
3.1	Stick-Slip experiment schematic for $\epsilon_y = 0.2$. The superposition of the numerical planar elastica solution (blue line), for the same value $\epsilon_y = 0.2$, illustrates the very good agreement between experiments and numerics	18
3.2	Master curve and experimental validation for the Stick-Slip test. Phase diagram where the master curve (in black) separates the sticking regime (in orange) from the slipping regime (in turquoise). The extended contact region (in gray) is not used in our validation protocol. The (numerical) black and (experimental) purple curves are plots of the ratio Q/P as a function of ϵ_y	20
3.3	ARGUS Results on Stick-Slip test with Adaptive Remeshing	22
3.4	ARGUS Results on Stick-Slip test without Adaptive Remeshing	23
3.5	Master curve and experimental validation for the Cantilever test. Left: snapshots of experimental rods under gravity, with 16 different Γ values. Right: our computed master curve (in black) representing the aspect ratio of the rod's shape as a function of Γ , together with experimental results from [Rom+21]	26
3.6	ARCSIM results on Cantilever experiment. Figure taken from our collective work [Rom+21]	27
3.7	ARCSIM results on Cantilever experiment with Grinspun <i>et al.</i> [Gri+03] Bending Energy. Figure taken from our collective work [Rom+21]	29
4.1	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 $\mu = 0.6$	32

4.2	Experimental images with various hard substrates for the same cloth material. From left to right: Aluminium, Aluminium-PET, Ceramic, Rough Glass, Smooth Glass, Polyester, and Stainless.	34
4.3	Experimental images with different textile substrates for the same cloth material	36
4.4	Experimental setup for real data-set creation	36
4.5	Sketch of the inclined plane setup for frictional characterisation	37
4.6	Distribution of baseline friction values (μ) for all material-substrate pairs. . .	38
4.7	Distribution of baseline friction values (μ) for all cloth-cloth pairs.	39
4.8	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.	40
4.9	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.	42
4.10	Cumulative error plots for all datasets	48
4.11	Cumulative error plots for dataset High-accurate-small-timestep (red, black, light blue, green) and simulated test data (dark blue).	49
4.12	Distribution of baseline friction values for cloth to cloth dataset. The orange bars are outside the range of our synthetic training data.	50
4.13	Extension of our synthetic dataset, with renderings more similar to the new cloth on cloth dataset	51
4.14	Prediction errors on Cloth to Cloth friction test dataset	51
4.15	Corresponding frames showing the cloth coming into contact with a rotating sphere substrate. Top Row: M02 (Polyester 95 %, Spandex 5 %) with $\mu = 0.3$, substrate: aluminium, Bottom Row M03 (Cotton 100 %) with $\mu = 0.6$, substrate: ceramic	52
4.16	Corresponding frames of a simulated character displaying cloth on cloth friction. Top Row: M02 (Polyester 95 %, Spandex 5 %) with $\mu = 0.3$, Bottom Row M06 (100 % Polyester) with $\mu = 0.7$	53
5.1	Experimental setup of Victor Romero, showing a manufactured suspended rod being captured.	57
5.2	Camera positions shown for capture of the experimental rod. Image courtesy of Victor Romero.	58
5.3	Silhouettes captured from experimental rods. Image courtesy of Victor Romero.	59
5.4	Capturing the centerline from unordered point clouds. Image courtesy of Victor Romero.	59
5.5	Samples from the synthetic dataset. Each data point is represented by 4 views and rendered with a randomly selected colour.	61

5.6	Curvature predictions for synthetic test data that the model has not seen. The left panel of each sub-figure shows the comparison of deformed centerline acquired after the inversion process following curvature prediction (in black) with the true centerline (in red). The right panel shows the deformed curvature prediction (in black) and the true deformed curvature (in red).	63
5.7	Curvature predictions for synthetic test data that the model has not seen. The left panel of each sub-figure shows the comparison of deformed centerline acquired after the inversion process following curvature prediction (in black) with the true centerline (in red). The right panel shows the deformed curvature prediction (in black) and the true deformed curvature (in red).	64
5.8	Curvature predictions for experimental test data. The left panel of each sub-figure shows the comparison of deformed centerline acquired after the inversion process following curvature prediction (in black) with the one acquired experimentally (in red). The right panel shows the deformed curvature predictions.	65
5.9	Curvature predictions for experimental test data. The left panel of each sub-figure shows the comparison of deformed centerline acquired after the inversion process following curvature prediction (in black) with the one acquired experimentally (in red). The right panel shows the deformed curvature predictions.	66

List of Tables

4.1	Experimental Cloth Materials	35
4.2	Parameter specifications of 3 datasets, generated for varying levels of accuracy.	40
4.3	Architecture details for a feature extractor block.	42
4.4	Results on Simulated Test Data	45
4.5	Results on unseen viewpoints	46
5.1	Architecture for baseline deep learning model.	61