

Universidad Rey Juan Carlos

Departamento de Ciencias de la Computación, Arquitectura de
la Computación, Lenguajes y Sistemas Informáticos y
Estadística e Investigación Operativa

Measurement-Based Model Estimation for Deformable Objects

Tesis doctoral

Autor: D. Eder Miguel Villalba
Director: Dr. D. Miguel A. Otaduy Tristán

Escuela Técnica Superior de Ingeniería Informática

Septiembre 2014

Dr. Miguel A. Otaduy Tristán, Profesor titular de la Universidad Rey Juan Carlos, con NIE 72447035W,

CERTIFICA: Que Don Eder Miguel Villalba, licenciado en Ingeniería de Telecomunicaciones, ha realizado en el Dpto. de Ciencias de la Computación, Arquitectura de la Computación, Lenguajes y Sistemas Informáticos y Estadística e Investigación Operativa, bajo su dirección, el trabajo correspondiente a la tesis doctoral titulada:

Measurement-Based Model Estimation for Deformable Objects

Revisado el presente trabajo, estima que puede ser presentado al tribunal que ha de juzgarlo. Y para que conste a efecto de lo establecido en la normativa reguladora del tercer ciclo de la Universidad Rey Juan Carlos, autoriza su presentación.

Móstoles, 24 de Septiembre de 2014.

Fdo.: Miguel A. Otaduy

Abstract

Deformable objects play a critical role in our life due to their compliance. Clothing and support structures, such as mattresses, are just a few examples of their use. They are so common that an accurate prediction of their behavior under a variety of environments and situations is mandatory in order to design products with the desired functionalities.

However, obtaining realistic simulations is a difficult task. Both, an appropriate deformation model and parameters that produce the desired behavior must be used. On one hand, there exist many deformation models for elasticity, but there are few capable of capturing other complex effects that are critical in order to obtain the desired realism. On the other hand, the task of estimating model parameters is usually performed using a trial-and-error method, with the corresponding waste in time.

In this thesis we develop novel deformation models and parameter estimation methods that allow us to increase the realism of deformable object simulations. We present deformation models that capture several of these complex effects: hyperelasticity, extreme nonlinearities, heterogeneities and internal friction. In addition, we design parameter estimation methods that take advantage of the structure of the measured data and avoid common problems that arise when numerical optimization algorithms are used.

First, we focus on cloth and present a novel measurement system that captures the behavior of cloth under a variety of experiments. It produces a complete set of information including the 3D reconstruction of the cloth sample under test as well as the forces being applied. We design a parameter estimation pipeline and use this system to estimate parameters for several popular cloth models and evaluate their performance and suitability in terms of quality of the obtained estimations.

We then develop a novel, general and flexible deformation model based on additive energy density terms. By using independent components this model allows us to isolate the effect that each one has on the global behavior of the deformable object, replicate existing deformation models and produce new ones. It also allows us to apply incremental approaches to parameter estimation. We demonstrate its advantages by applying it in a wide variety of scenarios, including cloth simulation, modeling of heterogeneous soft tissue and capture of

extreme nonlinearities in finger skin.

Finally, a fundamental observation extracted from the estimation of parameters for cloth models is that, in real-world, cloth hysteresis has a huge effect in the mechanical behavior and visual appearance of cloth. The source of hysteresis is the internal friction produced by the interactions between yarns. Mechanically, it can produce very different deformations in the loading or unloading cycles, while visually, it is responsible for effects such as persistent deformations, preferred wrinkles or history-dependent folds. We develop an internal friction model, present a measurement and estimation system that produces elasticity and internal friction parameters, and analyse the visual impact of internal friction in cloth simulation.

Acknowledgements

First of all, I would like to thank my advisor Miguel A. Otaduy. When I decided to accept your offer to do a PhD, I had no idea I would learn and enjoy it so much. Thank you for so many things, but above all for your guidance and patience, and for pushing me to never give up on deadlines.

I would like to thank all the collaborators I had during these years: Derek Bradley, David Miraut, Bernhard Thomaszewski, Rasmus Tamstorf, Sara C. Schvartzman, Bernd Bickel, Steve Marschner, Wojciech Matusik, Andrew Feng, Yuyu Xu and Ari Shapiro. I have learnt so much from you.

Some of the greatest times during my PhD. happened during my internships. I want to thank my supervisors for trusting me: Bernd Bickel, Wojciech Matusik, Rasmus Tamstorf and Ari Shapiro. Thank you for the opportunity to work in such amazing places, meet brilliant people, learn from them and enjoy my time there. I would like to thank Miquel, Alexey, Iker, Magda and Laura, with whom I had great times.

Finishing this PhD. would have been way harder without the help of the GMRV research group. It has been an amazing experience working with you all. Special mention to my office-mates: Sara, Juanpe, Loïc and Jaime, with whom I have shared many discussions and good times. And Angela, who has always helped us in so many ways.

I want to thank the members of my PhD. committee, for evaluating this thesis, and the reviewers that provided feedback for each one of our publications.

This thesis would not have been possible without the funding from the Spanish Science and Innovation Dept. (projects TIN2012-35840 and TIN2009-07942) and by the European Research Council (ERC-2011-StG-280135 Animetrics).

I would like to thank Lucky, Gorka, Luis, Pedro, Unai and Isra. You have always been there for the good and the not so good times, helped me look at things from different perspectives and supported me, even when I went missing during deadlines. I want to thank Evgenia, who encouraged me to keep working hard even when things seemed not to progress.

Finally, I would like to thank my parents, Maria del Carmen and Jose Antonio, and my sister Gurutze, for their support all this time.

Contents

Contents	ix
List of Figures	xiii
List of Tables	xv
Nomenclature	xv
1 Introduction	1
1.1 Data-Driven Elasticity Estimation	3
1.2 Internal Friction Modeling and Estimation	3
1.3 Contributions	4
1.4 Organization	5
2 Related Work	7
2.1 Modeling of Deformable Objects	7
2.1.1 Modeling of Hyperelasticity	8
2.1.2 Modeling of Heterogeneity	9
2.1.3 Modeling of Internal Friction in Cloth	9
2.1.4 Modeling of Plasticity	11
2.2 Measurement-Based Parameter Estimation	12
2.2.1 Parameter Estimation for Hyperelastic Heterogeneous Objects	12
2.2.2 Parameter Estimation for Internal Friction in Cloth	14
3 Data-Driven Estimation of Cloth Simulation Models	15
3.1 Introduction	15
3.2 Measurement system	16
3.2.1 Reconstruction	18

3.2.2	Measurements	20
3.2.3	Accuracy	20
3.3	Cloth Models	21
3.3.1	Strain Metrics	21
3.3.2	Nonlinear Stress-Strain Curves	24
3.4	Fitting	24
3.4.1	Optimization Problem	25
3.4.2	Incremental Parameter Fitting	27
3.5	Results	28
3.6	Discussion and Future Work	31
4	Data-Driven Modeling of Hyperelastic Objects	33
4.1	Introduction	33
4.2	A General Model of Hyperelasticity	35
4.2.1	Additive Energy Model	35
4.2.2	Interpolated Energy Functions	36
4.2.3	Convexity	37
4.2.4	Heterogeneity	38
4.2.5	Energy Constraints	39
4.2.6	FEM Simulation	39
4.3	Data-Driven Material Estimation	40
4.3.1	Parameter Estimation	41
4.3.2	Static Equilibrium	43
4.3.3	Convex Energy Projection	43
4.4	Results	44
4.5	Discussion and Future Work	48
5	Modeling and Estimation of Internal Friction for Cloth	51
5.1	Introduction	51
5.2	Dahl's Friction Model	52
5.2.1	Underlying Elasticity Model	53
5.2.2	Characterizing Hysteresis	53
5.2.3	Model and Interpretation	54
5.2.4	Parameters for Estimation	55
5.3	Estimation of Friction Parameters	57
5.3.1	General Rationale	57

Contents	xi
<hr/>	
5.3.2 Estimation of Stretch Friction	58
5.3.3 Estimation of Bending Friction	61
5.4 Dynamic Simulation of Internal Friction	64
5.5 Results	69
5.6 Discussion and Future Work	73
6 Conclusion	77
6.1 Summary of results	77
6.2 Future Work	78
References	81
Appendix A Parameter Estimation Results for Cloth Models	87

List of Figures

1.1	Examples of applications that use deformable objects.	2
2.1	Force-deformation plot for an experiment with real cloth.	10
2.2	Persistent deformations and hysteresis under friction and plasticity models.	11
3.1	Acquisition setup for the measurement system.	17
3.2	Force measurements for selected frames of a corner pulling sequence.	19
3.3	Selected frames from isolated measurements.	21
3.4	Selected frames from elaborate cloth manipulation.	22
3.5	Raw force measurements.	28
3.6	Fitting results for Sample #12 with the Soft Constraints model.	30
4.1	Nonconservative elastic forces produce uncontrolled energy growth.	34
4.2	Non-convex strain energy density function produces divergin motion.	38
4.3	Evaluation of the fitting quality of the additive energy model.	42
4.4	Evaluation of the fitting quality of our energy model on 3D MRI data of a knee.	43
4.5	Cross-sections of the knee MRI in the pose with largest bending.	45
4.6	Plots of contact force vs. contact area for a fingerpad.	47
4.7	Simulation of the finger deformation capture experiment.	48
5.1	Force-deformation curves for load-unload cycles of uniform stretch.	54
5.2	Schematic of friction and elastic forces interaction.	55
5.3	Fits of Dahl's model.	56
5.4	Acquisition of stretch for a piece of cotton.	58
5.5	Fitting of stretch friction parameters.	59
5.6	Acquisition of bending for a strip of doublecloth.	60
5.7	Schematic of a sampled profile and evaluation of torques at sampled points.	61
5.8	Spiral profiles and curvatures for bending estimation.	63

5.9	Results and validation of bending friction estimation.	64
5.10	Comparison of captured and simulated unrolling spirals for multilayer cloth.	65
5.11	Comparison of captured and simulated unrolling spirals for felt.	66
5.12	Comparison of captured and simulated unrolling spirals for paper.	67
5.13	Persistent wrinkles.	69
5.14	Wrinkles on a cotton cylinder with stretch friction.	70
5.15	Example of preferred wrinkles.	71
5.16	Images of a cylinder of cotton once a twist motion is undone.	72
5.17	Simulation of cotton pants with and without stretch friction.	74
5.18	A character performing jumping jacks.	75
A.1	Fitting results for Sample #4, Soft Constraints model.	89
A.2	Fitting results for Sample #4, St. V-K model.	90
A.3	Fitting results for Sample #4, Spring model.	91
A.4	Fitting results for Sample #12, Soft Constraints model.	92
A.5	Fitting results for Sample #12, St. V-K model.	93
A.6	Fitting results for Sample #12, Spring model.	94
A.7	Fitting results for Sample #14, Soft Constraints model.	95
A.8	Fitting results for Sample #14, Soft Constraints model, isotropic.	96
A.9	Fitting results for Sample #14, Soft Constraints model, linear.	97
A.10	Fitting results for Sample #14, Soft Constraints model, linear isotropic.	98
A.11	Fitting results for Sample #14, St. V-K model.	99
A.12	Fitting results for Sample #14, Spring model.	100
A.13	Fitting results for Sample #18, Soft Constraints model.	101
A.14	Fitting results for Sample #18, St. V-K model.	102
A.15	Fitting results for Sample #18, Spring model.	103
A.16	Stress-strain plots, Soft Constraints model.	104
A.17	Stress-strain plots, St. V-K model.	105
A.18	Stress-strain plots, Spring model.	106
A.19	Fitting results for Sample #04.2.	109
A.20	Fitting results for Sample #12.2.	110
A.21	Fitting results for Sample #14.2.	111
A.22	Fitting results for Sample #18.2.	112
A.23	Comparison of evaluation results for different cloth samples for each fabric.	113

List of Tables

3.1	Cloth and attachment clip masses.	18
3.2	RMS residuals of non-linear orthotropic fitting to Sample #12 with the Soft Constraints model.	31
4.1	Evolution of the position and force RMS error as the strain energy density is augmented with more addends.	44
4.2	Registration error for the knee MRI data	46
5.1	Fitting results for the warp direction on the three samples.	60
5.2	Fitting results for the bending on the three samples.	62
A.1	Results of fitting to Sample-04 (cotton satin).	107
A.2	Results of fitting to Sample-12 (rayon/spandex knit).	107
A.3	Results of fitting to Sample-14 (cotton denim).	108
A.4	Results of fitting to Sample-18 (wool/cotton blend).	108

Chapter 1

Introduction

Deformable objects are a fundamental part of our everyday life and, in many cases, it is thanks to their compliance that they play such an important role: grasping of objects relies heavily on the compliance of our fingers; comfortability of clothing, mattresses, etc, is usually related to their compliance; the automotive industry relies on the compliant, elastic and damping behaviors of the materials used for wheels as well as for the interior of cars.

For decades there has been a huge interest in simulating and analyzing the behavior of deformable objects in a variety of different fields. In ergonomics, simulation of deformable objects is used to design comfortable support structures for the human body. Medicine is another field in which simulation of deformable objects has been studied in depth with the growing interest in virtual surgery simulators. Video game and animation industries rely on deformable object simulation to create compelling visual effects. Fig. 1.1 shows examples of applications where simulation and analysis of deformable objects play a critical role.

Deformation models are key in order to model the behavior of deformable objects. Among the variety of existing deformation models we focus on continuum mechanics models, which produce physically-based accurate results. More specifically, we focus on the Finite Element Method (FEM) as the discretization for this continuum formulation.

Deformable objects are usually characterized by their elastic behavior. However, in the real world more complex effects are involved:

- **Nonlinearities:** popular deformation models in computer graphics produce elastic forces linear to the strain, but real world objects present a nonlinear behavior, with clearly nonlinear elastic forces. These nonlinearities can include both geometric nonlinearities, such as nonlinear relationships between strain and displacement or deformation, and material nonlinearities, such as nonlinear stress-strain laws or material properties that change with the applied loads.




Fig. 1.1 Examples of applications that use deformable objects. Left: ergonomics analysis (MADYMO). Right: virtual surgery simulator (De et al., 2006).

- Anisotropy: isotropic deformable objects present the same behavior regardless of the direction in which they are deformed. However, many real world objects behave differently depending on the deformation direction. For example, cloth usually presents an orthotropic behavior (more stretchable in one direction than in the orthogonal one) produced by differences in the underlying warp/weft yarns and the interlacing pattern.
- Heterogeneity: most real world objects, such as cushions, clothing or human flesh present a spatially varying or heterogeneous behavior, which is not commonly captured by popular deformation models.
- Internal friction: there exist many thread/fiber/yarn-based materials that are usually modeled as continuous objects obtaining accurate results, but there are side-effects that cannot be captured. The clearest example is cloth: interactions between yarns and threads produce internal friction effects that elastic models cannot reproduce.
- Plasticity: while pure elasticity describes the behavior of materials that deform under external loads and return to their original state once the load is removed, plasticity describes the behavior of materials that undergo non-reversible changes in shape (i.e. permanent deformations) due to the applied load.
- Viscosity: real deformable objects usually dissipate energy when a load is applied, hence showing a dependence on time. This is known as viscosity.

In order to realistically capture these behaviors, both appropriate deformation models and accurate model parameters are required.

1.1 Data-Driven Elasticity Estimation

Today’s deformable object simulators for animation, visual effects, engineering and medicine applications can mimic real soft materials and cloth to a high degree of fidelity. However, in order to obtain specific behaviors the user must choose the deformation model that produces the closest behavior to the desired one, and then tune the model parameters. Choosing the best deformation model for a given desired behavior can be challenging due to the large variety of deformation models available, while tuning parameters is a difficult and tedious task.

A possible solution to this problem is to acquire example deformations of real objects and estimate parameters of deformation models that best match the acquired examples. A fundamental part of this data-driven parameter estimation method is the interaction between the input data, the deformation model and the estimation procedure. Complex deformation models with many degrees of freedom are prone to falling into overfitting problems, while too few degrees of freedom may not be able to capture complex behaviors. The structure and the amount of input data is also critical when choosing the deformation model to estimate: sparse input data may only allow estimation of simple deformation models, while dense input data could be used with more complex models.

In this thesis we extend several existing cloth deformation models so that complex cloth behaviors can be captured. Then, we propose a novel more general deformation model that can be applied to cloth as well as volumetric deformable objects. Using the extended cloth models, we estimate material parameters based on measurements taken with a fully automatic capture system, and we test the novel deformation model estimating parameters with a variety of input data.

1.2 Internal Friction Modeling and Estimation

Clothing is a fundamental aspect of our world, hence computer animation research has put a lot of effort towards realism and efficiency in cloth simulation. Since the seminal paper by Terzopoulos et al. (Terzopoulos et al., 1987), a wide variety of deformation models have been proposed to capture the elastic behavior of cloth (Choi and Ko, 2002; Grinspun et al., 2003; Thomaszewski et al., 2009; Volino et al., 2009), sometimes with the addition of complex effects, such as plasticity (Bergou et al., 2007).

Previous works in mechanical engineering and computer graphics (Lahey, 2002; Ngo-Ngoc and Boivin, 2004) point out the existence of significant hysteresis observable in typical force-deformation plots of real cloth, as shown in Fig. 2.1. Researchers refer to internal fric-

tion due to the interaction between yarns and threads as the source of that hysteresis.

Not only is internal friction the source of hysteresis but, as we show in Chapter 5, it also plays a central role in the formation and dynamics of cloth wrinkles: it can induce the formation of ‘preferred’ wrinkles and folds, as shown in Fig. 5.15; persistent deformations, as shown in Fig. 5.17; history-dependent wrinkles, as shown in Fig. 5.14; and it may also make folds and wrinkles settle faster, as shown in Fig. 5.16.

In addition, existing elasticity estimation methods completely ignore internal friction, and either estimate parameters with loading data only, or fit an average of the loading and unloading cycles, leading to bias in the estimations. As we show in Chapter 5, results can be improved significantly by taking into account internal friction and estimating its model parameters.

In this thesis we present an internal friction model for cloth, analyse the visual impact of internal friction in cloth simulation and show an inexpensive measurement setup and a parameter estimation pipeline that produces good estimations for both elastic and internal friction parameters.

1.3 Contributions

The goal of this thesis is to define deformation models that can capture complex behaviors, design data-driven parameter estimation pipelines and estimate model parameters based on real world measurements. We propose deformation models that capture nonlinearities, anisotropy, heterogeneity and internal friction. Then, we estimate parameters and capture nonlinear, anisotropic and heterogeneous behaviors for both volumetric solids and cloth, and show that our internal friction model is able to effectively reproduce the observed hysteretic behavior of real cloth.

More specifically, in the next chapters we will present the following contributions:

1. A cloth deformation model that captures nonlinear, anisotropic behaviors, and a parameter estimation pipeline that produces accurate estimations for the model parameters. We extend existing cloth models using nonlinear strain-dependent parameters by interpolating control points. Then, we estimate the configuration of those control points using real world cloth measurements performed using an automatic capture setup.
2. A novel generic deformation model, based on additive energy components, capable of capturing extreme nonlinearities, anisotropy and heterogeneity applicable to both, volumetric objects and cloth. We use constraints to model the extreme nonlinearities and spatially localized energy-components to model heterogeneity. In addition, we show

how parameter estimation processes benefit from the additive definition of the deformation model in a variety of experiments using real world input data.

3. An internal friction model for cloth that is able to reproduce the hysteresis effect typical in force-deformation measurements for real cloth. We present a reparameterization of Dahl’s friction model that produces key hysteresis features observed in real world cloth. Then, we estimate parameters for both the elastic and the internal friction components using a simple parameter estimation process and data obtained with an inexpensive capture setup.

1.4 Organization

In the next chapter we present previous work on modeling of complex effects and estimation of parameters. The rest of the thesis will focus on describing our contributions. In [Chapter 3](#) we present an extension of existing cloth deformation models capable of capturing nonlinearities and anisotropy, and a parameter estimation pipeline that, together with an automatic measurement system, is capable of producing accurate parameter estimations for several deformation models. [Chapter 4](#) describes a generic deformation model that captures extreme nonlinearities, anisotropy and heterogeneity, that is suitable for parameter estimation based on a wide variety of measurement types. In [Chapter 5](#) we present an internal friction model for cloth, an inexpensive measurement setup and a parameter estimation pipeline that allows us to easily estimate both internal friction and elastic parameters. Finally, [Chapter 6](#) discusses the obtained results and future lines of work.

Chapter 2

Related Work

This chapter presents related work in modeling of deformable objects and measurement-based parameter estimation methods. The key component when simulating deformable objects is the underlying deformation model. Elasticity is the most common effect in deformable objects, but real-world deformable materials present many other sources of complexity: hyperelasticity, hysteresis, plasticity, viscosity, and heterogeneity, which produce characteristic behaviors that cannot be obtained with elasticity alone. In Section 2.1, we focus on related work in deformation models used in computer graphics for several of these effects: hyperelasticity, heterogeneity, internal friction and plasticity. In order to obtain realistic simulations, these models must capture the behavior of real-world objects. A common approach to obtain this realistic behaviors is to estimate the model parameters based on real-world measurements. Section 2.2 presents previous works in measurement-based parameter estimation.

2.1 Modeling of Deformable Objects

Deformable objects include many types of objects: thin shells, such as paper, soft tissue, such as human internal organs or skin, and cloth. All these types of objects can be modeled using a continuum formulation (Eitzmuß et al., 2003; Irving et al., 2007; Lee et al., 2009; Volino et al., 2009), but other methods such as mass-spring systems (Choi and Ko, 2002; Provot, 1995) and inextensibility constraint-based models (English and Bridson, 2008; Goldenthal et al., 2007) are also available. For cloth, due to its yarn-based nature, discrete yarn models (Kaldor et al., 2008) and mesostructure-based continuum models (Boisse et al., 1997; Parsons et al., 2010) have been used.

The goal when simulating these deformable objects is also a critical factor for deciding the type of deformation model to use. In computer graphics, applications such as video-games or

virtual surgery simulators prioritize real-time interactivity over accuracy, and therefore simpler models are preferred. Other applications, such as animation movies or visual effects favor accuracy and use models similar to the ones developed in computational mechanics. In this thesis, we focus on capturing complex effects with the goal of increasing the realism of the final simulation. Consequently, we use continuum-based deformation models and Finite Element Methods ([Bathe, 2006](#); [Hughes, 2000](#)) in order to obtain accurate results.

2.1.1 Modeling of Hyperelasticity

Hyperelastic materials are characterized by a stress-strain relationship derived from a strain energy density function. Linear elasticity is a special case of hyperelasticity, in which this stress-strain relationship is linear. However, linear elasticity is not capable of accurately modeling many real-world elastic behaviors. For example, materials such as rubber or biological soft tissue show nonlinear stress-strain relationships that linear elasticity cannot capture.

Hyperelasticity has been studied in computational mechanics for decades and, due to its complexity, it is still a very active research field. Several constitutive models have been designed to capture the behavior of hyperelastic materials, such as the general Ogden model, or other variants like neo-Hookean or Mooney-Rivlin ([Bonet and Wood, 1997](#); [Ogden, 1997](#)).

In computer graphics, an approach of growing popularity is to model hyperelastic materials by adapting continuum models borrowed from computational mechanics. Then, they are augmented with features to increase robustness and/or efficiency ([Barbič and James, 2005](#); [Bridson et al., 2003](#); [Grinspun et al., 2003](#); [Irving et al., 2004](#); [Müller and Gross, 2004](#); [Narain et al., 2012](#); [Patterson et al., 2012](#)).

Continuum-based approaches can accurately describe the directional variation of material properties, but regardless of the deformation model, a single set of material coefficients for the entire deformation range is not sufficient to faithfully capture the nonlinear response of many real world materials. Bi-phasic models, typically implemented as strain limiting methods ([Bridson et al., 2002](#); [Hernandez et al., 2013](#); [Thomaszewski et al., 2009](#); [Wang et al., 2010](#)), improve on this by splitting the material behavior into an initial, weakly elastic range and a stiff, quasi-inextensible limit. At the extreme, in cloth simulation, where the largest deformations are given by shear and bending components, the pure stretch elastic range can be replaced by inextensibility constraints ([English and Bridson, 2008](#); [Goldenthal et al., 2007](#)).

In order to obtain more realistic simulations, a possible strategy is to estimate the parameters of these models based on real-world measurements. An even better approximation to the true (potentially nonlinear) material response can be obtained by making the material pa-

rameters functions of the deformation, rather than constants, and by fitting these functions to measured data. This approach is described in detail in Chapter 3.

2.1.2 Modeling of Heterogeneity

Many real-world objects, such as human tissue, consist of heterogeneous hyperelastic materials. The complexity of modeling those materials goes beyond the choice of constitutive model, as material heterogeneity requires spatially varying material parameters. This substantially increases the number of model parameters and makes parameter estimation a high-dimensional nonlinear problem.

Virtual surgery is a very active area where heterogeneity is a critical feature. A possible approach to obtain heterogeneous behaviors is to define different objects for structures with different properties, such as bones, muscles and fat, and simulate the whole system handling the interactions between those independent homogeneous structures. However, this method translates heterogeneity modeling complexity into contact handling complexity since the described scenario is extremely demanding in terms of collision detection and response. The most common approach to model heterogeneity is to assign different parameters to each discretization element, as in (Hiller and Lipson, 2012; Lin et al., 2008).

In computer graphics, several works have presented constitutive models capable of capturing heterogeneities. Bickel et al. (2009) captured nonlinearity by interpolating stiffness parameters in strain space, and heterogeneity by defining such strain-space interpolation differently for each discretization element along the material domain. Our work in Chapter 4 follows a similar sampling strategy for heterogeneity, but we capture nonlinearity with an energy model, not stiffness parameters.

Cloth heterogeneity can be analyzed at two different levels. From a continuous domain view, seams are the typical source of heterogeneity. Pabst et al. (2008) modeled seams in cloth by precomputing a multiplicative factor on each discretization element (triangles in this case) based on measurements of real cloth samples. From a yarn-based view, thanks to the recent development of yarn-based models (Lim et al., 2003; Spillmann and Teschner, 2009), heterogeneity can be achieved by assigning different parameters to different sets of interlacing yarns.

2.1.3 Modeling of Internal Friction in Cloth

Two different observations motivate the interest in modeling internal friction in cloth. On one hand, measurements (Eberhardt et al., 1996; Lahey, 2002) have shown that the height of

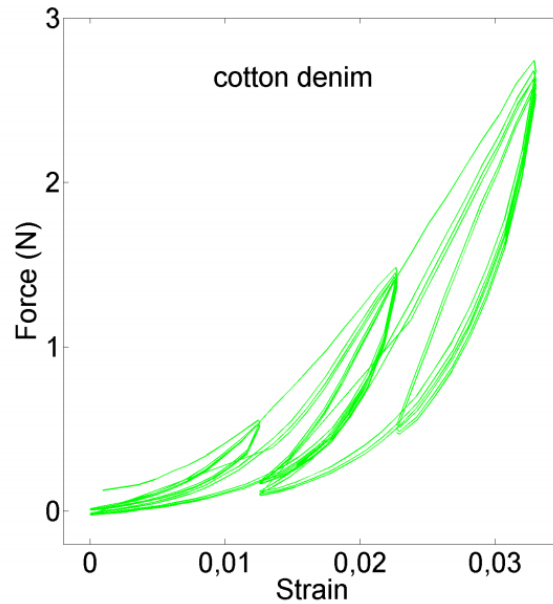


Fig. 2.1 Force-deformation plot for an experiment with real cloth. The height of the hysteresis cycle is almost 50% of the maximum applied force.

hysteresis cycles in force-deformation plots of experiments using real cloth can reach values of almost 50% of the maximum applied force (Fig. 2.1). On the other hand, different weaving patterns produce significantly different behaviors for the same yarn materials, which leads to the conclusion that the interactions between the yarns, which is the source of internal friction, have a huge effect on the cloth's overall behavior.

Friction is usually an external effect, taking place when two different objects interact. Cloth internal friction captures the frictional interactions that take place between interlacing yarns, and models small-scale resistance to inter-yarn motion. When cloth transitions from loading to unloading or vice versa, yarns realign and produce a force that acts against their relative motion. Unlike elastic forces, which oppose deformation (i.e., strain), friction forces tend to act against the change in deformation (i.e., strain rate). Friction and hysteresis have been long studied in mechanical engineering, and there is a large variety of available models. Padthe et al. (2008) survey several models, discuss their mathematical foundations, and analyze their effects.

In computer graphics, probably the most popular friction model is Coulomb's model. It states that dissipation should be maximized (i.e., strain rate should be minimized), subject to some constraint on the friction stress. For object-object contact, the constraint is posed as a relationship between tangential and normal forces (Baraff, 1991). Coulomb's model is also generalized to 3D, e.g., for the simulation of granular materials (Zhu and Bridson,




Fig. 2.2 Persistent deformations and hysteresis under friction and plasticity models. Left: hysteresis loops added to a linear-elastic model; Center: creep model with plasticity saturation (Müller and Gross, 2004); Right: bounded elasticity model (Bergou et al., 2007).

2005), by posing constraints on a 3D friction stress tensor. Coulomb’s model succeeds in capturing hysteresis effects, but it may undergo an instantaneous switch of force directions, not present in the data observed on real cloth. Sherburn (2007) simulated cloth at the yarn level using Coulomb’s model for inter-yarn frictional contact. Later, Kaldor et al. (2008) modeled inter-yarn frictional contact combining damping and velocity filters. More recently, Chen et al. (2013) have modeled and estimated friction effects between cloth and other deformable objects.

2.1.4 Modeling of Plasticity

Plastic materials are those that, in contrast to elastic materials, do not return to their original shape once the loading forces have been removed. This effect is similar to the persistent deformations obtained with internal friction, and it is therefore reasonable to consider plasticity models as an option to model hysteresis.

As with many other effects characteristic of deformable objects, plasticity has been thoroughly studied in mechanics. Good reviews of existing models and analysis methods can be found in (Han and Reddy, 2012; Simo and Hughes, 2000).

In computer graphics, Terzopoulos and Fleischer (1988) modeled plasticity, viscoelasticity and fracture based on models developed in mechanics. In the context of cloth simulation, Kim et al. (2011) designed a complex model of plasticity with ten parameters to produce persistent cloth wrinkles. Their model modifies the stiffness and rest angle of a bending spring when its strain rate exceeds a threshold value. The recent work of Narain et al. (2013) focuses on resolving the geometric detail in folded and crumpled sheets under plastic deformation.

In general, common plasticity models in computer graphics separate elastic strain (which defines elastic energy) from plastic strain (which does not change elastic energy). Two common models are a creep model that saturates to a maximum plastic strain (Müller and Gross, 2004) and a bounded elasticity model (Bergou et al., 2007), as shown in Fig. 2.2-center and

Fig. 2.2-right respectively. The curves evidence that plasticity models are not a good match for the hysteresis observed in cloth. A plastic material deforms elastically until the material's yield strength is exceeded, then abruptly gives way and begins deforming irreversibly. By contrast, hysteresis is prominent in cloth even for small deformations, and its effects increase gradually as deformation is increased. Moreover, even in the plastic regime, the response of plasticity models to reversed loads is purely elastic, whereas cloth exhibits local hysteresis.

Another important difference between plasticity and hysteresis is that persistent plastic deformations require a large reverse load to be undone, whereas persistent deformations due to hysteresis can be undone simply by applying small loads that produce narrow hysteresis.

In this discussion, we did not consider material hardening, but it would not eliminate the fundamental differences. Our model could be complemented with plasticity to capture irreversible wrinkles due to extreme deformations, but based on all previous observations we have discarded plasticity for modeling hysteresis in cloth.

2.2 Measurement-Based Parameter Estimation

Material parameter estimation is critical in the characterization of real-world deformable objects, specially biological soft tissues. However, it is a difficult task that involves designing the capture and measurement system as well as deciding the deformation model and estimation methods to use. Both tasks are tightly coupled since the density and structure of the input data will impose certain limitations on the estimation algorithms as well as the deformation models.

The choice of estimation methods is strongly affected by the dimensionality of the problem. Capturing nonlinearities may already require more than 30 parameters, as in (Wang et al., 2011b). If heterogeneity is included, the dimensionality explosion prevents the utilization of global optimization algorithms, due to its computational cost. Instead, local optimization algorithms are used and local minima problems arise. In order to minimize their impact, an appropriate design of the capture process and the input data is fundamental.

2.2.1 Parameter Estimation for Hyperelastic Heterogeneous Objects

In computer graphics, several previous works have attempted the estimation of material parameters of deformable objects. The pioneering work of Pai et al. (2001) introduced a system to capture and estimate shape, elasticity, and surface roughness by scanning a volumetric object. Becker and Teschner (2007) presented a method to estimate Young modulus and Poisson

ratio in a linear FEM formulation by solving a linear least squares problem. More recently, Bickel et al. (2009) proposed a data-driven model for which parameters are estimated based on a set of real-world example deformations.

For cloth, the traditional method to capture nonlinear elasticity has relied on the Kawabata Evaluation System (KES) (Kawabata, 1980). This system uses complex machinery to exert controllable uniform strain on cloth samples, measuring stretch, shear and bending, and then estimate parameters of elastic models (Breen et al., 1994; Eberhardt et al., 1996; Volino et al., 2009). However, despite the complexity of the capture setup, KES suffers a major shortcoming: by exciting cloth with uniform strain it fails to capture the interplay between different deformation modes.

In more specific contexts other devices have been used, such as the Picture Frame test (Culpin, 1979) for measuring shear properties and the Cantilever test (Clapp et al., 1990) for measuring bending properties (see also Pabst et al. (2008)).

Alternatively, recent approaches exploit computer vision to capture arbitrary cloth deformations and then estimate model parameters. Bhat et al. (2003) (and recently Kunitomo et al. (2010)) tried to extract model parameters from casually captured videos, avoiding the need for controlled conditions and using a simple and inexpensive acquisition process. More recently, Bouman et al. (2013) estimate cloth stiffness through the temporal analysis of texture patterns in video. These methods present the drawback of not being able to accurately separate internal (i.e. material-specific) and external (e.g. friction, air drag) parameters.

In a similar spirit, capture technology can be used to record time-varying geometry of complex cloth motions (Bradley et al., 2008b; Stoll et al., 2010; White et al., 2007). But while capturing can provide accurate deformation data, parameter fitting remains very difficult without explicit control over boundary conditions, in particular loading forces.

The work of Wang et al. (2011b) is particularly interesting due to the high estimation quality combined with simple acquisition procedures. Wang et al. propose a data-driven piecewise linear elastic cloth model comprising 39 material parameters, and advocate for a combination of semi-controlled deformations and computer-vision-based tracking to estimate stiffness parameters that depend on the value and direction of the principal membrane strain. These parameters are fitted to experimentally acquired data obtained from planar and bending deformations.

The closest research to our work is that by Wang et al. (2011b). While their capture setup is appealingly simple, in Chapter 3 we present a more general and powerful one: it produces a 3D surface, rather than a 2D deformation, and it measures all forces applied to the cloth as they change during a range of different deformations, including hysteresis as the

difference between loading and unloading paths. In addition, we apply our estimation pipeline with several popular deformation models and compare the obtained estimations, exposing the advantages and disadvantages of each deformation model and their suitability for realistic cloth modeling.

2.2.2 Parameter Estimation for Internal Friction in Cloth

Estimation of internal friction in cloth is important not only because it is the source of very characteristic behaviors, such as hysteresis, but also because ignoring it may lead to bias in the estimation of elasticity parameters (Volino et al., 2009; Wang et al., 2011b). Even if the elastic parameters are chosen to fit the average of loading and unloading behaviors, given observed hysteresis as high as 50% of the average force, ignoring internal friction may induce deformation errors of up to $\pm 25\%$ for a given load.

The works by Lahey (2002) and Ngo-Ngoc and Boivin (2004) modeled cloth hysteresis using internal friction, captured force-deformation data using KES and estimated (strain-independent) parameters of standard friction models. They considered a second-order Bliman-Sorine model (Bliman and Sorine, 1991), which can capture hysteresis, pre-sliding, and the Stribeck effect, i.e., a slight decrease in friction at the transition from static to dynamic regime. However, in their data they found only very subtle Stribeck effects, without which a first-order Bliman-Sorine model, equivalent to a simple Dahl model (Dahl, 1968) is sufficient.

Like Lahey (2002) and Ngo-Ngoc and Boivin (2004), we fit a friction model to deformation data. We use Dahl’s friction model, which produces a good fit to our measured force-deformation data. However, we have found that the standard strain-independent parameterization of Dahl’s model fails to capture the strain-dependent magnitude of hysteresis, which is present in our data. Therefore we reparameterize Dahl’s model to account for this effect. Another key difference, is that the KES data used in those papers provides dense sampling under uniform strain conditions, whereas our methods work with sparse sampling and non-uniform strain, enabling simpler setups. Furthermore, prior work on internal friction for cloth is limited to the estimation of parameters that fit hysteresis cycles well, while we also analyze the impact of internal friction on animations. Chapter 5 presents a detailed description of the internal friction model and the parameter estimation pipeline.

Chapter 3

Data-Driven Estimation of Cloth Simulation Models

Realistic simulation of cloth relies not only on the correct underlying deformation model but also on a set of parameters that produce realistic results. However, traditional methods to find adequate parameters for cloth models are time-consuming and require a trial-and-error process until the obtained behavior is considered close enough to the desired one. In this chapter, we aim to solve this problem by introducing new techniques to measure complete cloth behavior under controlled conditions and to estimate cloth deformation models from these measurements.

In this chapter we will present:

- A new, general system for observing cloth properties that measures more complete data than previous work in cloth capture or textile testing.
- A new method for fitting parametric models to this type of data.
- Results that illustrate the performance of several widely used cloth models in our estimation system.

The results of this work have been published in Computer Graphics Forum (Proceedings of Eurographics 2012) ([Miguel et al., 2012](#)).

3.1 Introduction

Most methods for testing cloth move the sample into a state of near-uniform strain, exercising one or at most two components of strain at once: pure stretching, pure shearing, or pure

bending. One or two forces are measured to quantify the cloth’s resistance to deformation, and the resulting force-displacement curves are valuable in studying the differences between materials. However, this approach has certain limitations. The inevitable deviations from uniform strain create modeling error that cannot be quantified without knowing the actual strain variation; and force-displacement curves can be used directly to tune a cloth model, but do not provide any way to validate the resulting fit.

Our measurement system applies forces to a sample of cloth using actuators and force sensors that let us know the complete applied force, in 3D. The resulting deformation is tracked by a stereo computer vision system that captures the complete deformation, also in 3D. Having deformation and force information makes our data well suited to model validation—the experiment measures the complete answer that should be predicted by a cloth simulator. Also, we do not need uniform strain, and in this chapter we illustrate a range of tests, some that mimic traditional tests and some with more complex deformations.

Our approach to model estimation is to numerically optimize nonlinear stress-strain curves to minimize errors in force and position compared to the measurement. We have designed a general fitting method, suited for the vast majority of existing cloth models, that leverages equilibrium conditions to guide the iteration. By estimating model parameters under a sequence of deformations of increasing complexity, we alleviate convergence problems in the presence of abundant local minima.

We have used our system to fit three membrane models and two bending models from the graphics literature, each based on a different strain measure, and to evaluate the resulting models against more complex motions.

3.2 Measurement system

The design goals of our measurement system are to create deformations in a sample of cloth that explore a substantial range of the material’s strain space, and to record complete information about the forces applied to the cloth and the deformation that it undergoes. Like other cloth testing systems, we focus primarily on tensile forces, because it is hard to repeatably produce and measure compression forces in a sheet that is inclined to buckle.

Tests are performed on 100 mm square cloth samples using two kinds of plastic clips: small, rounded clips that grab a localized area, and long clips that grip one whole side of the sample. We measure the weights of all cloth samples as well as the clips (see Table 3.1) and use these values in the optimization process. Forces are applied to the clips by fine wire cords that are pulled to defined displacements by eight linear actuators, and the tension in the cords is




Fig. 3.1 Acquisition setup for the measurement system.

monitored by miniature load cells located at the actuator ends (see Fig. 3.1). Our actuators and load cells are capable of applying and measuring tensions up to 45 N, but in our experiments the maximum force is typically on the order of 10 N.

The geometry of the cloth sample and the attached clips is monitored by a vision system composed of four high-resolution cameras. The location and orientation of the cords attached to the clips (which reveal the direction of the applied force) are also tracked. Each output frame of a measurement session contains:

- The configuration of the cloth sample, represented as a deformed mesh with 10K regularly sampled vertices.
- The positions and orientations of all clips attached to the cloth, including a list of clamped cloth vertices.
- The forces applied to all clips. The magnitudes are determined by the tension measure-

Cloth Sample	Id	Mass (g)
cotton satin	#4	1.2
rayon/spandex knit	#12	3.1
cotton denim	#14	4.6
wool/cotton blend	#18	2.4
plastic clips (3 sizes)		1.9, 10.1, 13.3

Table 3.1 Cloth and attachment clip masses.

ments, and the directions are determined by the observed directions of the cords.

Note that the actuator positions themselves are not part of the output, since they are superseded by the displacements measured at the clips. This prevents stretching of the cord, or other factors altering the distance between the clip and the actuator, from affecting displacement accuracy.

3.2.1 Reconstruction

Our vision system recovers the space-time geometry of the deforming cloth and attached rigid clips, as well as the directions of the forces applied to the clips.

Initialization. The cloth sample starts flat on a table and we capture the rest pose without applied tensile forces. This initial frame serves to compute the geometry of the cloth without any occlusion from clips. We then attach the clips, and the measurement process continues automatically, following a defined script of actuations, and recording images and forces. We typically deform the cloth by moving the actuators at 0.5 mm/sec, and every 2 seconds we allow the sample to reach equilibrium and capture a static frame.

Cloth Geometry Reconstruction. The raw data for a single deformation consists of 20 to 200 individual measurement frames, with a set of camera images and simultaneous force sensor readings for each frame.

We compute the per-frame geometry using a state-of-the-art stereo reconstruction technique (Bradley et al., 2008a), which was specifically tailored for reconstructing cloth geometry (Bradley et al., 2008b). If the inherent texture of the cloth is not sufficiently random, it is printed with a wavelet noise pattern (Atcheson et al., 2008) to provide texture that can be used




Fig. 3.2 Force measurements for selected frames of a corner pulling sequence. Forces are rendered as red vectors with magnitudes proportional to their values (in Newtons).

for stereo reconstruction and tracking. The pattern is printed with a flatbed inkjet printer and does not have a noticeable effect on the material behavior.

To represent inter-frame correspondence, we use optical flow to obtain a single triangle mesh that deforms over time, akin to the human face tracking method of Bradley et al. (2010). To start, the cloth vertices in the rest pose frame (frame 0) are projected onto the input images, where optical flow predicts the projection of each vertex at the next time step. Back-projecting onto the reconstructed geometry for the next frame gives new position estimates for the cloth vertices. The process is then repeated using the result from frame n to obtain frame $n + 1$. As with all sequential tracking methods, very small errors can accumulate over time and cause temporal drift in the reconstruction. To avoid drift, we subsequently match each frame independently back to the rest pose frame using the approach described in Bradley et al. (2010). The final solution is smoothed using Laplacian regularization to remove noise.

Tracking Clips and Cords. In order to measure the complete answer that a simulator should predict, we need to determine the interaction between the rigid clips, the cloth, and the cords. The clips are produced, using rapid prototyping, with embedded codes (Fiala, 2005) that allow us to determine their identity, position, and orientation automatically. The area of cloth occluded by the clips is used to automatically determine which cloth vertices are clamped by each clip and will therefore be constrained to it in the simulator.

The vision system also triangulates a 3D line for each cord in the images. A few user scribbles on an input image indicate which cords are affecting each clip. Fig. 3.2 illustrates the force measurements and clip locations for three different frames from one experiment. The forces are rendered as red vectors with lengths proportional to the force magnitudes.

3.2.2 Measurements

The set of deformations to measure is motivated by the goals of the parameter fitting stage (Section 3.4): to fit model parameters for stretch, shear and bending that best describe the cloth, and to validate the parameter fits by comparing against other measurements.

To reduce the risk of falling into local minima during parameter fits, we have designed deformation sequences that produce near-isolated strains, and allow estimating stretch, shear and bending properties in a separate and incremental manner. However, unlike standard textile evaluation practices (Kawabata, 1980), and thanks to our full 3D deformation capture solution, we relax the requirement of uniform strains.

To isolate stretching we perform a uni-axial tension experiment, with forces applied to two long bar clips attached to either side of the cloth (see Fig. 3.3, 2nd column). The cloth is slowly stretched until a maximum force is reached and then slowly released back. The process is repeated three times, in both weft and warp directions separately.

Shearing is captured using an approximate *picture-frame* experiment (Culpin, 1979), where four long clips fix the cloth boundaries and shear stress is applied as the cords pull on opposite corners (Fig. 3.3, 3rd column). To isolate bending deformation we slowly push the flat cloth sample off the edge of a table and measure its shape as it bends under its own weight (Fig. 3.3, 4th column), for both weft and warp directions. Thus we have a total of five measurements per cloth sample that will be used for parameter fitting (two stretch, one shear, and two bending).

We also capture two sequences with more complex deformation (Fig. 3.4) for validation after parameter fitting. In the first test, opposite edges of the cloth are pulled in opposite directions, causing shearing and buckling (Fig. 3.4, top). The second is a four-corner pulling test, where opposite pairs of corners are pulled in alternation, resulting in diagonal wrinkles (Fig. 3.4, bottom).

Fig. 3.3 and Fig. 3.4 show that our acquisition system is able to recover the 3D cloth geometry including temporal tracking (illustrated with an overlaid checkerboard), tracked 3D clip locations, and individual 3D force directions (shown as green lines). To our knowledge, our method presents the first system able to record such extensive information about the behavior of a cloth sample.

3.2.3 Accuracy

In the vision system, the camera calibration accuracy is within 0.3 pixels, or about 0.075 millimeters at the distance of the cloth. The multi-view stereo algorithm of Bradley et al. (2008a) is among the most accurate available according to the Middlebury evaluation benchmark. It is




Fig. 3.3 Selected frames from isolated measurements of stretching, shearing, and bending. The left column shows the cloth in its rest state. One input image is shown above each 3D reconstruction. The reconstruction includes parameterized cloth geometry, clip locations and the direction of the force vectors (shown as green lines).

difficult to quantify the accuracy of the temporal flow computation, but it can be visualized by compositing the reconstructed deformation on top of the input images.

The raw repeatability of our force sensors is about 3 millinewtons (RMS). The largest source of error in measuring the force indirectly through the cord is the internal friction in the cord as it bends around the pulleys, which introduces an artificial hysteresis of about 0.1 N.

3.3 Cloth Models

Our goal is to study the fidelity of constitutive models of cloth—models that predict the forces produced in the cloth in response to deformations. The input of such a model is the positions of the vertices $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}^3$ that define the deformation state of the sheet (analogous to strain in continuum mechanics) and the output is the forces that act between those vertices in response (analogous to stress). Although some of the models we look at are discrete in nature, we will use the convenient terms stress and strain to describe them.

3.3.1 Strain Metrics

Most elastic cloth models separate membrane (i.e., stretch and shear) and bending deformation energies. In both cases, deformation energy density can be described by the product of strain ($\boldsymbol{\varepsilon}$) and stress ($\boldsymbol{\sigma}$), i.e., $W = \frac{1}{2} \boldsymbol{\sigma} \cdot \boldsymbol{\varepsilon}$. Furthermore, most of these models define separable scalar




Fig. 3.4 Selected frames from more elaborate cloth manipulation, demonstrating complex deformations. Input image (left) and 3D reconstruction shown from two viewpoints.

stress components as linear functions of individual scalar strain metrics. In that case, the energy density of each deformation component i can be written as $W_i = \frac{1}{2}k_i\varepsilon_i^2$, where $k_i\varepsilon_i = \sigma_i$ and k_i is the stiffness coefficient corresponding to the deformation component ε_i . The force density due to each ε_i follows as $\mathbf{F}_i = -\nabla W_i = -\sigma_i\nabla\varepsilon_i = -k_i\varepsilon_i\nabla\varepsilon_i$. We have evaluated three models for membrane deformation that fit this description (spring systems, the soft constraint model by Baraff and Witkin (1998) and the diagonalized St.Venant-Kirchhoff (StVK) model by Volino et al. (2009)), and two bending models (spring systems and the edge-based bending model in Discrete Shells (Grinspun et al., 2003)).

Considering possible anisotropic behavior, we distinguish six different strain components on regularly triangulated cloth: weft-stretch ($\varepsilon_{s,u}$), warp-stretch ($\varepsilon_{s,v}$), shear ($\varepsilon_{s,uv}$), weft-bend ($\varepsilon_{b,u}$), warp-bend ($\varepsilon_{b,v}$), and diagonal-bend ($\varepsilon_{b,uv}$). Next, we describe in detail the strain metrics for the individual deformation components in the selected models. Note that not all force models define the quantities below explicitly as strains, as they often rely on the resolution of the discretization, or they differ simply by scale factors that can be embedded in the stiffness k_i . We use continuum strain definitions in all cases to fit them in a common formulation that allows us to easily compare the models. The details about the computation of strain gradients for the various cloth models, needed in the force computation, are given in their respective original papers.

Spring Membrane. All deformation components are modeled based on springs, with weft and warp ring-1 springs for stretch, and diagonal ring-1 springs for shear. The spring-based strain for each deformation component i can be defined as the relative change of edge length,

$$\varepsilon_{s,i} = l/l_0 - 1, \quad (3.1)$$

with l the current length of the spring, and l_0 its rest length.

Diagonalized StVK. The membrane deformation is defined using the Green-Lagrange strain tensor, a formulation introduced to computer graphics by Terzopoulos et al. (1987). Given a per-triangle mapping function \mathbf{w} from the undeformed 2D configuration $(\mathbf{x}_{a,0}, \mathbf{x}_{b,0}, \mathbf{x}_{c,0})$ to the deformed 3D configuration $(\mathbf{x}_a, \mathbf{x}_b, \mathbf{x}_c)$, the deformation gradient can be computed as

$$(\mathbf{w}_u \ \mathbf{w}_v) = (\mathbf{x}_b - \mathbf{x}_a \ \mathbf{x}_c - \mathbf{x}_a) (\mathbf{x}_{b,0} - \mathbf{x}_{a,0} \ \mathbf{x}_{c,0} - \mathbf{x}_{a,0})^{-1}. \quad (3.2)$$

Then, the components of the membrane Green-Lagrange strain are defined as:

$$\varepsilon_{s,u} = \frac{\|\mathbf{w}_u\|^2 - 1}{2}, \quad \varepsilon_{s,v} = \frac{\|\mathbf{w}_v\|^2 - 1}{2}, \quad \varepsilon_{s,uv} = \mathbf{w}_u^T \mathbf{w}_v. \quad (3.3)$$

Volino et al. (2009) approximate the standard StVK model zeroing out off-diagonal terms in the matrix that relates strain and stress, $\boldsymbol{\sigma} = \mathbf{E}\boldsymbol{\varepsilon}$. Then, in the diagonalized StVK, each membrane stress component depends only on its corresponding strain component, $\sigma_{s,i}(\varepsilon_{s,i})$.

Soft Constraints. Weft- and warp-stretch are measured through a subtle modification of the Green-Lagrange strain tensor, defining terms that are quadratic in positions instead of quartic:

$$\varepsilon_{s,u} = \|\mathbf{w}_u\| - 1, \quad \varepsilon_{s,v} = \|\mathbf{w}_v\| - 1, \quad \varepsilon_{s,uv} = \mathbf{w}_u^T \mathbf{w}_v. \quad (3.4)$$

Spring Bending. The deformation is measured based on weft and warp ring-2 springs for weft- and warp-bend, and diagonal ring-2 springs for diagonal-bend. Same as for membrane deformation, strain is measured as the relative change of edge length (Eq. 3.1).

Discrete Shells Bending. It is measured as the change of dihedral angle at edges:

$$\varepsilon_{b,i} = 1/h_0(\theta - \theta_0), \quad (3.5)$$

with θ the current angle, and θ_0 the undeformed angle. Grinspun et al. (2003) and Bridson et al. (Bridson et al., 2003) discovered concurrently the appropriate weighting of the angle change in order to model homogeneous bending on irregular triangle meshes with a homogeneous stiffness. Grinspun et al. define h_0 as a third of the average of the heights of the two triangles incident to the edge. This definition implies that bending energy density is integrated over edge-centered rectangles of size $l_0 \times h_0$. With our separation of weft-, warp- and diagonal-bending to capture anisotropy, the bending models in Discrete Shells and by Baraff and Witkin (1998) are equivalent up to a stiffness scale factor.

3.3.2 Nonlinear Stress-Strain Curves

The generic force density model $\mathbf{F} = -\boldsymbol{\sigma}\nabla\boldsymbol{\varepsilon}$ defined above assumes a linear stress-strain curve $\boldsymbol{\sigma} = k\boldsymbol{\varepsilon}$. However, stress-strain curves are potentially nonlinear functions. Then, for each deformation component, we model stress as a function $\boldsymbol{\sigma}_i = k_i(\boldsymbol{\varepsilon}_i)\boldsymbol{\varepsilon}_i$, with a strain-dependent stiffness k_i encoded using Hermite splines. We enforce non-negative constraints on the stiffness values at control points. The resulting nonlinear force density function, $\mathbf{F}_i = -k_i(\boldsymbol{\varepsilon}_i)\boldsymbol{\varepsilon}_i\nabla\boldsymbol{\varepsilon}_i$ yields a conservative force field, but note that the elastic energy density can no longer be defined simply as $\frac{1}{2}k\boldsymbol{\varepsilon}^2$, and would now require the integration of the stiffness function.

Although only Volino et al. (2009) propose a general nonlinear stress-strain relationship (though many systems use some form of strain limiting instead), the same construction can easily be built on any of our selected models. Because linear models fit the data poorly, we used the nonlinear model in all cases, resulting in a consistent set of models, parameterized by the number of spline control points, which reduces to the widely used linear models when each spline has a single control point.

3.4 Fitting

The key question of how well a given model describes a particular piece of cloth is answered by fitting the model to the measurement data: adjusting its parameters to minimize the difference between the model’s predictions and the measured behavior, both in position and force. We do this by solving an optimization problem, leveraging that the cloth is at static equilibrium at the measured configurations.

In principle all parameters of a cloth model can be fit to a sufficiently rich single deformation sequence, but this can result in a problem fraught with local minima. In order to achieve stable fits, we have designed an incremental optimization procedure that fits model parameters

a few at a time using the isolated deformations described in Section 3.2.2.

3.4.1 Optimization Problem

For each different cloth sample, we have created a simulated replica with the same mass, uniformly distributed, and the same 100mm square geometry, discretized with a regular 25×25 -node mesh, connected either with springs or with quadrilaterals split into triangles, depending on the model. In each measurement sequence, a different set of nodes is fixed to rigid bodies representing the clips. For the bending measurement sequences (see Fig. 3.3), we fix all cloth nodes above the edge of the table. The measured pulling forces of the cords are applied as point forces on the rigid bodies at known locations, with known magnitudes and orientations.

Given a set of captured static deformation frames, we wish to know the (nonlinear) stress-strain curves for the deformation components of a cloth model, such that a simulated cloth matches known positions and forces as well as possible. Specifically, we minimize the weighted error of cloth positions and clip forces over a sequence of measurement frames, subject to the constraint of static equilibrium on all frames. For the formulation of the objective function, we concatenate in vectors the positions, \mathbf{x}_n , and the net forces, \mathbf{F}_n , of free cloth nodes at all frames, as well as the forces, \mathbf{F}_c , applied by the cords on the clips. Due to equilibrium, the net force on the clips, produced by cord forces, gravity, and forces from fixed cloth nodes, must be zero. We indicate with $\tilde{\mathbf{x}}_n$ and $\tilde{\mathbf{F}}_c$, respectively, the known cloth node positions and clip forces, measured as described in Section 3.2.

We also concatenate in a vector \mathbf{k} the (unknown) stiffness values at the control points of the nonlinear stress-strain curves for the deformation components of the cloth. Since the pieces of cloth are homogeneous, we use a single curve for each deformation component for all frames and all cloth elements. Then, the computation of model parameters based on the minimization of position and force errors subject to the static equilibrium condition can be formulated as the following nonlinear constrained least-squares problem:

$$\begin{aligned} \mathbf{k} &= \arg \min \mu \|\mathbf{x}_n(\mathbf{k}) - \tilde{\mathbf{x}}_n\|^2 + \lambda \|\mathbf{F}_c(\mathbf{x}_n, \mathbf{k}) - \tilde{\mathbf{F}}_c\|^2, \\ \text{s.t. } \mathbf{F}_n(\mathbf{x}_n, \mathbf{k}) &= 0. \end{aligned} \quad (3.6)$$

In this optimization problem, we use the measured clip positions, $\tilde{\mathbf{x}}_c$, as known boundary conditions. For stretch tests, the objective function is based only on clip forces, i.e., $\mu = 0, \lambda = 1$, while for bend tests it is based only on cloth positions (since there are no measured forces), i.e., $\mu = 1, \lambda = 0$. For shear tests, the objective function is based only on clip forces parallel to the direction of the clips themselves. We observed that, in situations of near-homogeneous

shear, the clip-parallel forces are dominated by shear, while clip-orthogonal forces are dominated by stretch. Then, by fitting only clip-parallel forces we reduce the sensitivity to potential errors in stretch stiffness.

The optimization problem contains two unknowns: the parameter vector \mathbf{k} and cloth node positions \mathbf{x}_n . We solve the optimization in an iterative manner, refining \mathbf{k} and \mathbf{x}_n separately on two nested loops. In an outer loop, we refine \mathbf{k} by local minimization of the error function and, in an inner loop, we recompute \mathbf{x}_n to satisfy the equilibrium constraint.

Outer Loop. Assuming cloth positions that satisfy the (nonlinear) equilibrium constraints on all captured frames, we locally linearize those constraints w.r.t. both \mathbf{k} and \mathbf{x}_n . As a result, we obtain a linear expression that relates node positions to parameter values:

$$\frac{\partial \mathbf{F}_n}{\partial \mathbf{x}_n} \Delta \mathbf{x}_n + \frac{\partial \mathbf{F}_n}{\partial \mathbf{k}} \Delta \mathbf{k} = 0 \Rightarrow \Delta \mathbf{x}_n = -\frac{\partial \mathbf{F}_n^{-1}}{\partial \mathbf{x}_n} \frac{\partial \mathbf{F}_n}{\partial \mathbf{k}} \Delta \mathbf{k}. \quad (3.7)$$

We also locally linearize clip forces,

$$\Delta \mathbf{F}_c = \frac{\partial \mathbf{F}_c}{\partial \mathbf{x}_n} \Delta \mathbf{x}_n + \frac{\partial \mathbf{F}_c}{\partial \mathbf{k}} \Delta \mathbf{k}, \quad (3.8)$$

and we turn Eq. 3.6 into a linear least squares problem, which we solve to refine the parameters $\mathbf{k}(i+1) = \mathbf{k}(i) + \Delta \mathbf{k}$, with:

$$\begin{aligned} \Delta \mathbf{k} = \arg \min_{\mu} & \left\| \mathbf{x}_n(i) - \tilde{\mathbf{x}}_n - \frac{\partial \mathbf{F}_n^{-1}}{\partial \mathbf{x}_n} \frac{\partial \mathbf{F}_n}{\partial \mathbf{k}} \Delta \mathbf{k} \right\|^2 + \\ & \lambda \left\| \mathbf{F}_c(i) - \tilde{\mathbf{F}}_c + \left(\frac{\partial \mathbf{F}_c}{\partial \mathbf{k}} - \frac{\partial \mathbf{F}_c}{\partial \mathbf{x}_n} \frac{\partial \mathbf{F}_n^{-1}}{\partial \mathbf{x}_n} \frac{\partial \mathbf{F}_n}{\partial \mathbf{k}} \right) \Delta \mathbf{k} \right\|^2. \end{aligned} \quad (3.9)$$

We terminate the outer loop (and hence the overall optimization) when the residual is reduced by less than 1% between two consecutive iterations. To ensure convergence of the Newton-like iterations and to enforce non-negativity constraints on the components of \mathbf{k} , we execute a line search from $\mathbf{k}(i)$ to the solution of Eq. 3.9 if the residual grows or if the solution violates some constraint. The solution to the linear least squares problem requires solving a system $\mathbf{A}\mathbf{k} = \mathbf{b}$, where the size of \mathbf{A} is given by the number of unknown stiffness values, $|\mathbf{k}|$. In our test examples, this number was always below 10, and we solved the linear systems using LDL factorization. The formulation of \mathbf{A} , on the other hand, requires solving $|\mathbf{k}|$ linear systems of type $\frac{\partial \mathbf{F}_n}{\partial \mathbf{x}_n} \mathbf{y} = \mathbf{b}$, which we did using the conjugate gradient method.

Inner Loop. Once the parameter values $\mathbf{k}(i+1)$ are refined, we bring the cloth to a static equilibrium position, $\mathbf{x}_n(i+1)$. We do this by solving quasi-static simulations until convergence on all captured frames, starting always from the measured configuration \mathbf{x}_n and using the measured clip positions $\tilde{\mathbf{x}}_c$ as boundary conditions. We consider that a piece of cloth has converged to equilibrium when $\|\mathbf{F}_n\| < 10\mu\text{N}$. The quasi-static simulations involve linear-system solves with the cloth stiffness matrix $\frac{\partial \mathbf{F}_n}{\partial \mathbf{x}_n}$. We found that, during intermediate iterations, the stiffness matrix may not always be well conditioned, therefore we have solved the quasi-static equilibrium problems using additive Levenberg-Marquardt, which effectively produces a modified stiffness matrix of the form $\frac{\partial \mathbf{F}_n}{\partial \mathbf{x}_n} + \mu \mathbf{I}$. For improved conditioning, we also use this modified stiffness matrix in the outer loop.

3.4.2 Incremental Parameter Fitting

The nonlinearity of cloth deformation, together with the complex interplay of various deformation components in the resulting forces and positions, make the optimization problem above extremely complex in the general case, prone to falling in local minima and sensitive to initialization values. However, we largely alleviate these issues with the design of the five isolated deformation measurements described in Section 3.2.2, which allow us to separately fit stiffness curves for the six deformation components described in Section 3.3.1, following an incremental parameter fitting procedure.

First, we fit in parallel the weft-stretch stiffness curve, $k_{s,u}(\epsilon_{s,u})$, for the weft-stretch sequence, and the warp-stretch stiffness, $k_{s,v}(\epsilon_{s,v})$, for the warp-stretch sequence. We ignore shear and bend parameters for stretch fits, as we have observed that they have little effect. Second, using known stretch stiffness curves, we fit the shear stiffness $k_{s,uv}(\epsilon_{s,uv})$, for the shear sequence. Third, we fit in parallel the weft-bending stiffness $k_{b,u}(\epsilon_{b,u})$, for the weft-bending measurement sequence, and the warp-bending stiffness $k_{b,v}(\epsilon_{b,v})$, for the warp-bending sequence. Finally, we fit the diagonal-bending stiffness curve $k_{b,uv}(\epsilon_{b,uv})$, using both weft- and warp-bending measurements. To better account for cross-influence of shear and bending, we use their estimated values as initial guesses and run another fitting iteration.

To fit each stiffness curve $k_i(\epsilon_i)$, we iteratively subdivide the Hermite spline adding more control points until the residual error function Eq. 3.6 is reduced by less than 1% or a specified maximum number of points, usually 4 or 5, is reached. First, we evaluate the strain histogram for the corresponding measurement sequence, and we determine maximum and minimum strains after removing outliers. We initialize the stiffness curve with one control point (i.e., constant stiffness), and subsequently we subdivide the strain range with equidistant control

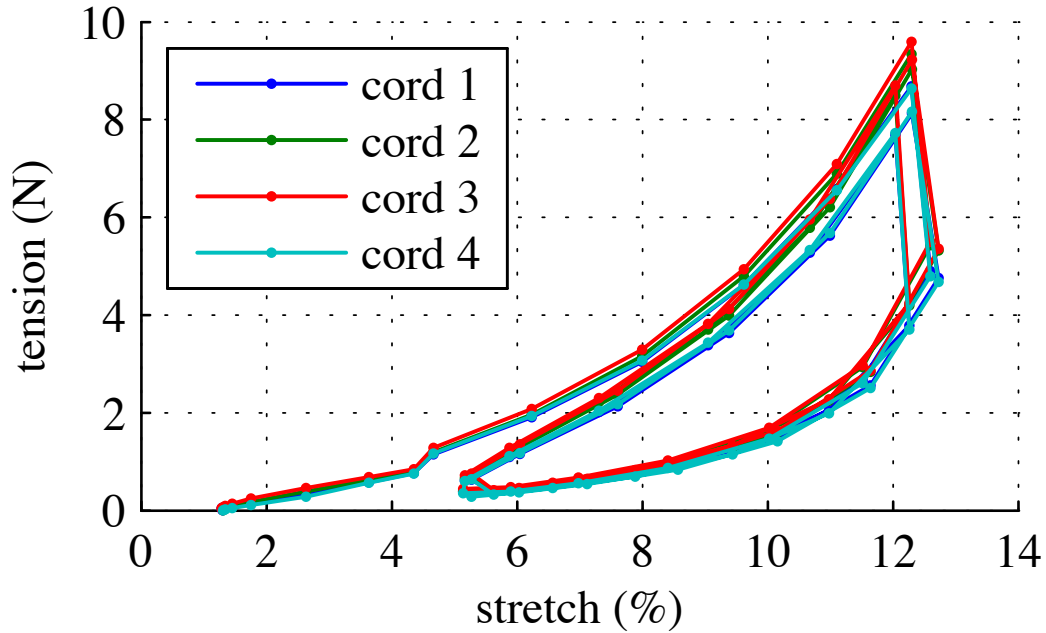


Fig. 3.5 Raw force measurements for X stretching in Sample #18. Cloth is held by two bar clips, with two cords attached to each clip.

points.

3.5 Results

We tested our system on four fabric samples, including a knit and the three common weave patterns (plain weave, twill, and satin), and three fiber types (cotton, wool, and synthetic): cotton satin (#4), rayon/spandex knit (#12), cotton denim (#14), and wool/cotton blend (#18). Each fabric was tested with seven deformations (see Section 3.2.2): for fitting, stretch in X and Y, simple shear, and bending in X and Y; and for evaluation, complex shearing and corner pulling. Fig. 3.5 shows the four individual tension measurements for a typical stretching test, plotted as a function of the overall extension of the cloth. The measurement shows the typical behavior of a woven fabric: a nonlinear curve with increasing stiffness for higher strain, and large hysteresis. The test repeats three times, retracing the same loop each time after the initial extension from rest.

We worked with three cloth models built from the components described in Section 3.3. The Springs model uses the spring membrane model with the spring bending model; the Soft Constraints model uses Baraff and Witkin’s membrane model with the Discrete Shells bending model; and the StVK model uses the diagonalized St. Venant-Kirchoff membrane model with

the Discrete Shells bending model. We fit all the models in four variants: linear (constant stiffness for each deformation mode), isotropic (identical stiffness in warp and weft), linear and isotropic (the simplest variant), and nonlinear orthotropic (the most general variant). The results are too numerous to include in this chapter; we refer the reader to Appendix A, which illustrates the behavior of the nonlinear orthotropic variant of all three models for all four fabrics, and the behavior of the variants of the Soft Constraints model for denim, a largely nonlinear and anisotropic material.

Observations. In Fig. 3.6 and Table 3.2 we present results for fitting the Soft Constraints model to Sample #12. For each test we show a selected frame (near maximum distortion) with renderings illustrating the captured and fitted cloth geometry and forces. To illustrate the fitting residuals more quantitatively, we show a force-displacement plot comparing a summary of the measured forces to the predictions of the fitted model and a vector-field plot illustrating the position error over the geometry of the fitted mesh (see caption for details). Similar tables are shown for all cloth samples and several model variants in the supplementary document.

The four selected fabrics span a large range of possible cloth behaviors. In a nutshell, #12 is isotropic and very compliant in stretch and bending; #4 is also isotropic, very stiff in stretch but compliant in bending; #14 is stiff and quite isotropic in stretch, but extremely anisotropic in bending (with 33/1 stiffness ratio in weft and warp); and #18 is anisotropic both in stretch (with 10/1 stiffness ratio) and in bending (with 13/1 stiffness ratio). The maximum stretch stiffness for #4 is 250 times higher than for #12, while #14 is 10 times stiffer in shear than any other fabric. All four fabrics show similar hysteresis behavior, with loading-to-unloading stretch stiffness ratios ranging from 1.4/1 to 1.8/1. Sample #12 is nearly linear in the test deformation range, while all other three fabrics exhibit nonlinearity. Interestingly, nonlinearity may arise in some deformation modes but not in others, with no clear pattern.

For stretching, all three cloth models fit nicely to the average of the hysteresis bands, even in highly nonlinear cases. The fitting residual is larger for stiffer fabrics, and the nonlinear orthotropic model variants fit anisotropic fabrics best, as expected, while linear and/or isotropic variants reach a reasonable compromise but are not always able to remain inside the hysteresis band. For shearing, the fitting force residual is larger for #14, the stiffest fabric. Across models, the Soft Constraints and StVK models fit to the average of the shearing hysteresis band, while the Springs model deviates at times. For bending, no forces are available, and we evaluate the position residual as well as profiles of sample curves orthogonal to the support plane. The fitting residual is similar for all fabrics, but distinctly higher for the Springs model. Often, the residual is dominated by a difference in curl near the edge of the sample, while the

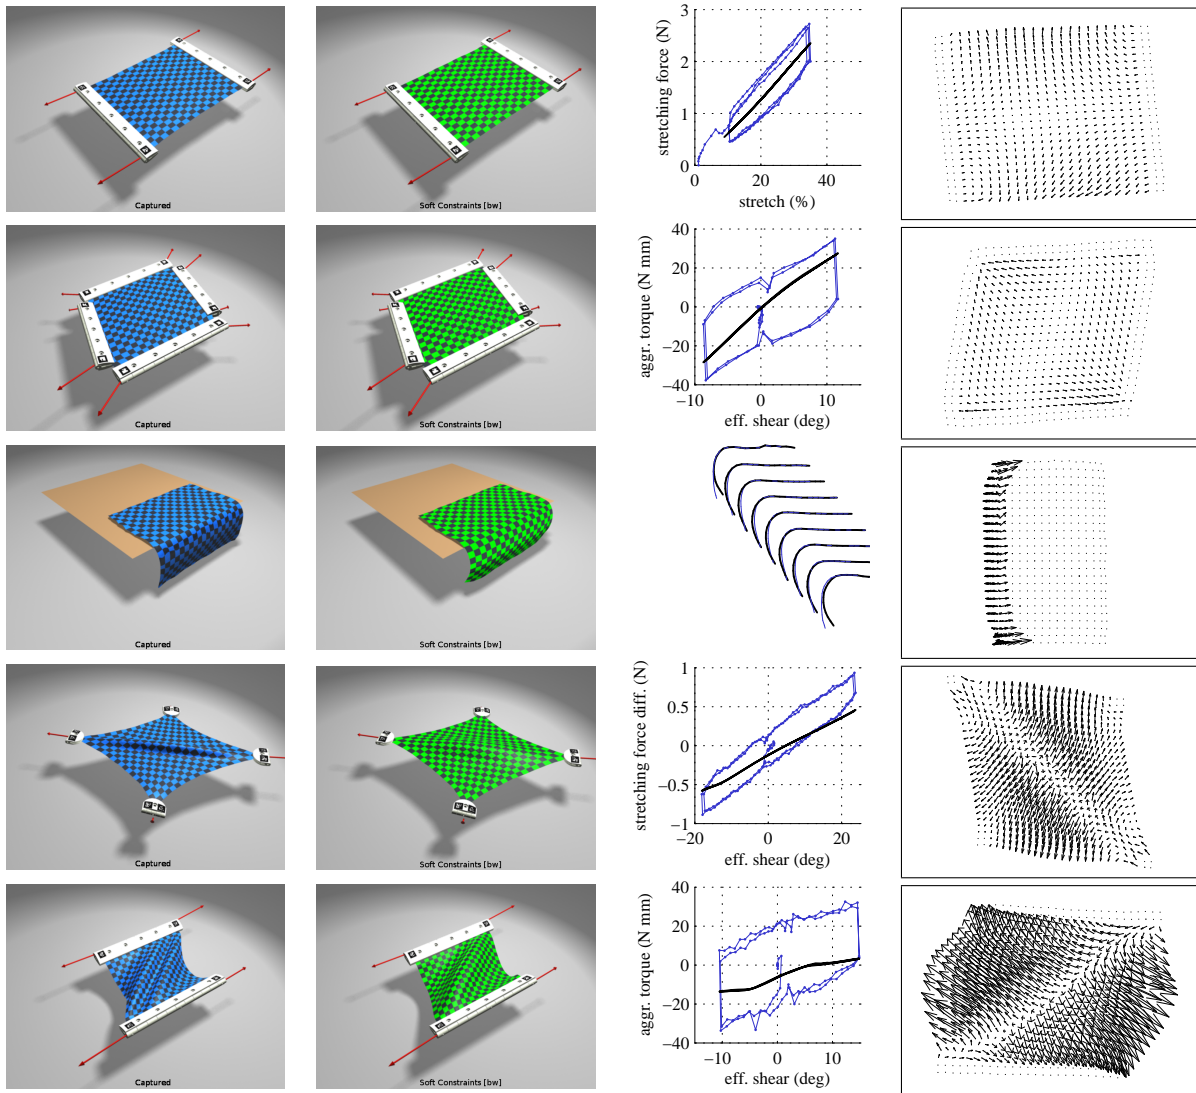


Fig. 3.6 Fitting results for Sample #12 with the Soft Constraints model. Top to bottom; Stretch-X, Simple shear, Bend-X, Corner pull, Complex shear. Left to right: captured geometry, equilibrium of fitted model, force comparison (thin line: measurement; thick line: model), position residual (vertex position minus corresponding measured position, magnified 5x). “Effective shear” is the shear angle of the best-fit transformation to the motion of the clamped cloth vertices. “Aggregated torque” is the torque about the center of the cloth applied by the clamps, with every other clamp counted negatively. “Stretching force difference” is the difference in the two corner-to-corner stretching forces.

overall shape is well fit.

The last two rows of Fig. 3.6 show the evaluation tests; these data were not used in fitting the models. The behavior of sample #12, the most linear fabric, is predicted well in all cases, as seen in the force-displacement plots, the buckling behavior in corner pulling, and the (lower)

Scenario	Stretch-X	Stretch-Y	Shear	Bend-X	Bend-Y
Residual	0.53N	0.62N	0.22N	21.3mm	22.0mm

Table 3.2 RMS residuals of non-linear orthotropic fitting to Sample #12 with the Soft Constraints model.

effective shear stiffness of the sheet when allowed to buckle in the complex shear test. In the three other samples, however, the force-displacement plot of the corner pulling test reveals that the stiffness is underestimated for the Soft Constraints and StVK models, and overestimated for the Springs model. Visually, the mismatch is more apparent in the complex shear test, where models with underestimated stiffness exhibit wider folds than the real fabrics.

We have also evaluated the fitted models on new test samples of each fabric, to validate their generality. Specifically, we have tested stretching on new samples of rayon/spandex knit (#12.2) and cotton denim (#14.2), and shearing on new samples of cotton satin (#4.2) and wool/cotton blend (#18.2). The force-displacement plots of the real cloth samples, shown in the supplementary document, indicate very similar behavior between fitting and test samples for #12 and #14, and a larger disparity for #4 and #18. The evaluation plots for the simulation models behave similarly for the test and fitting cases, but the matching quality depends on the actual disparity across cloth samples.

3.6 Discussion and Future Work

In this chapter, we have presented a novel system for observing cloth behavior, including complete information about deformation and forces, and a new method for fitting and evaluating cloth models using the measurements. Our system is different from standard textile testing systems because it captures detailed geometry information; it is different from previous cloth capture systems in that it captures complete force information and measures deformations of a 3D surface. The combination of very complete position *and* force information provides an unprecedented view into the complex behavior of cloth.

Our measurement setup offers very accurate control over membrane deformations, but the bending tests require manual intervention and are thus less precise. Furthermore, the bending tests are most accurate for samples with straight edges, but some cloth materials (in particular knit) tend to curl up at free boundaries. In order to eliminate these problems, we would like to investigate alternative ways of controlling bending deformations in the future.

The data from our experiments shows some of the limitations of current models. The most obvious of these is hysteresis—all widely used cloth models are elastic, but cloth is clearly far from elastic, resulting in quite large errors for any given point in the experiment. There are many paths for future work in measurement, including more complete exploration of strain space (including compression) and capture of dynamic properties, and in fitting, where new ways of evaluating fitting error are needed that can work when the cloth’s equilibrium state is unstable or non-deterministic.

While overall force-displacement behavior is nicely matched, the actual folding shapes of simulated cloth may deviate largely from the captured cloth, because even a small change in material properties may lead to distant stable configurations in the L^2 sense. For this reason, the traditional L^2 metric is not appropriate for evaluating error in this case. The discontinuity of stable configurations is also the cause of flickering and twitches in some of our examples.

The Springs model exhibits the worst fitting quality in shearing force-displacement curves, and the highest fitting residual for bending. This is probably due to the inherent coupling of stretch and bending deformation components in this model. Nevertheless, the overall deformations in complex shearing fit reasonably well. In contrast to continuum models, complex parameter tuning has often been regarded as a caveat of mass-spring models; but our results indicate that satisfactory parameter estimation is possible by incorporating anisotropy and nonlinearity into the model. The Soft Constraints and StVK models produce results with very similar quality, which is expected as the models present only subtle differences as described in Section 3.3.1.

At least three effects are missed by the tested models: hysteresis, Poisson effect (due to the diagonalization of the standard StVK model), and cross-modal stiffening (e.g., shear stiffening due to stretching). We indeed identified stretch stiffening in the shearing deformations, therefore we chose clip-parallel forces as objective function to minimize the effect of stretch errors on shear optimization. We conjecture that missing cross-modal stiffening may also be, to a large extent, the reason for stiffness underestimation in the corner pulling test for the Soft Constraints and StVK models. An extension to the nonlinear model of Wang et al. (Wang et al., 2011b) could help alleviate these problems.

Chapter 4

Data-Driven Modeling of Hyperelastic Objects

In Chapter 3, we have presented a cloth capture and measurement setup that provides complete information on geometry and boundary conditions (applied loads) in the experiments. Then, we have used that information to estimate strain-dependent parameters of several popular cloth models. However, we identify two major limitations of this approach. One is that through strain-space interpolation of stiffness parameters it is difficult to enforce fundamental properties of elastic forces, such as integrability and energy convexity. The other is that it turns out complicated to estimate parameters in a progressive manner.

In this chapter, we present:

- A general hyperelastic material model based on separable energy terms that can capture extreme nonlinearities and heterogeneity.
- A parameter estimation system that takes advantage of the energy-based deformation model and works with a variety of input data types.

4.1 Introduction

A growing interest in computer graphics is to produce simulation models that match the elastic behavior of deformable objects in the real world. One possibility to reach this goal is to acquire example deformations of real objects, and estimate parameters of constitutive models that best match the acquired examples. Unfortunately, real materials show a high degree of complexity, in the form of heterogeneity, anisotropy, and nonlinearity.




Fig. 4.1 Nonconservative elastic forces may produce uncontrolled energy growth, which leads to stability.

A common temptation in previous work was to produce nonlinear materials by making the coefficients of a linear strain-stress relationship strain-dependent, thus making the resulting model nonlinear [Bickel et al. \(2009\)](#); [Wang et al. \(2011a\)](#). Unfortunately, this choice makes forces non-integrable, as the stress is no longer guaranteed to be the derivative of some energy field.

Fig. 4.1 shows an example of a triangle with two fixed vertices and the third one moved in circles. The plot shows the growing energy, computed as the work done to move the vertex with a Saint Venant-Kirchhoff (StVK) material model with strain-dependent Young modulus $E = \varepsilon_1 + 2\varepsilon_2 + 3\varepsilon_{12}$, where ε_1 and ε_2 are stretches, and ε_{12} is shear strain. As we discuss in the next sections, with our energy-based parameterization of hyperelastic materials we achieve force integrability by construction, and we enforce energy convexity thanks to a constraint-based parameter estimation.

We propose to define general hyperelastic materials using an additive model of separable energy terms, each of them obtained by interpolating strain-dependent energy parameters. With this formulation, elastic models are conservative by construction. In addition, the additive model allows progressive estimation of parameter values, as well as the definition of energy terms that depend on diverse deformation metrics. We also enable the addition of constraints to the energy formulation, allowing the efficient implementation of extreme nonlinearities.

We show that, with the right choice of energy terms and basis functions, our model can exactly match common standard models. But, with our parameter optimization, the general

elastic model can match more complex nonlinear, anisotropic, heterogeneous objects. The optimization takes as input force-deformation examples in static equilibrium, and estimates material parameters such that the deformations and forces best match the measured ones.

Our formulation is general, and can be used to model volumetric elastic objects as well as thin shells. In this chapter, we show the application of our model to the estimation of complex nonlinearities in cloth, extreme nonlinearities of finger skin deformation, and combined mechanical model estimation and non-rigid registration of internal human anatomy in medical imaging data.

4.2 A General Model of Hyperelasticity

In this section, we present our model of hyperelasticity based on the addition of strain-dependent energy terms. Our model is general, and it can model the elastic behavior of volumetric solids, as well as shells through a decomposition of membrane and bending energies. We start with the formal definition of the additive energy model, followed by a description of strain-dependent energy addends. We discuss the connections between our model and several standard hyperelasticity models as well. We continue with a discussion of additional features, such as the enforcement of energy convexity, heterogeneity, and the use of energy constraints. Finally, we conclude with remarks concerning the use of our model in the context of FEM simulations.

4.2.1 Additive Energy Model

The elastic forces of a hyperelastic material can be defined using a generic strain energy density function. Given a point \mathbf{X} in rest position, deformed to position \mathbf{x} , the deformation gradient is defined as $\mathbf{F} = \frac{\partial \mathbf{x}}{\partial \mathbf{X}}$. Then, the strain energy density can be defined as a generic function of the deformation gradient, $\Psi(\mathbf{X}) = f(\mathbf{F}(\mathbf{X}))$. Alternatively, we can express the energy density as a function of the Green strain tensor $\boldsymbol{\varepsilon} = \frac{1}{2}(\mathbf{F}^T \mathbf{F} - \mathbf{I})$, i.e., $\Psi(\mathbf{X}) = f(\boldsymbol{\varepsilon}(\mathbf{X}))$. Finally, elastic forces can be obtained by differentiating the integral of strain energy density. With an energy-based definition of the hyperelastic material, forces are integrable by construction.

The complete strain energy density depends on all terms of the strain tensor, but we find that it is convenient to express the energy using an additive model, where some addends depend on individual strain terms, and others complete the remaining energy as a function of multiple strain terms. This additive decomposition of the energy enables the estimation of individual addends, plus a progressive estimation of the full energy, as we will show in Sec-

tion 4.4 for several examples. The strain energy density can be defined in an additive manner as:

$$\Psi(\mathbf{X}) = \sum_i \Psi_i(\varepsilon_i(\mathbf{X})). \quad (4.1)$$

In this representation, the addend-wise dependency of strain, ε_i , denotes the possibility to use a different strain metric for each energy addend. In this way, our additive energy definition is very general, and it can represent as additive models popular hyperelastic models, such as StVK, Ogden, Neo-Hookean, etc. (Ogden, 1997). In the StVK model the energy addends are given by quadratic terms of the Green strain tensor, in the Ogden model they are given by various powers of principal stretches, and in the Neo-Hookean model the energy depends on the first invariant of the left Cauchy-Green strain tensor, which corresponds to the sum of squared principal stretches.

4.2.2 Interpolated Energy Functions

Since our goal is to estimate hyperelastic materials in a data-driven manner, we design a representation of each energy addend Ψ_i that allows capturing arbitrary nonlinearities (up to hyperelasticity, not other nonlinearities such as time-dependency). We design each energy addend as the interpolation of energy samples in the addend-dependent strain domain:

$$\Psi_i(\mathbf{X}) = \sum_j \phi(\varepsilon_i(\mathbf{X}) - \varepsilon_{ij}) \Psi_{ij}, \quad (4.2)$$

where ϕ denotes some basis function, Ψ_{ij} its corresponding weight, and ε_{ij} a particular sample of the strain component ε_i .

We have considered two general ways of interpolating energy samples. For one-dimensional strain domains, we distribute energy control points evenly, and interpolate them using cubic Hermite splines. For high-dimensional strain domains, we use Gaussian radial basis functions to interpolate scattered control points. By using basis functions with local support, during data-driven parameter estimation each energy control point is influenced only by a fraction of the input data, thereby simplifying the structure of the optimization problem.

But the choice of basis functions also allows us to replicate standard models exactly. Let us take, for instance the StVK model, with strain energy density $\Psi = \frac{\lambda}{2} \text{tr}(\varepsilon)^2 + \mu \text{tr}(\varepsilon^2)$, λ and μ Lamé parameters, and ε Green strain. In the 2D case, with stretches ε_1 and ε_2 and shear

strain ε_{12} , this strain energy density can be rewritten in an additive manner as:

$$\Psi = \underbrace{\left(\frac{\lambda}{2} + \mu\right) \varepsilon_1^2}_{\Psi_1} + \underbrace{\left(\frac{\lambda}{2} + \mu\right) \varepsilon_2^2}_{\Psi_2} + \underbrace{\lambda \varepsilon_1 \varepsilon_2}_{\Psi_3} + \underbrace{2\mu \varepsilon_{12}^2}_{\Psi_4}. \quad (4.3)$$

The energy addends Ψ_1 , Ψ_2 and Ψ_4 are unimodal and can be represented in the format of Eq. 4.2 using just one quadratic basis function each. The addend Ψ_3 , on the other hand, is bimodal and can be represented using one bilinear basis function. We have tested the ability of our additive interpolated model and optimization framework to match in practice test deformation data generated with a pure StVK model, and we found that energies are matched exactly.

But once the correspondence with standard models is analyzed, the power of our proposed additive model lies in the versatility to introduce additional features. Some of them are:

- Anisotropy, which could be added, e.g., to Eq. 4.3 simply by using different functions for the stretch energy addends Ψ_1 and Ψ_2 .
- Additional strain-dependency (e.g., material hardening), which could be obtained by adding further control points to each energy addend.
- Cross-stiffening, which could be modeled with an additional addend dependent on the principal strain.
- Volume/Area preservation or strain limiting, which could be modeled as energy addends that weight the deviation w.r.t. default volume, area, or strain values. Patterson et al. (2012) also formulate volume preservation as an energy addend.

4.2.3 Convexity

In addition to the integrability of elastic forces, another desirable property of a strain energy density function is convexity, which will ensure passivity of the mechanical model and stability of numerical integration under some choice of time step. An energy function is convex iff the eigenvalues of its Hessian are always positive. The definition of an additive energy model with separable addends might suggest enforcing convexity of each addend individually, but this limits the versatility of the optimization procedure in practice. Taking again as example the 2D StVK model in Eq. 4.3, the bimodal addend Ψ_3 , which captures area preservation, is concave. With a regular StVK model, enforcing convexity of the added energy is simple, as




Fig. 4.2 A large negative λ , that creates a nonconvex strain energy density function, produces a diverging motion instead of restoring the triangle to its rest configuration.

it reduces to imposing conditions on the Lamé parameters. The shear modulus μ is always positive, but a large negative value of λ can turn the energy concave along some direction. Fig. 4.2 shows diverging motion of a deformed triangle due to large negative λ . With a strain-dependent model, energy convexity is harder to enforce, as it may be locally violated.

Fortunately, we relax the need to explicitly constrain the energy Hessian in the full strain domain in the following way. In practice, we found that unimodal energy addends (e.g., stretch, shear, cross-stiffening, or strain limiting) can be assumed to be convex. Hence, we enforce convexity of these addends separately, simply in their respective one-dimensional domains. For bimodal energy addends, since all other addends are convex, it turns out sufficient to enforce convexity of the full energy in the two-dimensional domain of the bimodal addend. In Section 4.3.3, we describe the enforcement of convexity in the context of our optimization framework for parameter estimation.

4.2.4 Heterogeneity

To model heterogeneous materials, we make use of the flexibility of our additive energy model. Given a deformable object, we sample it in material space at the desired resolution to capture material heterogeneity, and we define a separate additive energy model Ψ_k at each sample point \mathbf{X}_k . Then, to obtain the strain energy density at an arbitrary test point \mathbf{X} in material space, we first evaluate energies at the sample points using the deformation at the test point \mathbf{X} , and then interpolate these energies:

$$\Psi(\mathbf{X}) = \sum_k \phi(\|\mathbf{X} - \mathbf{X}_k\|) \Psi_k(\boldsymbol{\varepsilon}(\mathbf{X})). \quad (4.4)$$

In our examples, we interpolate strain energy densities in material space using as basis

functions ϕ normalized Gaussian radial basis functions.

4.2.5 Energy Constraints

Modeling extremely stiff materials using strain energy density functions implies two major difficulties. One is the need to solve stiff systems of equations during simulation, which usually requires tiny time steps and slows down the convergence rate of commonly used solvers. The other is that the generation of reliable force-deformation data requires expensive setups.

As an alternative to stiff energy models, we propose to augment the additive energy model described so far with energy constraints. Several authors have enforced in the past constraints on strain (Goldenthal et al., 2007; Perez et al., 2013; Provot, 1995; Tomaszewski et al., 2009; Wang et al., 2010) or volume (Irving et al., 2007) to model extreme nonlinearities. As a small variant, we express constraints on strain energy density, which allows us to constrain either individual addends, combinations, or the total energy. In general, energy constraints can be expressed as:

$$C = \Psi_{\max} - \Psi_i(\mathbf{X}) \geq 0. \quad (4.5)$$

4.2.6 FEM Simulation

In our examples, we have used FEM to discretize elastic deformation problems, and we have tested tetrahedral elements with linear interpolation and hexahedral elements with trilinear interpolation. To compute elastic forces, we first integrate the strain energy densities over each element using quadrature points (1 point for tetrahedra and 8 points for hexahedra), and then differentiate these energies w.r.t. element nodes. We have implemented elastic force computation using finite differences, which reduces to performing multiple evaluations of each element's strain energy (13 for tetrahedra and 25 for hexahedra). Each strain energy density evaluation on a quadrature point requires: first, the evaluation of each energy addend for each nearby material sample, according to Eq. 4.2; second, the computation of the total energy for each nearby material sample, according to Eq. 4.1; and third, the interpolation of energies from material samples, according to Eq. 4.4. Implicit integration or (quasi-)static simulation also require the Hessian of the energy, which again we compute through finite differences of forces.

4.3 Data-Driven Material Estimation

In this section, we describe our general optimization framework for the estimation of material parameters in our energy model. Later in Section 4.4, we discuss how we use this framework iteratively in several applications with very diverse deformations and data, ranging from sparse to dense measurements.

Let us consider a general deformable model with a vector \mathbf{q} that concatenates all its nodal positions, and a vector \mathbf{f} that concatenates all nodal forces. Our optimization framework receives as input a set of N example deformations in static equilibrium. Each of the example deformations is produced from some known boundary conditions (forces \mathbf{f}_c and/or positions \mathbf{q}_c), and contains some known measurements $\bar{\mathbf{m}}$. We consider diverse types of measurements, such as positions, forces, or image intensities.

The goal of the optimization framework is to estimate a vector \mathbf{p} of material parameters consisting of energy control points, such that the least-squares error between the known measurements $\bar{\mathbf{m}}$ and their estimated values $\mathbf{m}(\mathbf{f}, \mathbf{q})$ is minimized, subject to static equilibrium and energy convexity. We formally define the objective function of the optimization by summing the error over all examples:

$$f = \frac{1}{2} \sum_{i=1}^N w_i \|\mathbf{m}_i([\mathbf{q}_i, \mathbf{f}_i](\mathbf{p}, \mathbf{f}_{c,i}, \mathbf{q}_{c,i})) - \bar{\mathbf{m}}_i\|^2. \quad (4.6)$$

We enable the possibility to weight examples differently with $\{w_i\}$. In the expression above, we explicitly indicate that estimates of measured variables \mathbf{m} depend on simulated positions and forces, and through these they depend on estimated parameters and boundary conditions. Note that the measured variables and the nodes where boundary conditions are applied may differ across examples.

Our optimization framework iterates the following three steps until convergence:

1. Update of parameter estimates.
2. Convex energy projection of parameters.
3. Static equilibrium on all examples.

Next, we discuss these three steps in detail.

4.3.1 Parameter Estimation

In each iteration of parameter estimation, we minimize Eq. 4.6 subject to energy convexity and static equilibrium. We treat these two types of constraints differently. We do not enforce energy convexity during parameter updates, and instead, we compute an unconstrained update and project the parameters to a convex configuration afterwards, as described in Section 4.3.3. On the other hand, we enforce static equilibrium implicitly during parameter updates. We do this by computing a Jacobian of positions w.r.t. parameters that respects static equilibrium. After the parameter update, it is anyway necessary to project all examples to static equilibrium, as they may have slightly deviated.

For each input example, we formulate the Jacobian of positions w.r.t. parameters $\frac{\partial \mathbf{q}_i}{\partial \mathbf{p}}$ through the application of the implicit function theorem on the static equilibrium constraints. In its most general formulation, our energy model includes energy constraints; therefore, static equilibrium forces include energy constraint forces: $\mathbf{f}_i + \mathbf{J}_i^T \lambda_i = \mathbf{0}$. In this expression, we assume that energy constraints are enforced through Lagrange multipliers, and $\mathbf{J}_i = \frac{\partial \mathbf{C}_i}{\partial \mathbf{q}_i}$ is the Jacobian of active energy constraints \mathbf{C}_i for the i^{th} example. The enforcement of hard energy constraints complicates the application of the implicit function theorem, as it requires to first solve for Lagrange multipliers.

However, we observe that the purpose of the implicit function theorem is just to provide a suitable Jacobian of positions w.r.t. parameters. For this purpose, we found that it is sufficient to approximate energy constraints as soft constraints with energy $\frac{1}{2} k \|\mathbf{C}\|^2$, for some large value of the stiffness k . Taking only active constraints, the static equilibrium in the i^{th} example can then be reformulated as:

$$\mathbf{f}_i - k \mathbf{J}_i^T \mathbf{C}_i = \mathbf{0}. \quad (4.7)$$

Differentiating the static equilibrium constraints w.r.t. the vector of parameters, we obtain:

$$\frac{\partial \mathbf{f}_i}{\partial \mathbf{p}} - k \mathbf{J}_i^T \frac{\partial \mathbf{C}_i}{\partial \mathbf{p}} + \frac{\partial \mathbf{f}_i}{\partial \mathbf{q}_i} \frac{\partial \mathbf{q}_i}{\partial \mathbf{p}} - k \mathbf{J}_i^T \mathbf{J}_i \frac{\partial \mathbf{q}_i}{\partial \mathbf{p}} = \mathbf{0}. \quad (4.8)$$

And from this we obtain the Jacobian of positions that satisfies equilibrium:

$$\frac{\partial \mathbf{q}_i}{\partial \mathbf{p}} = \left(\frac{\partial \mathbf{f}_i}{\partial \mathbf{q}_i} - k \mathbf{J}_i^T \mathbf{J}_i \right)^{-1} \left(\frac{\partial \mathbf{f}_i}{\partial \mathbf{p}} - k \mathbf{J}_i^T \frac{\partial \mathbf{C}_i}{\partial \mathbf{p}} \right). \quad (4.9)$$

In practice, we evaluate all terms in this expression using finite differences, but accounting for sparsity patterns for efficiency. If the energy model does not include energy constraints, it is sufficient to drop the related terms from the expression.

By enforcing static equilibrium constraints implicitly in the Jacobian of positions, the op-




Fig. 4.3 Evaluation of the fitting quality of our energy model on force-deformation data of two cloth samples. Left half: sample A, rayon/spandex knit. Right half: sample B, cotton/denim. The first and third columns show one frame of a *corner-pull* deformation, and the second and fourth columns show one frame of a *complex shear* deformation. The top row shows the input data (not used for training), the middle row shows simulated results with our fitted model, and the bottom row shows per-vertex error.

timization problem Eq. 4.6 turns into an (unconstrained) nonlinear least-squares problem. We solve this problem iteratively using Gauss-Newton followed by a line-search. Gauss-Newton requires the evaluation of the derivative of the estimates of measurements \mathbf{m}_i , which can be written as:

$$\frac{\partial \mathbf{m}_i}{\partial \mathbf{p}} = \frac{\partial \mathbf{m}_i}{\partial \mathbf{q}_i} \frac{\partial \mathbf{q}_i}{\partial \mathbf{p}} + \frac{\partial \mathbf{m}_i}{\partial \mathbf{f}_i} \frac{\partial \mathbf{f}_i}{\partial \mathbf{p}}. \quad (4.10)$$

In this expression, we plug in the Jacobian of constrained positions Eq. 4.9. The Jacobians of estimates of measurements w.r.t. positions $\frac{\partial \mathbf{m}_i}{\partial \mathbf{q}_i}$ and w.r.t. forces $\frac{\partial \mathbf{m}_i}{\partial \mathbf{f}_i}$ are trivial when those measured variables are positions and/or forces, as is the case in some of our experiments. In cases where the measured variables depend in a complex way on simulated forces and positions, such as image intensity in one of our experiments, we opt to compute their derivatives w.r.t. the parameters directly through finite differences.




Fig. 4.4 Evaluation of the fitting quality of our energy model on 3D MRI data of a knee. The left image and the three top images constitute the input MRI data [Rhee et al. \(2011\)](#). The bottom three volume-rendering images have been obtained by deforming the left image to best match the top images, in a combined hyperelastic model estimation and image registration.

4.3.2 Static Equilibrium

After each update of material parameters, we reproject all input examples to static equilibrium. We use a different method depending on the type of deformable object and energy model.

For volumetric objects whose energy model does not contain energy constraints, we found that Newton-Raphson-type root-finding performs best. For thin shells (e.g., cloth) with no energy constraints, we found that dynamic relaxation with kinetic damping ([Volino and Magnenat-Thalmann, 2007](#)) performs best.

For objects with energy constraints, static equilibrium should be posed as the minimization of potential energy (elastic plus gravitational) subject to the energy constraints. We found that, instead of imposing constraints on quasi-static iterative updates, we obtain better convergence by imposing constraints on dynamic solves in the context of dynamic relaxation. For the enforcement of energy constraints in dynamic solves, we use a Lagrange multiplier formulation and an LCP solver ([Cottle et al., 1992](#)), analogous to constraint-based solvers for contact problems.

4.3.3 Convex Energy Projection

In Section [4.2.3](#), we have discussed the importance of enforcing convexity of strain energy density functions. For our additive energy model, a sufficient convexity condition is to impose convexity of unimodal addends in their one-dimensional strain domain, together with convexity of the full energy in the strain domain of multimodal addends.

Dataset	StVK	A. Stretch	Nl. Stretch	A. Shear	Area Pres.	Cross-Modal
A-Pos. (mm)	0.6299	0.6229	0.6242	0.6177	0.6199	0.6194
B-Pos. (mm)	0.8549	0.8584	0.8261	0.7983	0.8031	0.8031
A-Force (N)	0.3484	0.3482	0.3244	0.3196	0.3174	0.3170
B-Force (N)	2.5816	2.5276	2.2807	2.2605	2.2569	2.2569

Table 4.1 Evolution of the position and force RMS error of the two cloth samples as the strain energy density is augmented with more addends. From left to right: Isotropic StVK, Anisotropic Stretch, Nonlinear Stretch, Anisotropic Shear, Anisotropic Area Preservation, Cross-Modal Stiffening.

Given a possibly concave energy function resulting from a parameter update, we project the energy to a convex space by computing the most similar energy function that satisfies convexity constraints at a discrete set of strain samples. Let us denote as $\bar{\mathbf{p}}$ the parameter vector of energy control points resulting from the optimization step. We wish to compute a new parameter vector \mathbf{p} that minimizes the deviation $\|\mathbf{p} - \bar{\mathbf{p}}\|^2$. We formulate constraints based on eigenvalues λ_j of the Hessian of energy addends Ψ_i (or the full energy), evaluated at strain samples ε_k :

$$C_{i,j,k} = \lambda_j(\mathbf{H}(\Psi_i(\mathbf{p}, \varepsilon_k))) \geq 0. \quad (4.11)$$

We solve this constrained optimization using Interior-Point method in Matlab. This method handles well complex nonlinear constraints, and it does not become a bottleneck in our experiments.

4.4 Results

We have tested our modeling and estimation framework on a variety of deformable objects and applications. Each benchmark allows us to test certain aspects of our approach, and altogether they allow us to validate its generality. In all benchmarks we fit our energy model using real-world deformations measurements, with high nonlinearity and possibly measurement error. All experiments were executed on a 2.67 GHz Intel Core i7 920 CPU with 12 GB RAM.

Cloth Models from Force-Deformation Data For this experiment, we have used data captured by the measurement system described in Section 3.2. Specifically, we have used data from cloth samples 12 (A: rayon/spandex knit) and 14 (B: cotton/denim).

We have evaluated our additive energy model and its effect on parameter estimation by




Fig. 4.5 Cross-sections of the knee MRI in the pose with largest bending. Left: original MRI; Middle: deformation of the straight MRI with our estimated heterogeneous model; Right: deformation with a homogeneous model. This is one of the most challenging slices, and it shows the improved registration with the heterogeneous model.

fitting an increasingly complex energy model to the training data. We discretize the cloth geometry using an irregular triangle mesh with 520 triangles. To formulate the objective function Eq. 4.6, we use as measured variables the positions of mesh vertices and the forces on clips attached to the cloth samples. Table 4.1 shows the evolution of position and force RMS error as the energy model is augmented with more addends. We start with an isotropic StVK model for membrane deformations and the discrete shells model for bending (Grinspun et al., 2003), and then add stretch anisotropy, stretch nonlinearity, shear anisotropy, anisotropic area preservation, and cross-modal stiffening based on the principal strain.

Finally, we have also validated the fitted energy model on the two test datasets. Fig. 4.3 shows representative frames of the two datasets for both cloth samples, with a plot of per-vertex error.

Internal Tissue Model and Image Registration For this benchmark, we have used 3D MRI data of a knee obtained by Rhee et al. (2011). The dataset includes volumetric images for 4 different poses of the knee, shown in Fig. 4.4. This scenario is particularly challenging because the underlying mechanical model is heterogeneous and there is no correspondence information between images. The scenario constitutes a classic problem for non-rigid image registration, for which many approaches exist. However, due to the large deviation between the input images, standard methods fail, as we have attested using ITK. To avoid the problem of large initial deviations, an alternative is to first design a patient-specific FEM model, then deform the source image into the target image using this FEM model, and finally execute the

registration algorithm (Han et al., 2011). Leveraging our data-driven estimation approach, we have tested an extension that executes progressive registration plus hyperelastic model estimation.

We started by segmenting the leg in the first pose, and we meshed it with a tetrahedral mesh (with 822 tets in our example). We then manually defined a crude correspondence between the 4 input poses (15 point correspondences in our case). We used a subset of 8 correspondences on the outer skin as boundary conditions during the fitting process, and 7 correspondences in the interior for the definition of the objective function as the sum of squared position deviations. With this crude data, we estimated a homogeneous elastic model.

Then, we progressively added material samples throughout the model, until the optimization converged. For subsequent registration and material estimation steps, we defined as objective function Eq. 4.6 the sum of image differences between the input poses and the deformed ones. To evaluate the objective function, we rasterized the deformed tetrahedral mesh using the source volume as 3D texture map (Gascon et al., 2013), and simply computed image differences w.r.t. the destination volumes for all voxels inside the tetrahedral mesh.

Fig. 4.4 shows in its bottom row the poses obtained by deforming the source volume with our final fitted energy model. This model contains 16 material samples to capture tissue heterogeneity. The fitted model is not perfect, and it provides a least-squares solution that weights error in all target poses. As an extra spin to our method, we have also tested to use its deformed output as input to free-form deformation methods for image registration in ITK (Mattes et al., 2003; Rueckert et al., 1999). Table 4.2 shows the improvement in RMS image error from the homogeneous model to the final heterogeneous model, as well as the error using ITK’s registration initialized with both models. Fig. 4.5 also compares slices of target and deformed volumes resulting from the heterogeneous model.

Method	Homog.	Homog. + ITK	Heter.	Heter. + ITK
Error	65.88	60.80	56.48	52.62

Table 4.2 Registration error (RMS of voxel intensities in a scale from 0 to 255) for the knee MRI data in Fig. 4.4, using several registration methods: homogeneous (hom.) or heterogeneous (het.) energy model, with/without ITK post-registration.

Nonlinear Skin from Finger Pressing Experiments Our final benchmark involves the estimation of a hyperelastic model of finger skin from controlled force-deformation data provided by collaborators from the European project Wearhap. In this benchmark, in particular, we have

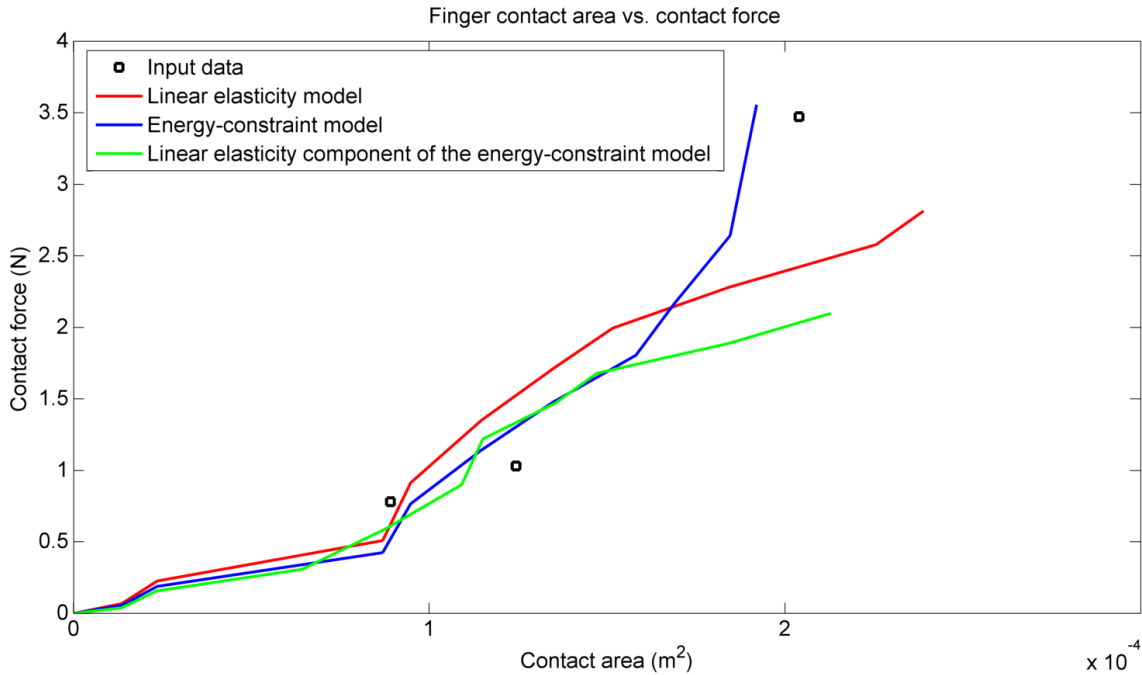


Fig. 4.6 Plots of contact force vs. contact area for a fingerpad. The dots indicate the input examples, and the solid lines the results produced by our hyperelastic model with energy constraints (blue), the elastic component of this hyperelastic model (green) and the linear elastic model (red).

evaluated the estimation of energy constraints formulated in Section 4.2.5 to capture extreme nonlinearity. The input data, shown in Fig. 4.6, consists of paired values of contact area vs. total normal force for one subject’s index fingerpad. The data was collected keeping the subject’s finger fixed, with the nail-side resting on a fixed surface, while a flat transparent moving platform was pressed against the fingerpad. The platform was equipped with a force sensor, and the contact area was estimated by capturing the image of the pressed finger through the transparent platform.

Using this data, we have estimated a homogeneous hyperelastic model for the skin of the fingertip. We have replicated the capture scenario, modeling the finger with a 347-tetrahedra mesh fixed on the nail area, as shown in Fig. 4.7. We have clustered the input data into 4 force-deformation examples. For each of the examples, we move a fixed platform against the fingerpad until the contact area matches the input value. We formulate the objective function as the difference between the simulated contact force on the fingerpad and the input force value, and we update the model parameters as described in Section 4.3.1. In this benchmark, once a new set of parameters is computed, we enforce energy convexity and recalculate static equilibrium as described in Section 4.3.2, but we also reposition the moving platform for each




Fig. 4.7 Simulation of the finger deformation capture experiment. Left: finger in rest pose. Right: finger deformed due to the platform pressing against the fingerpad.

example to match the input contact area.

We have estimated energy control points for a regular StVK model with the addition of an energy constraint, alternating the estimation of the StVK model and the energy constraint until convergence. The force-area result produced by the resulting model is compared to the input data in Fig. 4.6.

4.5 Discussion and Future Work

In this chapter, we have presented a general formulation of hyperelasticity based on energy addends that allows modeling and estimating various nonlinear elasticity effects in a separable manner. We couple our formulation with an optimization algorithm that enables data-driven estimation of complex hyperelasticity in diverse applications. Two of the main features of our model are the enforcement of force integrability and energy constraints, overlooked in previous approaches.

The performance of the optimization algorithm could be further improved. One of its limitations is that, in practice, we do not guarantee global enforcement of energy convexity. We enforce convexity only at discrete strain samples, for efficiency reasons. It is worth looking into simpler ways of enforcing convexity, perhaps by formulating energy addends that are convex in nature, and then global convexity could be enforced through individual convexity. Another limitation of our current enforcement of energy convexity is that convexity constraints are not satisfied implicitly during parameter updates, only through projection afterwards. The efficiency of the optimizer would increase by handling convexity constraints implicitly, as currently done for static equilibrium constraints.

The current bottleneck in the optimization is the solution of static equilibrium conditions

after every parameter update. Our model would benefit from faster static equilibrium solvers, as well as more drastic minimization methods, which would reduce the number of static equilibrium solves. Concerning the optimization algorithm, one last limitation is that we rely on gradient-based optimization, which does not prevent falling in local minima. We partially avoid such local minima thanks to the incremental increase of complexity of the parameter space, but with no absolute guarantees.

In addition to hyperelasticity, the example data used in our experiments is likely to exhibit other nonlinearities, such as hysteresis. Even though modeling such nonlinear effects is orthogonal to our contribution, the fact that we do not account for them may bias the estimation of hyperelasticity in some cases.

We have demonstrated the applicability of our energy model in diverse settings, but further work would be needed to achieve more accurate material estimations. In particular, in the medical imaging benchmark, our combined parameter estimation and registration algorithm is not at the level of specialized registration methods. We instead introduce an interesting spin to regular methods, which handle registration and model estimation as two separate tasks. For model estimation, the results could be largely improved by segmenting the data, meshing it according to the segmentation, and accounting for contact between anatomical structures. In the finger skin estimation benchmark, the input data is rather crude, just a 1D function, hence it is insufficient for estimating an accurate model of skin.

To conclude, we think that an interesting extension to our energy model would be to identify representative energy addends that both minimize the number of parameters and ease their estimation, while allowing simple enforcement of convexity. One possibility would be to perform extensive data mining in force-deformation data, to identify such representative energy addends.

Chapter 5

Modeling and Estimation of Internal Friction for Cloth

Previous chapters have focused on modeling hyperelasticity in cloth, always ignoring internal friction. Several works in mechanical engineering have studied cloth hysteresis, but there is very little work on it in computer graphics and, in particular, no knowledge about its visual effects. In this chapter we present:

- An internal friction model, based on Dahl’s friction model ([Dahl, 1968](#)), that is able to produce the observed hysteresis effect.
- A reparameterization of the standard model to account for strain-dependent behaviors of cloth.
- A capture and measurement setup that allows us to estimate both elasticity and internal friction parameters.
- An analysis of the visual impact of internal friction in cloth simulation.

This work has been published in ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia) ([Miguel et al., 2013](#)).

5.1 Introduction

Internal friction in cloth is the source of its characteristic hysteresis behavior. While several previous works ([Lahey, 2002](#); [Ngo-Ngoc and Boivin, 2004](#)) have used complex second order models to capture internal friction, from our observations of real cloth samples, we conclude

that simpler first order models, such as Dahl’s friction model, can effectively produce realistic results. We apply this friction model on cloth, but it can be easily extended to solids as long as the constitutive model used to capture elasticity is able to produce scalar strain measurements. As we will discuss in Section 5.2.4, we also identify features that are not well captured by the standard model, and propose an augmented reparameterization that includes a strain-dependent definition of hysteresis.

In contrast to previous work, we will also analyze the visual impact of internal friction in cloth simulation, describing several interesting visual effects:

- ‘Preferred’ wrinkles and folds, as in Fig. 5.15.
- Persistent deformations, as in Fig. 5.17.
- History-dependent wrinkles, as in Fig. 5.14.
- Faster settling of folds and wrinkles, as in Fig. 5.16.

We also propose simple and inexpensive data-driven procedures for friction parameter estimation, described in Section 5.3. Previous work relied on complex force-deformation measurement systems with uniform strain, controlled deformation velocity, and dense data acquisition (Kawabata, 1980). Our procedures, in contrast, need only sparse data and can deal with non-uniform strain. We demonstrate procedures for stretch and bending acquisition that lead to simple optimization problems for parameter estimation.

Finally, in Section 5.4, we present a simple algorithm for the efficient simulation of internal friction using implicit integration methods. We validate our parameter estimation for several fabrics by simulating the acquisition procedures, and we also show cloth animation examples that exhibit disparate behavior with and without internal friction.

5.2 Dahl’s Friction Model

In this section we will first introduce the underlying elasticity model and then discuss the connection between Dahl’s model and the observed features of hysteresis. Most importantly, we highlight features that are not captured by the standard Dahl model, and propose an augmented reparameterization that is well suited for parameter estimation given force-deformation measurements.

5.2.1 Underlying Elasticity Model

To model cloth elasticity we use a thin shell model, similar to the one described in Chapter 3, that separates membrane and bending energies. For membrane elasticity, we adopt the St. Venant-Kirchhoff (StVK) model [Ogden \(1997\)](#), with strain-energy density $W(\mathbf{E}) = \frac{\lambda}{2} \text{tr}(\mathbf{E})^2 + \mu \text{tr}(\mathbf{E}^2)$. \mathbf{E} is the Green-Lagrange strain tensor, and λ and μ are the Lamé constants. With a triangle-based discretization, strain is considered constant per triangle, energy density is integrated over the triangle area A , and membrane force is computed as $\mathbf{F} = A \nabla W$.

For bending elasticity, we adopt the discrete shells bending model [Grinspun et al. \(2003\)](#). Under this model, the discrete strain associated to a deformed edge can be expressed as $\varepsilon = 3 \frac{\theta - \bar{\theta}}{h}$, where θ (resp. $\bar{\theta}$) is the deformed (resp. undeformed) dihedral angle, and h is the undeformed average height of the incident triangles. Bending forces are computed as $\mathbf{F} = A \sigma \nabla \varepsilon$, where $A = \frac{1}{3} e h$ is the area associated to an edge, e the rest length of the edge, $\sigma = k \varepsilon$ the bending stress, and k the bending stiffness.

For our exposition, it is convenient to employ a unified representation of scalar strain and stress components. In the remaining, we denote by ε (resp. σ) an arbitrary scalar strain (resp. stress), either the bending strain (resp. stress) or an individual component of the membrane strain (resp. stress). Then we express the force due to a scalar stress on a deformed element with area A as $\mathbf{F} = A \sigma \nabla \varepsilon$, with $\sigma = k \varepsilon$, and k the stiffness of the stress on the element.

5.2.2 Characterizing Hysteresis

[Fig. 5.1](#) shows plots of measured end-to-end distance vs. applied force for 10cm \times 10cm square pieces of rayon/spandex knit (left) and denim (right). We used the capture system described in [Section 3.2](#) and applied loading-unloading cycles between 20% and 60% stretch for rayon/spandex and between 0% and 3% for denim, centered at different offset deformations. We repeated the tests under deformation speeds ranging from 0.5mm/s to 2mm/s, obtaining very similar force values. While these experiments relied on complex hardware, our estimation procedures, described in [Section 5.3](#), rely on very simple setups.

We draw the following conclusions from the force-deformation data in [Fig. 5.1](#). First, hysteresis is a strong effect, with a ratio $\frac{F_{\max} - F_{\min}}{F_{\text{avg}}}$ at a given deformation of up to 50%. Second, the amount of hysteresis depends on the deformation range of the load-unload cycle and the offset deformation. Third, the transition from the load to unload path and back exhibits pre-sliding, i.e., the opposing friction grows smoothly, not sharply. In our work, we show that a model of internal friction based on Dahl's model can replicate this behavior.

Another interesting conclusion about hysteresis is that it produces persistent deformations.




Fig. 5.1 Force-deformation curves for load-unload cycles of uniform stretch on a rayon/spandex knit (left) and denim (right).

Fig. 2.2-left shows an ideal linear elastic force with ideal hysteresis cycles. In the unload path, the displacement stops when the total internal stress is canceled, producing a persistent deformation ϵ_{per} . We have observed such persistent deformations in stretch and bending experiments with real cloth, and we show that Dahl's model can replicate this behavior too.

In the rest of this chapter, a positive internal force will indicate a force that opposes deformation, hence the total force acting on cloth should be evaluated as $F_{\text{external}} - F_{\text{internal}}$.

5.2.3 Model and Interpretation

For any scalar strain metric ϵ (e.g., weft or warp stretch, shear, or discrete-shells bending strain), Dahl's model (Dahl, 1968) can be used to define a scalar friction stress σ in terms of the change of stress as a function of strain, as follows:

$$\frac{d\sigma}{d\epsilon} = \alpha \left(1 - s \frac{\sigma}{\sigma_{\max}} \right)^\gamma. \quad (5.1)$$

σ_{\max} is the maximum friction stress, α is the slope of the stress ratio $\frac{\sigma}{\sigma_{\max}}$ at $\sigma = 0$, and $s = \text{sign}(\dot{\epsilon})$ indicates the sign of the rate of deformation. In the rest of the chapter, we restrict the discussion to a Dahl model with $\gamma = 1$. With constant parameters σ_{\max} and α , this model corresponds to a first-order system with time constant $\tau = \frac{\sigma_{\max}}{\alpha}$, and limiting stress σ_{\max} when strain is increasing, or $-\sigma_{\max}$ when strain is decreasing. Recall that stress is transformed to forces through multiplication by the element's area A and the strain gradient $\nabla\epsilon$, as discussed in Section 5.2.1.

To interpret the parameters of Dahl's model, it is convenient to analyze the behavior of a

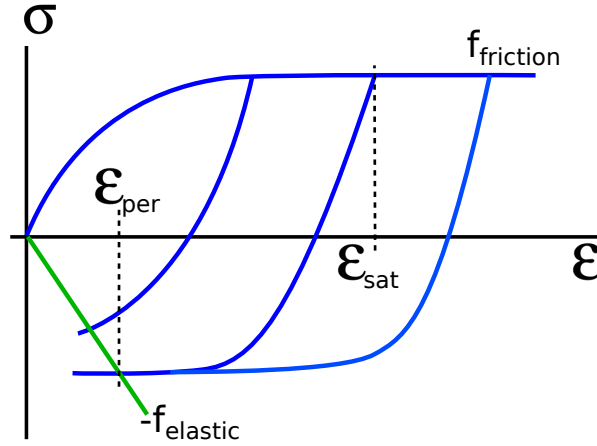


Fig. 5.2 Three different unload paths for friction forces and negative elastic forces. Intersection defines equilibrium configuration and a remaining deformation.

linear elastic material with friction in the absence of external forces. Fig. 5.2 shows friction forces under three possible unload paths, as well as (negative) elastic forces. In the loading path, friction adds to the elastic force and saturates in all three cases. In the unloading path, friction gradually decreases and eventually changes sign and opposes elastic forces until an equilibrium is reached. With constant maximum stress, σ_{\max} , the persistent deformation for a linear elastic material with stiffness k saturates at a value $\epsilon_{\text{per}} = \frac{\sigma_{\max}}{k}$. Then, an effective way to tune the maximum friction stress is to specify a desired (maximum) persistent deformation, ϵ_{per} , and the maximum stress is set as $\sigma_{\max} = k \epsilon_{\text{per}}$.

The unload paths do not reach saturation in all cases though, only when the deformation in the loading path reaches a saturation deformation ϵ_{sat} . For practical reasons, we consider that Dahl's friction has saturated if it reaches a value $\|\sigma\| \geq c \sigma_{\max}$ (with, e.g., $c = 0.95$). Then, based on the analytic solution of first-order systems, the saturation deformation can be computed as $\epsilon_{\text{sat}} = \epsilon_{\text{per}} + \frac{\sigma_{\max}}{\alpha} \ln \frac{2}{1-c}$. In other words, an effective way to tune the friction saturation rate is to specify the saturation deformation, ϵ_{sat} , and then the friction saturation rate is automatically set as $\alpha = \frac{\sigma_{\max}}{\epsilon_{\text{sat}} - \epsilon_{\text{per}}} \ln \frac{2}{1-c}$.

5.2.4 Parameters for Estimation

The force-deformation data from Fig. 5.1 suggests that the magnitude of hysteresis grows as the deformation grows. A possible explanation for this growth is that higher strain produces higher inter-yarn normal forces, and hence higher inter-yarn friction. This effect may be captured by a strain-dependent maximum stress $\sigma_{\max}(\epsilon)$. Indeed, as shown in Fig. 5.3, we

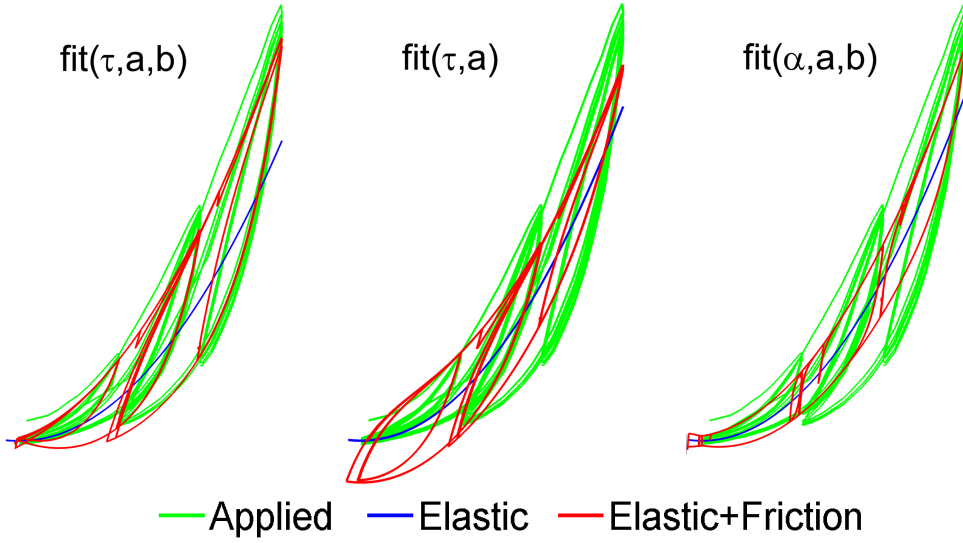


Fig. 5.3 Fits of Dahl's model to the data in Fig. 5.1-right using as parameterizations: (left) linear σ_{\max} and constant τ ; (center) constant σ_{\max} and constant τ ; (right) linear σ_{\max} and constant α .

found that a maximum stress linear w.r.t. strain, i.e., $\sigma_{\max} = a + b\varepsilon$, rendered an excellent match to the observed hysteresis data.

With linear maximum stress σ_{\max} and constant slope of the stress ratio α , the time constant τ of Dahl's model becomes strain-dependent. As shown in Fig. 5.3, a constant α does not capture well the profile of hysteresis across the range of deformations in the measured data, while a constant τ produces an excellent match.

With these two conclusions in mind, we generalize and reparameterize Dahl's model as follows:

$$\frac{d\sigma}{d\varepsilon} = \frac{\sigma_{\max}(\varepsilon) - s\sigma}{\tau}, \quad \text{with } \sigma_{\max}(\varepsilon) = a + b\varepsilon. \quad (5.2)$$

For the purpose of parameter estimation, given initial conditions ε_0 and σ_0 , the analytic solution can be approximated as

$$\sigma(\varepsilon) = s\sigma_{\max}(\varepsilon) + (\sigma_0 - s\sigma_{\max}(\varepsilon_0))e^{-s\frac{\varepsilon - \varepsilon_0}{\tau}}. \quad (5.3)$$

For constant maximum stress, the solution is accurate. For linear maximum stress, it omits a varying offset term bounded by $\|b\tau\|$. In practice, we found this term to be small, as well as a source of local minima for parameter optimization; therefore, we resorted to the approximate analytic solution for parameter estimation.

5.3 Estimation of Friction Parameters

Based on the understanding of Dahl’s friction model, in this section we propose simple procedures to estimate its parameters for different cloth deformation modes. Our procedures each include an acquisition step and a numerical optimization step. By defining simple principles for the acquisition step, we can leverage the analytic definition of friction forces, and thus largely simplify the optimization step. The friction model and the capture principles are general, and we demonstrate them by designing estimation procedures for stretch and bending friction. In both cases, the hardware used for acquisition is extremely simple.

5.3.1 General Rationale

Our parameter estimation procedures share the high-level strategy of many previous approaches for estimating elasticity parameters, namely: apply known position or force boundary conditions, measure the resulting static deformation, and estimate parameters that minimize force and/or position errors for the given conditions. Model estimation is cast as an optimization problem with friction parameters (and possibly elastic parameters) as unknowns.

Elastic forces are defined as explicit functions of elasticity parameters and deformations; therefore, it is straightforward to define an objective function based on force errors, and use optimization methods such as Newton’s or its multiple extensions, which make use of force gradients w.r.t. the parameters. Dahl’s friction forces from Eq. 5.2, however, are defined in differential form and the derivative is discontinuous. Friction forces are not simply a function of parameters and deformations. They depend on the path followed by those deformations; therefore, the application of Newton-type optimization methods is not straightforward. A possible solution would be to apply slow deformations to the captured material, sample forces and positions densely along the path of deformation, and locally linearize the friction force at each sample.

We propose instead friction estimation procedures that rely on simple capture setups and sparse data. We apply *anchor deformations* where friction forces are saturated, and then reverse the deformation gradually along a monotonic path. Friction forces depend on the path taken, not its duration; therefore, the friction forces at any position along a monotonic path starting at an anchor deformation are defined solely by the current deformation and the anchor deformation. Following this rationale, we can define friction forces as explicit functions of friction parameters and deformations, using the analytic expression in Eq. 5.3, with the strain and stress at an anchor deformation as initial conditions: $\epsilon_0 = \epsilon_{\text{anchor}}, \sigma_0 = \sigma_{\text{anchor}} = \pm \sigma_{\text{max}}(\epsilon_{\text{anchor}})$. Then it is straightforward to compute gradients of friction forces and apply




Fig. 5.4 Acquisition of stretch for a piece of cotton. From left to right: at rest, deformed under a weight of 6N, and with a persistent deformation back at rest.

Newton-type optimization methods.

With linear maximum stress σ_{\max} , the approximate expression of friction stress Eq. 5.3 corresponding to a measured deformation with strain ε_i , anchor deformation $\varepsilon_{\text{anchor},i}$, and strain-rate sign $s_i \in \{1, -1\}$ can be rewritten as

$$\sigma_i = s_i(a + b\varepsilon_i) - 2s_i(a + b\varepsilon_{\text{anchor},i})e^{-s_i \frac{\varepsilon_i - \varepsilon_{\text{anchor},i}}{\tau}}. \quad (5.4)$$

For practical purposes, in the expression above we assume that strain does not switch signs in the range $[\varepsilon_{\text{anchor},i}, \varepsilon_i]$. For stretch, we perform only tensile experiments, hence the assumption is trivially satisfied. For bending, we suggest performing separate experiments with positive and negative curvatures, and estimating separate parameters in each case.

5.3.2 Estimation of Stretch Friction

To estimate parameters of stretch friction (independently for warp and weft directions), we have devised a procedure to capture force-deformation pairs under quasi-uniform stretch. We place a 10cm×10cm cloth sample on a vertical plane and we hang various weights from the lower side as shown in Fig. 5.4. We perform incremental loading and unloading cycles, measuring deformations for weights $[0, 1, \dots, n, \dots, 1, 0]$ N, and n varying from 1 to 6 in our experiments. We repeat the full deformation sequence three times. For each measurement on a loading (resp. unloading) path, we define as anchor deformation the last measurement with a weight of 0N (resp. n N).

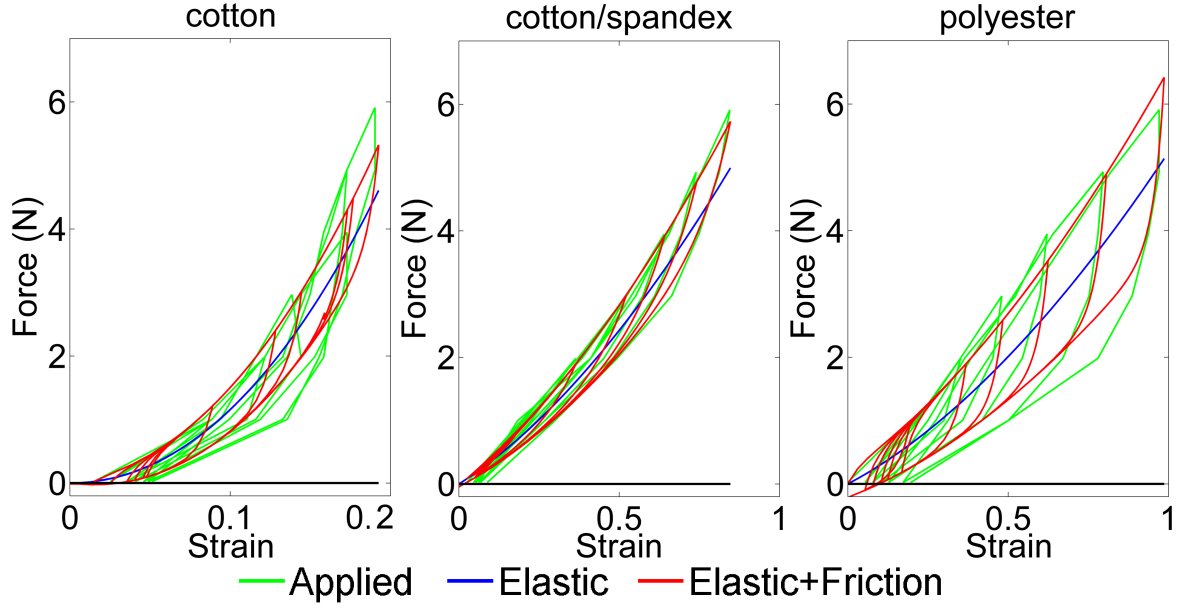


Fig. 5.5 Fitting of stretch friction parameters for cloth samples of cotton (left), cotton/spandex (center), and polyester (right).

For the i^{th} force-deformation measurement, we define a force error as the deviation from force equilibrium between the weight F_i (which includes the external weight, the clip, and the cloth itself), the elastic force $k \varepsilon_i A \nabla \varepsilon_i$, and the friction force $\sigma_i A \nabla \varepsilon_i$, i.e.,

$$f_i = F_i - k \varepsilon_i A \nabla \varepsilon_i - \sigma_i A \nabla \varepsilon_i. \quad (5.5)$$

k is the elastic stiffness, $A = L_0^2$ is the area of the undeformed cloth of width and length L_0 , and the friction stress σ_i is defined as in Eq. 5.4 based on the current and anchor strains. We use the Green-Lagrange strain $\varepsilon = \frac{1}{2} \left(\left(\frac{L}{L_0} \right)^2 - 1 \right)$, where L is the current length of the cloth. The strain gradient can be computed as $\nabla \varepsilon = \frac{L}{L_0^2}$.

Given all the force-deformation measurements, we jointly estimate elastic and friction parameters by solving the following least-squares problem:

$$(k, \tau, a, b) = \arg \min \sum_i f_i^2. \quad (5.6)$$

In addition to linear maximum friction stress $\sigma_{\max} = a + b \varepsilon$, we have also estimated a linear stiffness k . We have solved the non-linear optimization in Matlab using its built-in trust-region reflective algorithms. To avoid local minima, we initialize the parameters in the following way: first, we fit the stiffness k with no internal friction, then we fit σ_{\max} with very low τ (which

Material	ρ	k_1	k_2	a	b	τ
cotton	0.143	135.6	64.03	2.22	3.16	0.006
spandex	0.157	4.42	32.19	0.40	4.84	0.088
polyester	0.121	26.62	3.64	2.10	5.41	0.031

Table 5.1 Fitting results for the warp direction on the three samples.




Fig. 5.6 Acquisition of bending for a strip of doublecloth. From left to right: rolled around a cylinder, after removing the cylinder, and final spiral shape after unrolling.

approximates Eq. 5.4 to $\sigma_i = s_i(a + b\varepsilon_i)$, and then we fit τ .

Results and Discussion We have tested our stretch friction estimation procedure on cloth samples of cotton, cotton/spandex, and polyester. Fig. 5.5 compares the measured force-deformation data and the results of our parameter estimation for the three materials. As a reference for the fitting results, Table 5.1 shows density (ρ , in kg/m^2), stiffness ($k_1 + k_2\varepsilon$, in N/m), maximum friction stress ($a + b\varepsilon$, in N/m), and time constant (τ , dimensionless) for the warp direction on the three samples.

Thanks to the use of sparse data, our parameter estimation procedures work well with coarser capture resolution, and hence allow the use of less expensive hardware. For compliant cloth such as the one in the examples, inexpensive desktop hardware turns out to be sufficient. For stiffer cloth our desktop setup is not sufficient, but our estimation procedures could reduce the cost of instrumented capture setups and broaden the range of measurable materials.

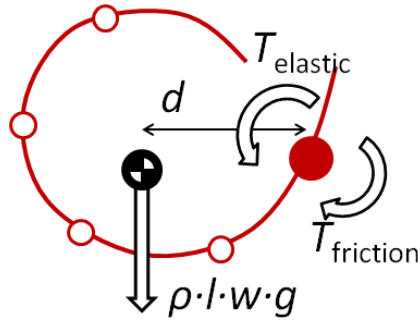


Fig. 5.7 Schematic of a sampled profile and evaluation of torques at sampled points.

5.3.3 Estimation of Bending Friction

The key aspect of our procedure to estimate parameters of bending friction is a simple way to produce anchor deformations with controlled strain. We roll a $10\text{cm} \times 2.5\text{cm}$ strip of cloth around a cylinder of known radius as shown in Fig. 5.6, and then we gradually unroll the cloth. We have used cylinders with radii of $[3, 4, 5, 6, 7, 7.5, 10, 16, 23, 27.5, 31.5, 50]\text{mm}$, and we have repeated the experiment with each cylinder three times. Before each rolling operation, we flatten the strip of cloth to remove possible persistent deformations.

For each experiment, we fit a cubic spline to the profile of the spiral, and we sample it at n points ($n = 5$ in our experiments), as shown in Fig. 5.7. For each of these samples, we define a torque error as the deviation from equilibrium between the gravity torque produced by the cloth's weight, the elastic bending torque, and the bending friction torque. With our simple procedure, the anchor deformation for all n samples in the same experiment is given by the same initial rolling radius. Next, we will define the elastic, friction, and gravity torques in detail.

The deformed cloth can be regarded as a ruled surface obtained by sweeping the spiral profile along the width of the cloth w . For the derivation of the elastic torque, we assume an infinitely fine regular triangulation with triangles of height dh . At each sample, we assume a hinge formed by edges of the triangulation running across the width of the cloth. Following the discrete-shells bending model described in Section 5.2.1, the strain of each edge is $\varepsilon = 3\frac{d\theta}{dh} = 3\kappa$, where κ is the curvature of the spiral profile at the sample. The elastic torque can then be computed as $T_{\text{elastic}} = k\varepsilon dA \nabla_{\theta}\varepsilon = 3kw\kappa$. The friction torque is defined as $T_{\text{friction}} = \sigma dA \nabla_{\theta}\varepsilon = \sigma w$, where σ is defined as in Eq. 5.4 based on the current and anchor strains. Finally, the gravity torque can be computed based on the density of the cloth ρ as $T_{\text{gravity}} = \rho l w g d$, where l is the length of the hanging piece of cloth, and d is the horizontal

Material	ρ	k_1	k_2	a	b	τ
felt	0.208	145.3	0.15	5.12	0.15	29.9
doublecloth	0.255	62.19	0.07	12.84	0.05	7.67
paper	0.075	479.7	0.47	0.04	0.64	10.2

Table 5.2 Fitting results for the bending on the three samples.

distance from its center of mass to the sample point on the spline.

We can now express a torque error for each spline sample as the deviation from equilibrium, $T_{\text{gravity}} - T_{\text{elastic}} - T_{\text{friction}}$. Substituting the expressions defined above, and factoring out the cloth width w , we obtain an expression of bending stress error for each sample,

$$f_i = \rho l_i d_i g - 3k \kappa_i - \sigma_i. \quad (5.7)$$

Note that the final expression is independent of the triangulation.

Same as for the estimation of stretch friction parameters, we formulate a least-squares problem with the form of Eq. 5.6. However, we use a cantilever experiment (Clapp et al., 1990) to estimate stiffness parameters, because in the spiral shapes of light materials the contribution of gravity is small, and there is not sufficient data to jointly estimate elasticity and friction parameters. To increase the optimization robustness, we first fit σ_{max} with very low τ , and then τ , before fitting both together. In addition, we remove sample points where gravity and elastic torques oppose each other, as they do not guarantee a monotonic deformation.

Results and Discussion We have tested our bending friction estimation procedure on materials with a moderately large bending-stiffness-to-density ratio. In particular, we report results for two cloth fabrics, felt and a synthetic doublecloth, as well as paper. Fig. 5.8 shows all captured spirals for one bending direction for the three materials, as well as the measured curvatures along the spiral profiles. As a reference for the fitting results, Table 5.2 shows density (ρ , in kg/m^2), stiffness ($k_1 + k_2 \epsilon$, in $(N \cdot m) \times 10^{-7}$), maximum stress ($a + b \epsilon$, in $N \times 10^{-4}$), and time constant (τ , in $1/m$), for the warp direction on the three samples.

Fig. 5.10, Fig. 5.11 and Fig. 5.12 show photographs of representative spirals for the three materials (for 4mm, 7.5mm, and 23mm initial radii). These figures show the estimated hysteresis curves with unload paths for the three initial radii, and the torque optimization residual at the sample points. Notice that the residual is overall low compared to the measured gravity-plus-elasticity torque.

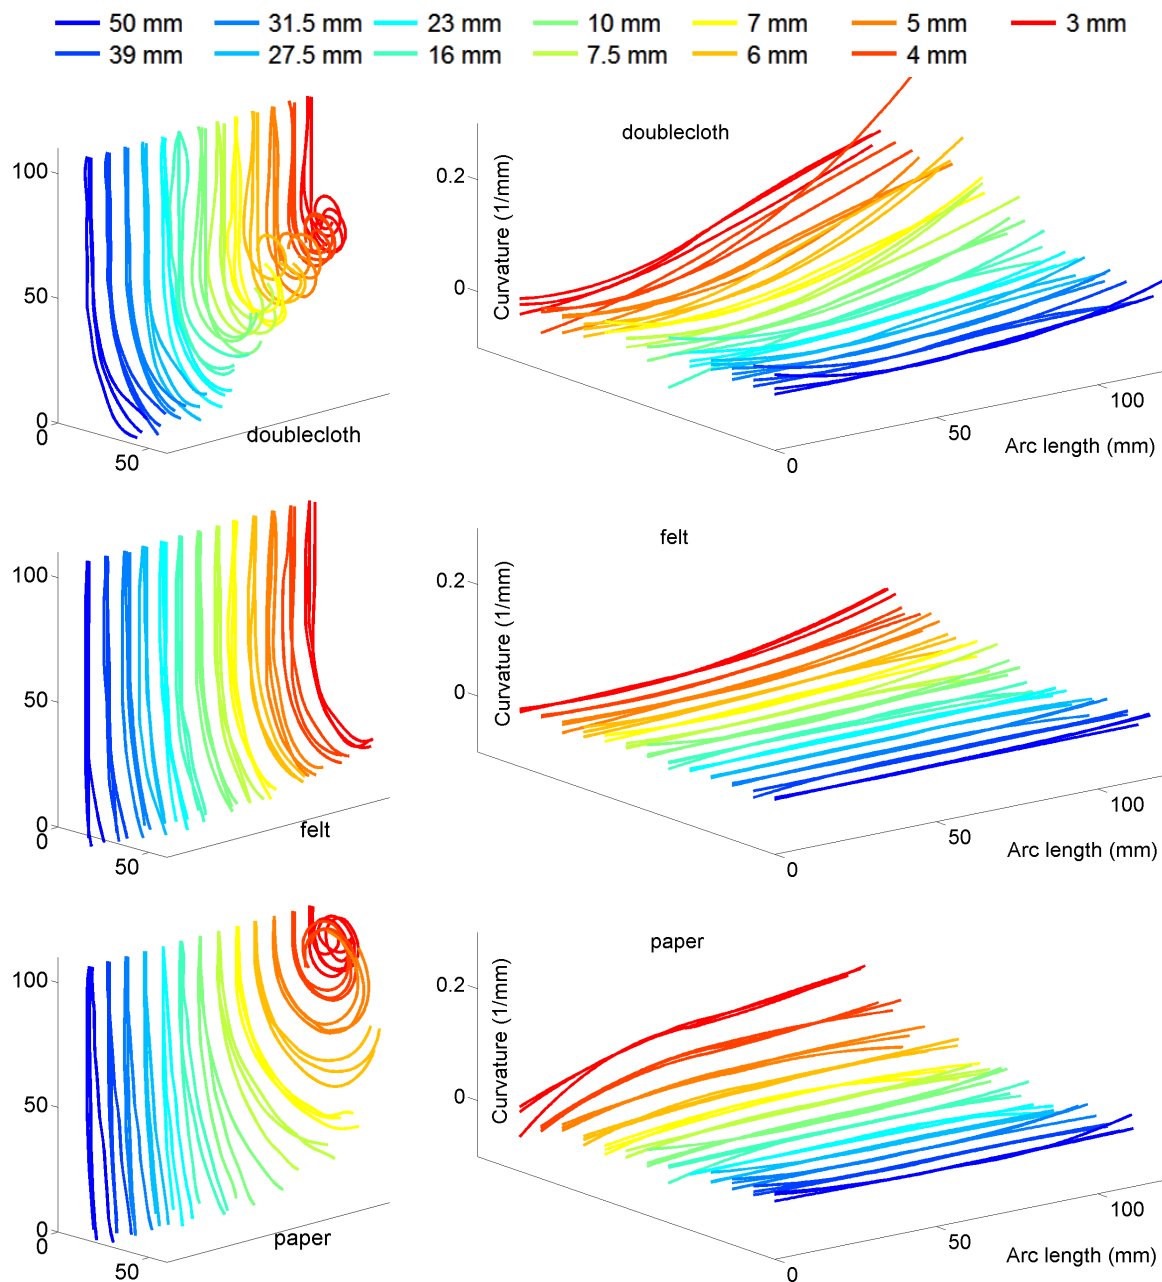


Fig. 5.8 Left: Spiral profiles for bending estimation; Right: Curvatures along the spiral profiles. The plots show data for 12 different initial radii and 3 samples per radius. The materials are, from top to bottom, synthetic doublecloth, felt, and paper.

Fig. 5.10, Fig. 5.11 and Fig. 5.12 also show final snapshots of unrolling simulations executed with the estimated parameters. The simulation results demonstrate characteristic features of both final spiral profiles and dynamic unrolling, with good qualitative matches to real-world results, and strong differences across materials. Improving accuracy would require

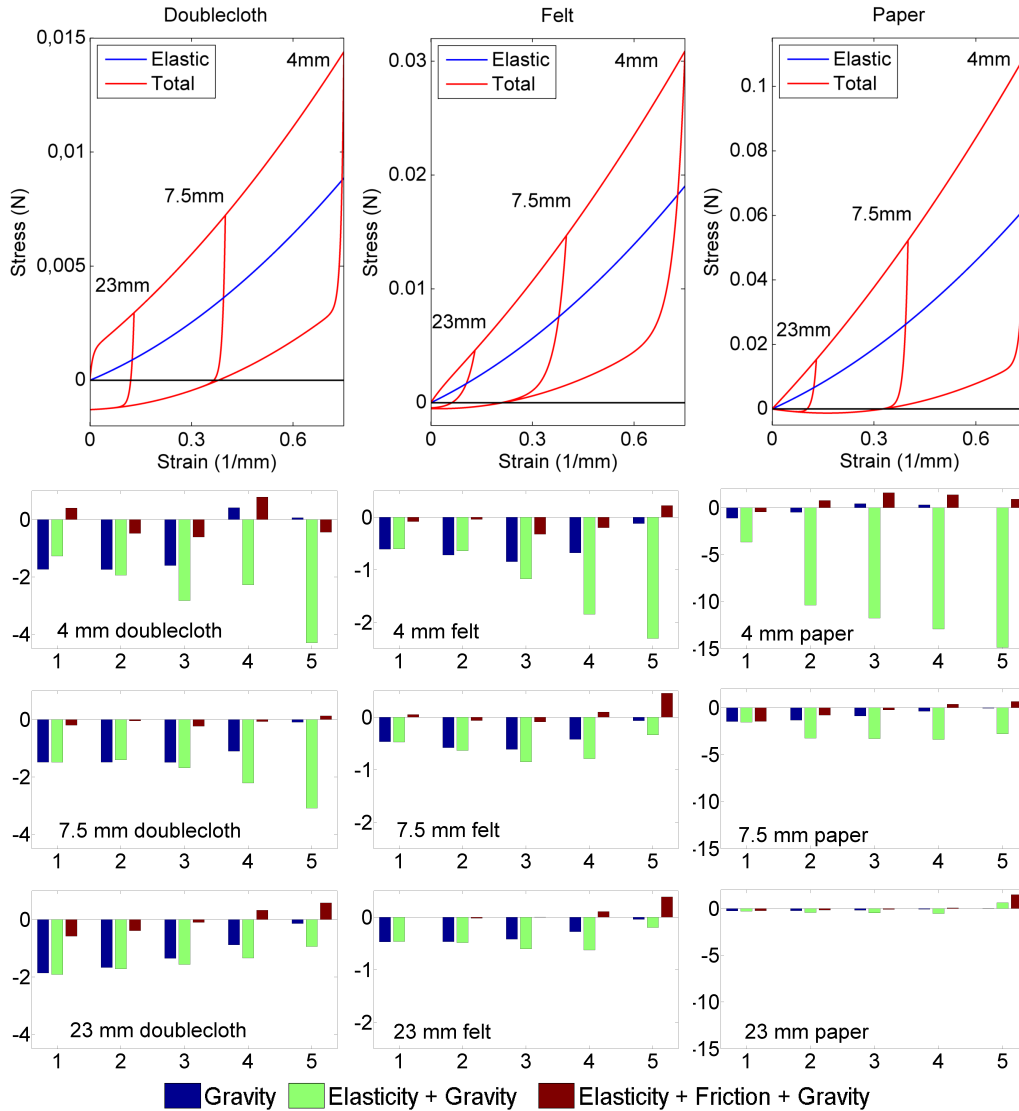


Fig. 5.9 Results and validation of bending friction estimation on doublecloth (left), felt (middle), and paper (right). Top row: Estimated hysteresis curves with unload paths for three initial radii (4mm, 7.5mm, and 23mm). Bottom 3 rows: Stress error residuals evaluated at samples on the spiral profiles (shown in Fig. 5.10, Fig. 5.11 and Fig. 5.11). The error in total stress is compared to the gravity and elastic stress.

not only more accurate internal friction, but also accurate damping and contact friction.

5.4 Dynamic Simulation of Internal Friction

Modeling internal friction in the context of dynamic simulation requires answering several challenges. First, we seek an easy integration into existing cloth simulators, which implies that




Fig. 5.10 Comparison of captured and simulated unrolling spirals for initial radii of 4mm, 7.5mm, and 23mm for doublecloth. The dots and numbers indicate samples for estimation residual evaluation in Fig. 5.9.

Dahl's model should be formulated as an explicit force expression (and possibly its derivatives w.r.t. positions and velocities) to be evaluated every time step. However, Dahl's model is expressed in differential form, hence a straightforward integration in dynamic simulators would require the inclusion of new state variables. Instead, we propose an algorithm that allows handling Dahl's friction force just as any other force expression.




Fig. 5.11 Comparison of captured and simulated unrolling spirals for initial radii of 4mm, 7.5mm, and 23mm for felt. The dots and numbers indicate samples for estimation residual evaluation in Fig. 5.9.

The second challenge concerns stability and performance of the simulation. The stiffness of Dahl's model is given directly by its differential expression in Eq. 5.2. At zero strain, the maximum stiffness is $\frac{2a}{\tau}$ (for $\sigma = -\sigma_{\max}$). Based on our estimated values, the effective stiffness of Dahl's friction can be in the order of 100 times larger than the elastic stiffness. Given this high stiffness, it is crucial to design a simulation algorithm for implicit integration,

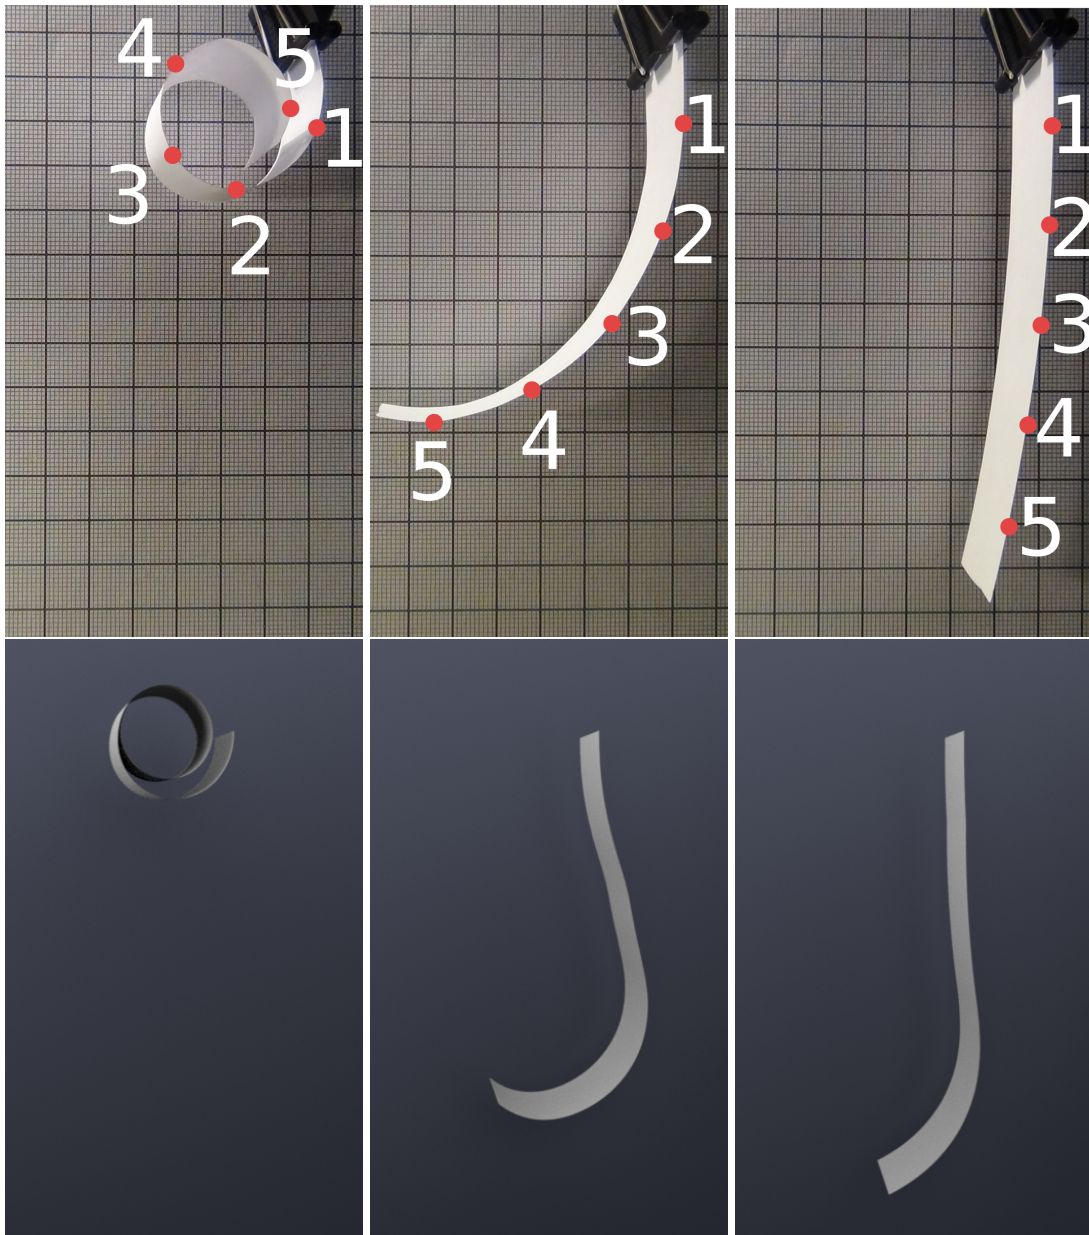


Fig. 5.12 Comparison of captured and simulated unrolling spirals for initial radii of 4mm, 7.5mm, and 23mm for paper. The dots and numbers indicate samples for estimation residual evaluation in Fig. 5.9.

stable under reasonably large time steps.

To formulate an explicit expression of friction force, we start by expressing the time derivative of the friction stress. Applying the chain rule to Eq. 5.2,

$$\dot{\sigma} = \frac{d\sigma}{d\varepsilon} \dot{\varepsilon} = \frac{1}{\tau} (\sigma_{\max}(\varepsilon) - s\sigma) \dot{\varepsilon}. \quad (5.8)$$

Dahl's model becomes an ODE, and the friction stress becomes part of the dynamic state of the cloth (together with positions and velocities). This growth of the state would complicate the inclusion of Dahl's model in existing cloth simulators, with the need of major changes to numerical solvers for implicit integration.

However, we note that the friction stress due to a strain element is a function of only the strain, strain rate, and stress at the element itself. Then, with backward Euler integration and a simple local linearization of Eq. 5.8 w.r.t ϵ , $\dot{\epsilon}$, and σ at every time step, we turn the (dynamic) friction stress into a linear function dependent only on the strain rate:

$$\sigma = \sigma_0 + \Delta t \frac{a + b \epsilon_0 - s_0 \sigma_0 + \Delta t b \dot{\epsilon}_0}{\tau + \Delta t s_0 \dot{\epsilon}_0} \dot{\epsilon}. \quad (5.9)$$

This function can be trivially evaluated just like any other force in the dynamics simulator, and its derivatives can be added to the system Jacobian for implicit integration. All values with a subindex 0 indicate values evaluated at the beginning of the time step (i.e., at the point of linearization), and Δt is the time step.

The simulation algorithm for implicit integration with our internal friction model is:

1. Evaluate friction stress values at the beginning of the time step using the analytic expression in Eq. 5.3, with the values from the previous time step as initial conditions.
2. Evaluate the Jacobians of friction stress according to Eq. 5.9.
3. Evaluate other forces and their Jacobians.
4. Formulate a linear system based on the implicit integration method of choice, and solve for cloth positions and velocities.

The algorithm above assumes one Newton iteration for (linearized) implicit integration, but it can be extended to iterate until convergence. In this case, friction stress values may only be re-evaluated once the solution has converged, not on every iteration of the solver, to ensure that friction follows the correct path. The same principle is applied as part of contact handling: friction stress values may only be re-evaluated once collisions are resolved, not inside iterative contact solvers.

As discussed in the next section, in our examples we found no difference between the time step requirements with and without friction, as the main source of time step restrictions turned out to be handling of self-collisions. Without collision handling, we found that implicit integration of friction forces enables stable integration under large time steps (e.g., more than 15ms for the cylinder demos in Figs. 5.13-5.16).




Fig. 5.13 Persistent wrinkles for cotton (left) and doublecloth (right) after a twist is quickly undone. The first and third images show wrinkles with friction and damping. The second and fourth images show more subtle wrinkles without damping, because vibrations partially eliminate the persistent deformations.

5.5 Results

We have executed multiple simulation tests to evaluate the effect and performance of our internal friction model and simulation algorithm on cloth simulations. From the purely mechanical point of view, our measurements indicate that the lack of friction may cause force and/or position deviations of up to 50%. From a visual point of view, force errors are not directly relevant, but we have observed that internal friction affects in multiple ways the formation and dynamics of folds and wrinkles. Next, we start with a discussion of friction effects on simple examples, where those effects are easier to isolate. Then, we discuss friction effects and the performance of our model and algorithm on larger character animation examples. We have used the damping model by Baraff and Witkin (1998), with damping values between one tenth and one hundredth of the stiffness at unit strain.

Persistent Deformations As already discussed, bending friction is the key phenomenon that produces the distinct persistent spiral profiles shown in Fig. 5.10, Fig. 5.11 and Fig. 5.12. In these benchmarks, the overall dynamic behavior is influenced by other factors such as damping and contact friction, but the explosiveness of the unrolling effect is highly determined by internal forces, both elastic and friction forces. Paper turned out to be the most challenging material to simulate in these examples, due to the fast motion and high stiffness-to-mass ratio.




Fig. 5.14 Wrinkles on a cotton cylinder with stretch friction after 1, 2, and 3 twist cycles. With internal friction, wrinkles are qualitatively similar, but arise at different locations after each twist cycle. Without friction, wrinkles are repetitive.

We executed the simulations with a time step of 0.25ms to resolve all self-collisions robustly.

We have also compared the formation of persistent deformations under different settings. Fig. 5.13 shows examples of a sheet of cloth tightly wrapped around a cylinder. The cloth undergoes a quick twist-untwist motion, which is similar to the motion of a sleeve when the wrist is rotated. The figure compares persistent deformations for cotton (using stretch friction only) and doublecloth (using bending friction only) with and without damping. Without damping, cloth exhibits vibrations when the twist motion ends, and as a result the persistent wrinkles are more subtle.

Preferred Folds and Wrinkles The benchmark shown in Fig. 5.15 is a clear example of the influence of internal friction on the existence of ‘preferred wrinkles’. A localized stretch deformation is applied to a piece of cloth, and the persistent deformation induces a preferred wrinkle when the cloth is compressed in the transversal direction. With our internal friction model we achieve a behavior that closely matches the real world. Under different angles and directions of compression, the simulation with internal friction tends to produce similar wrinkles, whereas the frictionless simulation exhibits highly varying wrinkles.

History-Dependent Folds Internal forces with friction are not conservative, hence the deformation under given boundary conditions is not defined as a simple energy minimum in the space of cloth configurations. In practice, this observation implies that the folds and wrinkles




Fig. 5.15 We produce a wrinkle on a piece of cotton (top). With internal friction (left), it becomes a ‘preferred wrinkle’ and arises repeatedly under diverse deformations. Without internal friction (right), folds and wrinkles show no clear similarity. Middle column shows real-world deformations for the same experiment.

of cloth may differ largely for the same boundary conditions, and depend on history, i.e., the path traveled to reach those boundary conditions. Fig. 5.14 shows cloth wrinkles on a cylinder of cotton (with stretch friction) after several cycles of twist motions. With purely elastic deformations (not shown), the wrinkles are repetitive, but with the addition of friction wrinkles are different after each twist cycle.

Settling of Wrinkles Friction resists motion, and sometimes it even stops it. In connection with this property, internal friction of cloth helps wrinkles settle faster when external motion stops. Fig. 5.16 compares wrinkles on a cylinder of cloth when a twist motion is undone, with and without friction. It shows snapshots 0.5sec apart, once the external motion is over. Without friction, wrinkles slide longer, whereas friction stops them more rapidly.

Character Clothing We have simulated two examples of characters wearing clothing. The pants shown in Fig. 5.17 are represented using a 20K triangle mesh, and are attached to the character at the waist and ankles. The character shown in Fig. 5.18 wears pants represented by a 17K triangle mesh, and again attached at the waist and ankles. The T-shirt is represented by a 12K triangle mesh and is not attached. In both examples, we used for the cloth model




Fig. 5.16 Images of a cylinder of cotton once a twist motion is undone, without friction (left) and with friction (right). The first (resp. third) and second (resp. fourth) snapshots are 0.5sec apart. Without friction, wrinkles slide along the surface, while friction stops them more rapidly.

parameters estimated from the cotton sample described in Section 5.3.

The pants in Fig. 5.17 show clear persistent deformations induced by friction. When the character lifts the knee, the fabric around the knee is strongly stretched. The persistent local stretch that remains when the knee is lowered produces a more irregular surface than the frictionless case. The visual appearance of this persistent deformation is different from thin localized wrinkles typically induced by plastic deformations.

The shirt of the character in Fig. 5.18 shows more stable folds and wrinkles with friction. The images highlight folds on the left side of the character's chest. Without friction, the folds disappear half-way through a jumping jack motion. Also, folds and wrinkles settle faster with friction when the character performs squats.

We used a time step of 1ms for the simulation in Fig. 5.17, and 0.5ms for the simulation in Fig. 5.18, both for the friction and frictionless cases. In all cases, the time step restrictions were imposed by robust constraint-based handling of collisions and self-collisions. Despite having a numerical stiffness up to 100 times larger than the frictionless case, our implicit integration algorithm provides stable simulation for the internal friction model. The simulations were computationally expensive, mainly due to the implicit integration of nonlinear elasticity and friction coupled with LCP-type contact handling. The character simulation took an average of 6sec per time step on a single core of a 2.67 GHz Intel Core i7 920 CPU with 12 GB RAM.

5.6 Discussion and Future Work

In this chapter, we have presented a model of internal friction that captures cloth hysteresis, whose parameters can be estimated through simple procedures, and can be easily integrated into existing cloth simulation frameworks. It constitutes a step forward toward higher realism in cloth animation and opens up several avenues of further work.

Our current friction model and estimation procedures present several limitations. The bending friction estimation procedure is not sufficiently sensitive for fabrics with very low bending-stiffness-to-density ratio, because the fabric's weight clearly dominates elastic and friction forces together. We anticipate, however, that an estimation procedure based on anchor deformations might be designed for such fabrics too. Shear hysteresis appears to be at least as large as stretch or bending hysteresis, and estimation of shear friction parameters would require the design of a suitable estimation procedure. The challenge stems from the difficulty to impose controlled shear deformations that can be employed as anchor deformations in our estimation procedures.

Our friction model is conveniently expressed as a scalar stress dependent on an individual scalar strain, and this formulation cannot capture potential cross-modal effects. However, such potential cross-modal effects are hardly understood to date. In general, with the availability of good internal friction models, it would be convenient to revisit the existing methods for the estimation of elasticity, and address the joint estimation of elasticity and friction. Based on our experiments, this need has become particularly evident for bending estimation on light materials. Following with a further understanding of internal friction, it would also be necessary to evaluate the effect of external factors such as moist on friction forces.



Fig. 5.17 Simulation of cotton pants with stretch friction (bottom) and without (top). When the character lifts the knee, the fabric is stretched, and internal friction produces persistent deformations in the stretched areas. They are visible particularly between the knee and ankle on the rightmost image.




Fig. 5.18 A character performing jumping jacks. Cloth is simulated without friction on the top, and with friction on the bottom. Internal friction produces more stable wrinkles and folds. This effect is particularly visible in the folds on the left side of the character’s chest. Without friction, these folds disappear half-way through the motion, but they remain under friction.

Chapter 6

Conclusion

The main goal of this thesis is to develop models and techniques to improve the realism of deformable object simulation. We achieve this goal by developing novel deformation models that capture complex behaviors, such as nonlinearities, heterogeneity and internal friction, design measurement systems and estimate model parameters that reproduce the observed behaviors. In the following sections we provide a summary of the results of this thesis and possible future research directions.

6.1 Summary of results

In Chapter 3 we presented a novel capture system for cloth and an estimation pipeline that outputs model parameters that produce realistic simulations based on real-world measurements. In contrast to standard textile testing and capture systems, our measurement setup produces complete 3D geometry and force information, offering a detailed view of the behavior of cloth. We estimate parameters for three different cloth models, namely Soft Constraints model, StVK model and Springs model. We make these models nonlinear by using strain-dependent stiffness parameters and compare them in terms of quality of the obtained estimations. Our results point out that the Springs model exhibits the worst fitting quality, probably due to the coupling of membrane and bending deformation modes, while the Soft Constraints and StVK models produce similar, good quality fits.

In Chapter 3, we observed that Poisson and cross-modal stiffening effects cannot be captured with the proposed models. Consequently, in Chapter 4, we presented a novel deformation model based on additive energy density terms, that is able to capture these two effects. In addition, by using an additive energy model we can guarantee integrability of elastic forces, which could not be guaranteed with the deformation models in Chapter 3. This model also fa-

cilitates parameter estimation by minimizing the impact of local minima problems, since it is suitable for incremental parameter estimation strategies. We demonstrate the applicability of our novel deformation model in three different simulation and parameter estimation scenarios: modeling of complex nonlinearities in cloth, non-rigid registration of internal heterogeneous human anatomy and extremely nonlinear finger skin deformation.

In Chapter 5, we presented an internal friction model that is able to produce the typical hysteresis behavior of cloth. We estimate parameters for this friction model using a simple and inexpensive measurement setup and local optimization algorithms, and show that this model can be easily integrated into existing simulation algorithms. Our model is based on Dahl's friction model, but we apply a reparameterization that produces a better fit to real-world measurements. Finally, in contrast to previous work, we provide an analysis of the visual impact of internal friction in cloth simulation.

6.2 Future Work

The capture system described in Chapter 3 provides very detailed information, but it has several limitations. For bending measurement, manual intervention is still required, which reduces the precision of the measurement system. In addition, cloth samples tend to curl up at edges, which may create errors and biases in the estimation. There are still many open research questions in measurement and capture of cloth: exploration of larger strain space, maybe including compression; capture and fitting of dynamic properties; and improvement of shear measurement.

Several limitations arise from current estimation procedures. A common step in the estimation pipelines described in Chapter 3 and Chapter 4 is the resolution of the quasi-static problem. We have observed that it is difficult to navigate the energy space towards a minimum-energy configuration, which has made this step the bottleneck in the optimization performed for parameter estimation. We would greatly benefit from faster quasi-static solves since it would allow us to use other optimization algorithms, maybe with more error function evaluations but less prone to falling into local minima.

Another limitation of parameter estimation is that we rely on gradient-based local optimization algorithms, which may suffer local minima problems. Our progressive estimation process, which incrementally increases the complexity of the parameter space, helps us minimize this problem, but there are no absolute guarantees. Also, it would be interesting to investigate simpler ways to enforce convexity of the energy addends from Chapter 4.

An interesting extension to our general hyperelasticity model presented in Chapter 4 would

be to smartly choose other energy addends, such that the total number of energy addends is reduced, the number of parameters is minimized and the estimation process is eased. An interesting approach could be to apply Principal Component Analysis (PCA) in force-deformation data, look for the strain components that could have the most impact in the error function and define energy addends that depend on those strain components.

Our internal friction model (Chapter 5) presents a serious limitation in that it takes as input a single scalar strain value, which prevents it from capturing potential cross-modal effects. After observing the results in Section 5.5, it becomes clear that current elasticity estimation procedures can be improved by taking into account internal friction during the estimation process.

Finally, the internal friction estimation procedures in Section 5.3 are far from perfect. Estimation of bending friction turns out to be extremely difficult with materials showing low bending-stiffness-to-density ratio, due to the fabric's weight dominating elastic and friction forces together. Also, there is no clear solution to estimate shear internal friction due to the difficulty of designing measurement setups that can impose controlled deformations that serve as anchor deformations.

References

- Atcheson, B., Ihrke, I., Heidrich, W., Tevs, A., Bradley, D., Magnor, M., and Seidel, H.-P. (2008). Time-resolved 3d capture of non-stationary gas flows. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)*, 27(5):132.
- Baraff, D. (1991). Coping with friction for non-penetrating rigid body simulation. In *Computer Graphics (Proceedings of SIGGRAPH 91)*, pages 31–40.
- Baraff, D. and Witkin, A. (1998). Large steps in cloth simulation. In *Proc. of ACM SIGGRAPH*, pages 43–54.
- Barbič, J. and James, D. (2005). Real-time subspace integration for St. Venant-Kirchhoff deformable models. *ACM Trans. Graph.*, 24(3):982–990.
- Bathe, K. (2006). *Finite Element Procedures*. Prentice Hall.
- Becker, M. and Teschner, M. (2007). Robust and efficient estimation of elasticity parameters using the linear finite element method. In *SimVis*, pages 15–28.
- Bergou, M., Mathur, S., Wardetzky, M., and Grinspun, E. (2007). Tracks: Toward directable thin shells. *ACM Transactions on Graphics*, 26(3):50:1–50:10.
- Bhat, K. S., Twigg, C. D., Hodgins, J. K., Khosla, P. K., Popović, Z., and Seitz, S. M. (2003). Estimating cloth simulation parameters from video. In *Proc. ACM SIGGRAPH/Eurographics SCA*, pages 37–51.
- Bickel, B., Bäcker, M., Otaduy, M. A., Matusik, W., Pfister, H., and Gross, M. (2009). Capture and modeling of non-linear heterogeneous soft tissue. *ACM Trans. Graph.*, 28(3):89:1–89:9.
- Bliman, P. A. and Sorine, M. (1991). Friction modelling by hysteresis operators. application to dahl, stiction and stribeck effects. *Proc. Models of Hysteresis*.
- Boisse, P., Borr, M., Buet, K., and Cherouat, A. (1997). Finite element simulations of textile composite forming including the biaxial fabric behaviour. *Composites Part B: Engineering*, 28(4):453–464.
- Bonet, J. and Wood, R. D. (1997). *Nonlinear Continuum Mechanics for Finite Element Analysis*. Cambridge University Press.
- Bouman, K., Xiao, B., Battaglia, P., and Freeman, W. (2013). Estimating the material properties of fabric from video. In *Computer Vision (ICCV), 2013 IEEE International Conference on*, pages 1984–1991.

- Bradley, D., Boubekur, T., and Heidrich, W. (2008a). Accurate multi-view reconstruction using robust binocular stereo and surface meshing. In *Proc. CVPR*.
- Bradley, D., Heidrich, W., Popa, T., and Sheffer, A. (2010). High resolution passive facial performance capture. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 29(4):41:1–41:10.
- Bradley, D., Popa, T., Sheffer, A., Heidrich, W., and Boubekur, T. (2008b). Markerless garment capture. *ACM Trans. Graph. (Proc. of SIGGRAPH)*, 27(3):99:1–99:9.
- Breen, D., House, D., and Wozny, M. (1994). Predicting the drape of woven cloth using interacting particles. In *Proc. of ACM SIGGRAPH*, pages 365–372.
- Bridson, R., Fedkiw, R., and Anderson, J. (2002). Robust treatment of collisions, contact and friction for cloth animation. In *Proc. of ACM SIGGRAPH*, pages 594–603.
- Bridson, R., Marino, S., and Fedkiw, R. (2003). Simulation of clothing with folds and wrinkles. In *Proc. ACM SIGGRAPH/Eurographics SCA*, pages 28–36.
- Chen, Z., Feng, R., and Wang, H. (2013). Modeling friction and air effects between cloth and deformable bodies. *ACM Transactions on Graphics*, 32(4).
- Choi, K.-J. and Ko, H.-S. (2002). Stable but responsive cloth. In *Proc. of ACM SIGGRAPH*, pages 604–611.
- Clapp, T. G., Peng, H., Ghosh, T. K., and Eischen, J. W. (1990). Indirect measurement of the moment-curvature relationship for fabrics. *Textile Research J.*, 60(9).
- Cottle, R., Pang, J., and Stone, R. (1992). *The Linear Complementarity Problem*. Academic Press.
- Culpin, M. F. (1979). The shearing of fabrics: A novel approach. *J. Textile Institute*, 70(3):81–88.
- Dahl, P. R. (1968). A solid friction model. Technical report, The Aerospace Corporation.
- De, S., Lim, Y.-J., Manivannan, M., and Srinivasan, M. A. (2006). Physically realistic virtual surgery using the point-associated finite field (paff) approach. *Presence: Teleoper. Virtual Environ.*, 15(3):294–308.
- Eberhardt, B., Weber, A., and Straßer, W. (1996). A fast, flexible, particle-system model for cloth draping. *IEEE Computer Graphics and Applications*, 16(5):52–59.
- English, E. and Bridson, R. (2008). Animating developable surfaces using nonconforming elements. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 27(3):66:1–66:5.
- Etzmuß, O., Keckeisen, M., and Straßer, W. (2003). A Fast Finite Element Solution for Cloth Modelling. In *Proc. Pacific Graphics*, pages 244–251.
- Fiala, M. (2005). Artag, a fiducial marker system using digital techniques. In *IEEE CVPR*, pages 590–596.

- Gascon, J., Espadero, J. M., Perez, A. G., Torres, R., and Otaduy, M. A. (2013). Fast deformation of volume data using tetrahedral mesh rasterization. In *Proceedings of the 12th ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 181–185.
- Goldenthal, R., Harmon, D., Fattal, R., Bercovier, M., and Grinspun, E. (2007). Efficient simulation of inextensible cloth. *ACM Trans. Graph. (Proc. of SIGGRAPH)*, 26(3):49:1–49:7.
- Grinspun, E., Hirani, A., Desbrun, M., and Schröder, P. (2003). Discrete shells. In *Proc. ACM SIGGRAPH/Eurographics SCA*, pages 62–67.
- Han, L., Hipwell, J., Mertzaniidou, T., Carter, T., Modat, M., Ourselin, S., and Hawkes, D. (2011). A hybrid FEM-based method for aligning prone and supine images for image guided breast surgery. In *Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on*, pages 1239–1242.
- Han, W. and Reddy, B. D. (2012). *Plasticity: mathematical theory and numerical analysis*, volume 9. Springer.
- Hernandez, F., Cirio, G., Perez, A. G., and Otaduy, M. A. (2013). Anisotropic strain limiting. In *Proc. of Congreso Español de Informática Gráfica*.
- Hiller, J. and Lipson, H. (2012). Dynamic simulation of soft heterogeneous objects. *arXiv preprint arXiv:1212.2845*.
- Hughes, T. (2000). *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis*. Dover Publications.
- Irving, G., Schroeder, C., and Fedkiw, R. (2007). Volume conserving finite element simulations of deformable models. *Proc. of ACM SIGGRAPH*.
- Irving, G., Teran, J., and Fedkiw, R. (2004). Invertible finite elements for robust simulation of large deformation. In *Proc. ACM SIGGRAPH/Eurographics SCA*, pages 131–140.
- Kaldor, J. M., James, D. L., and Marschner, S. (2008). Simulating knitted cloth at the yarn level. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 27(3):65:1–65:9.
- Kawabata, S. (1980). The standardization and analysis of hand evaluation. *Textile Machinery Soc. Japan*.
- Kim, B.-C., Oh, S., and Wohn, K. (2011). Persistent wrinkles and folds of clothes. *International Journal of Virtual Reality*, 10(1):61–66.
- Kunitomo, S., Nakamura, S., and Morishima, S. (2010). Optimization of cloth simulation parameters by considering static and dynamic features. In *ACM SIGGRAPH Posters*, page 15:1.
- Lahey, T. J. (2002). Modeling hysteresis in the bending of fabrics. Master’s thesis, University of Waterloo.
- Lee, S.-H., Sifakis, E., and Terzopoulos, D. (2009). Comprehensive biomechanical modeling and simulation of the upper body. *ACM Transactions on Graphics (TOG)*, 28(4):99.

- Lim, C., Shim, V., and Ng, Y. (2003). Finite-element modeling of the ballistic impact of fabric armor. *International Journal of Impact Engineering*, 28(1):13–31.
- Lin, S., Narayan, R., and Lee, Y.-S. (2008). Heterogeneous deformable modeling and topology modification for surgical cutting simulation with haptic interfaces. *Computer-Aided Design and Applications*, 5(6):877–888.
- Mattes, D., Haynor, D. R., Vesselle, H., Lewellen, T. K., and Eubank, W. (2003). Pet-ct image registration in the chest using free-form deformations. *Medical Imaging, IEEE Transactions on*, 22(1):120–128.
- Miguel, E., Bradley, D., Thomaszewski, B., Bickel, B., Matusik, W., Otaduy, M. A., and Marschner, S. (2012). Data-driven estimation of cloth simulation models. *Computer Graphics Forum (Proc. of Eurographics)*, 31(2).
- Miguel, E., Tamstorf, R., Bradley, D., Schwartzman, S. C., Thomaszewski, B., Bickel, B., Matusik, W., Marschner, S., and Otaduy, M. A. (2013). Modeling and estimation of internal friction in cloth. *ACM Trans. on Graphics (Proc. of ACM SIGGRAPH Asia)*, 32(6).
- Müller, M. and Gross, M. (2004). Interactive virtual materials. *Proc. of Graphics Interface*.
- Narain, R., Pfaff, T., and O’Brien, J. F. (2013). Folding and crumpling adaptive sheets. *ACM Transactions on Graphics*, 32(4).
- Narain, R., Samii, A., and O’Brien, J. F. (2012). Adaptive anisotropic remeshing for cloth simulation. *ACM Trans. Graph.*, 31(6):152:1–152:10.
- Ngo-Ngoc, C. and Boivin, S. (2004). Nonlinear cloth simulation. Technical report, INRIA.
- Ogden, R. W. (1997). *Non-Linear Elastic Deformations*. Courier Dover Publications.
- Pabst, S., Krzywinski, S., Schenk, A., and Thomaszewski, B. (2008). Seams and bending in cloth simulation. In *Proc. of VRIPHYS*.
- Padthe, A., Drincic, B., Oh, J., Rizos, D., Fassois, S., and Bernstein, D. (2008). Duhem modeling of friction-induced hysteresis. *Control Systems, IEEE*, 28(5):90–107.
- Pai, D. K., van den Doel, K., James, D. L., Lang, J., Lloyd, J. E., Richmond, J. L., and Yau, S. H. (2001). Scanning physical interaction behavior of 3d objects. In *Proceedings of ACM SIGGRAPH*, pages 87–96.
- Parsons, E. M., Weerasooriya, T., Sarva, S., and Socrate, S. (2010). Impact of woven fabric: Experiments and mesostructure-based continuum-level simulations. *Journal of the Mechanics and Physics of Solids*, 58(11):1995–2021.
- Patterson, T., Mitchell, N., and Sifakis, E. (2012). Simulation of complex nonlinear elastic bodies using lattice deformer. *ACM Transactions on Graphics*.
- Perez, A. G., Cirio, G., Hernandez, F., Garre, C., and Otaduy, M. A. (2013). Strain limiting for soft finger contact simulation. *Proc. of World Haptics Conference*.
- Provot, X. (1995). Deformation constraints in a mass-spring model to describe rigid cloth behavior. In *Graphics Interface '95*, pages 147–154.

- Rhee, T., Lewis, J., Neumann, U., and Nayak, K. S. (2011). Scan-based volume animation driven by locally adaptive articulated registrations. *IEEE Transactions on Visualization and Computer Graphics*, 17(3):368–379.
- Rueckert, D., Sonoda, L. I., Hayes, C., Hill, D. L. G., Leach, M. O., and Hawkes, D. (1999). Nonrigid registration using free-form deformations: application to breast mr images. *Medical Imaging, IEEE Transactions on*, 18(8):712–721.
- Sherburn, M. (2007). *Geometric and Mechanical Modelling of Textiles*. PhD thesis, University of Nottingham.
- Simo, J. and Hughes, T. (2000). *Computational Inelasticity*. Interdisciplinary Applied Mathematics. Springer New York.
- Spillmann, J. and Teschner, M. (2009). Cosserat nets. *Visualization and Computer Graphics, IEEE Transactions on*, 15(2):325–338.
- Stoll, C., Gall, J., de Aguiar, E., Thrun, S., and Theobalt, C. (2010). Video-based reconstruction of animatable human characters. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)*, 29(5).
- Terzopoulos, D. and Fleischer, K. (1988). Modeling inelastic deformation: viscoelasticity, plasticity, fracture. In *ACM Siggraph Computer Graphics*, volume 22, pages 269–278. ACM.
- Terzopoulos, D., Platt, J., Barr, A., and Fleischer, K. (1987). Elastically deformable models. In *Proc. of ACM SIGGRAPH*, pages 205–214.
- Thomaszewski, B., Pabst, S., and Straßer, W. (2009). Continuum-based strain limiting. *Computer Graphics Forum (Proc. of Eurographics)*, 28:569–576.
- Volino, P. and Magnenat-Thalmann, N. (2007). Stop-and-go cloth draping. *Vis. Comput.*, 23(9):669–677.
- Volino, P., Magnenat-Thalmann, N., and Faure, F. (2009). A simple approach to nonlinear tensile stiffness for accurate cloth simulation. *ACM Trans. Graph.*, 28(4).
- Wang, H., O’Brien, J., and Ramamoorthi, R. (2010). Multi-resolution isotropic strain limiting. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)*, 29:156:1–156:10.
- Wang, H., O’Brien, J., and Ramamoorthi, R. (2011a). Data-driven elastic models for cloth: Modeling and measurement. *ACM Transactions on Graphics*.
- Wang, H., Ramamoorthi, R., and O’Brien, J. (2011b). Data-driven elastic models for cloth: Modeling and measurement. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 30(4):71.
- White, R., Crane, K., and Forsyth, D. A. (2007). Capturing and animating occluded cloth. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 26(3).
- Zhu, Y. and Bridson, R. (2005). Animating sand as a fluid. *ACM Transactions on Graphics*, 24(3):965–972.

Appendix A

Parameter Estimation Results for Cloth Models

This appendix contains additional results, in this order:

- Fitting results for Sample #4 using the Soft Constraints model.
- Fitting results for Sample #4 using the St. V-K model.
- Fitting results for Sample #4 using the Spring model.
- Fitting results for Sample #12 using the Soft Constraints model.
- Fitting results for Sample #12 using the St. V-K model.
- Fitting results for Sample #12 using the Spring model.
- Fitting results for Sample #14 using the Soft Constraints model.
- Fitting results for Sample #14 using the linear Soft Constraints model.
- Fitting results for Sample #14 using the isotropic Soft Constraints model.
- Fitting results for Sample #14 using the linear isotropic Soft Constraints model.
- Fitting results for Sample #14 using the St. V-K model.
- Fitting results for Sample #14 using the Spring model.
- Fitting results for Sample #18 using the Soft Constraints model.

- Fitting results for Sample #18 using the St. V-K model.
- Fitting results for Sample #18 using the Spring model.
- Stress-strain plots using the Soft Constraints model.
- Stress-strain plots using the St. V-K model.
- Stress-strain plots using the Spring model.
- Fitting residuals for Sample #4.
- Fitting residuals for Sample #12.
- Fitting residuals for Sample #14.
- Fitting residuals for Sample #18.
- Evaluation results for Sample #4.2 using non-linear orthotropic models fit to #4.
- Evaluation results for Sample #12.2 using non-linear orthotropic models fit to #12.
- Evaluation results for Sample #14.2 using non-linear orthotropic models fit to #14.
- Evaluation results for Sample #18.2 using non-linear orthotropic models fit to #18.
- Comparison of evaluations between fitting and test samples of each fabric.

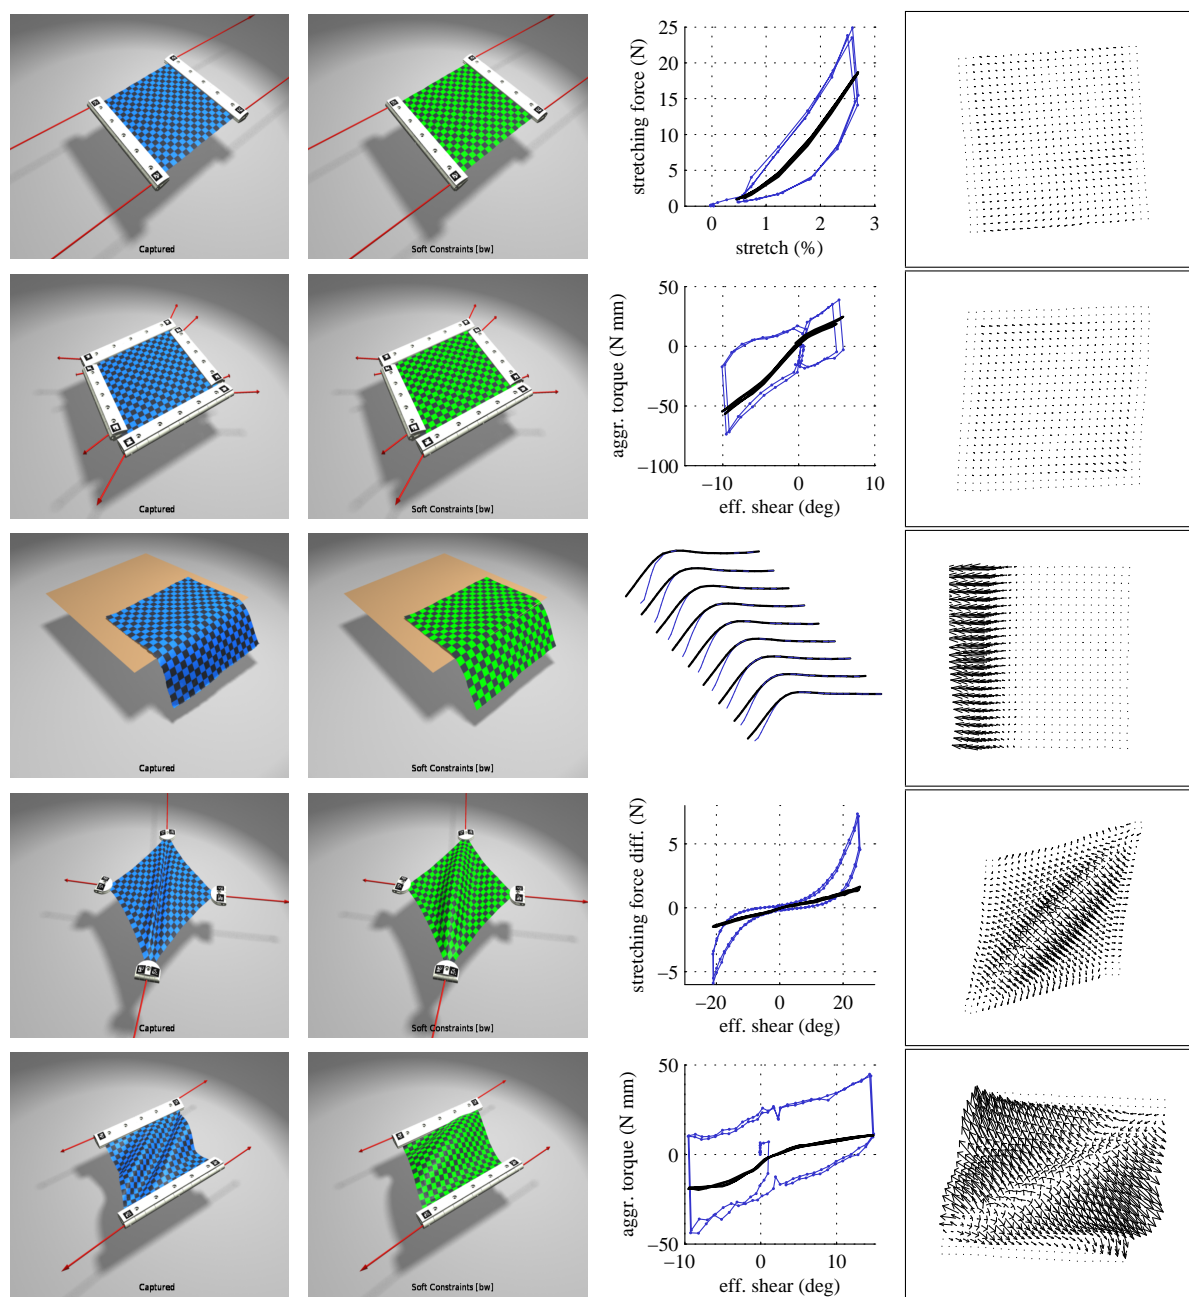


Fig. A.1 Fitting results for Sample #4, Soft Constraints model. Top to bottom; Stretch-X, Simple shear, Bend-X, Corner pull, Complex shear. Left to right: captured geometry, equilibrium of fitted model, force comparison (thin line: measurement; thick line: model), position residual (vertex position minus corresponding measured position, magnified 5x).

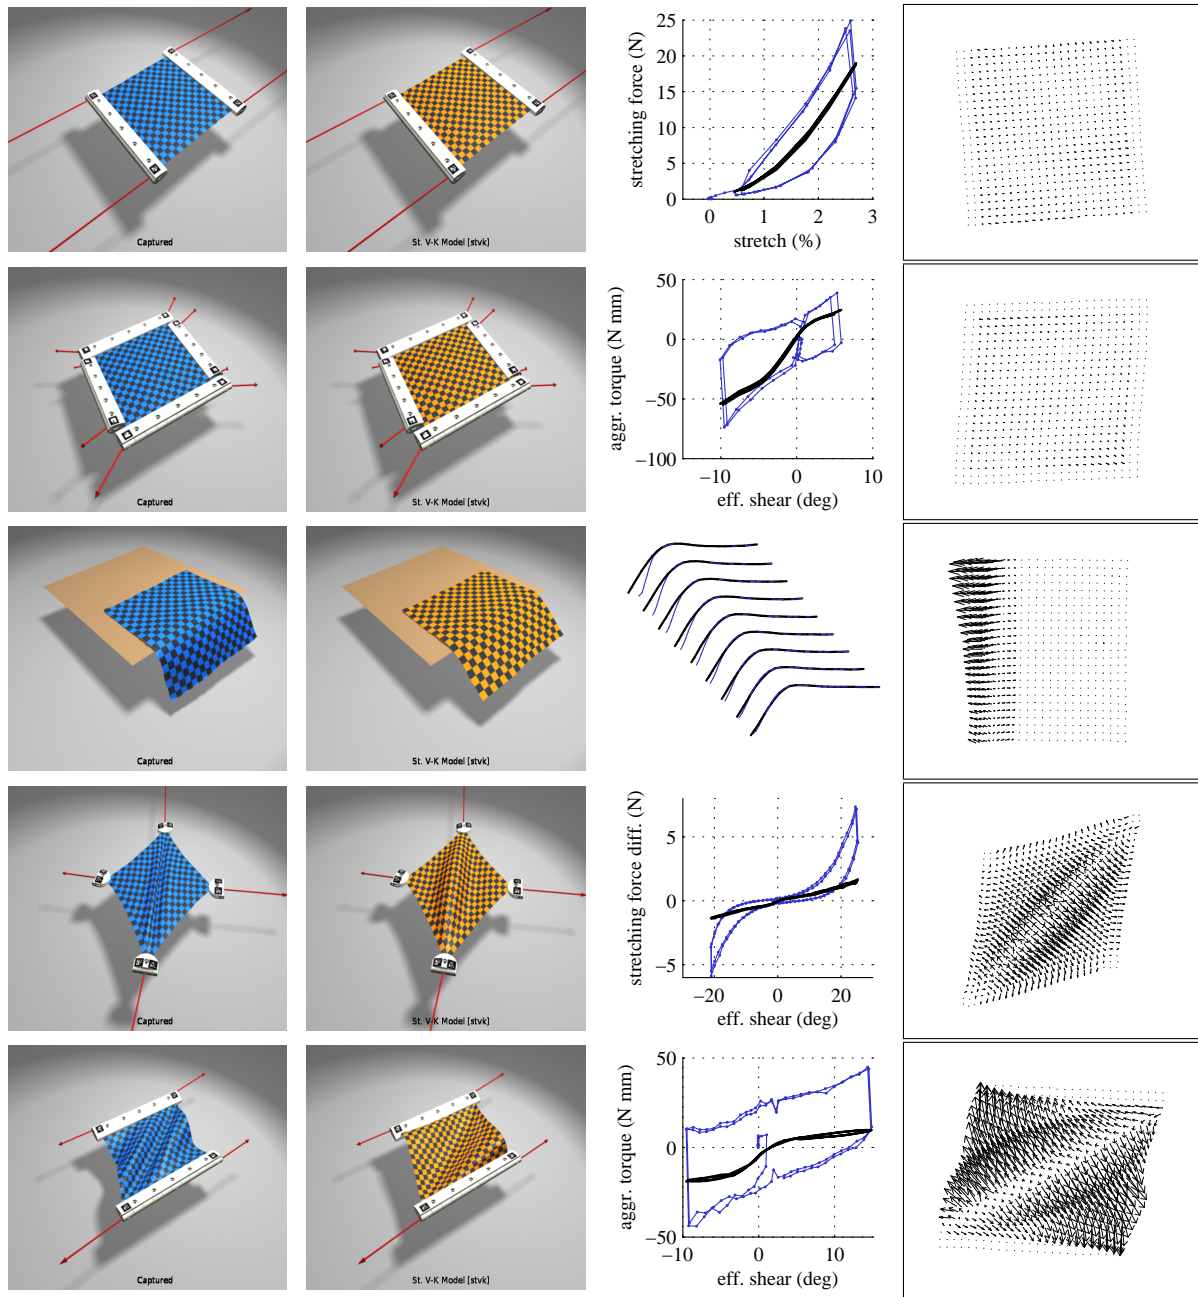


Fig. A.2 Fitting results for Sample #4, St. V-K model. Top to bottom; Stretch-X, Simple shear, Bend-X, Corner pull, Complex shear. Left to right: captured geometry, equilibrium of fitted model, force comparison (thin line: measurement; thick line: model), position residual (vertex position minus corresponding measured position, magnified 5x).

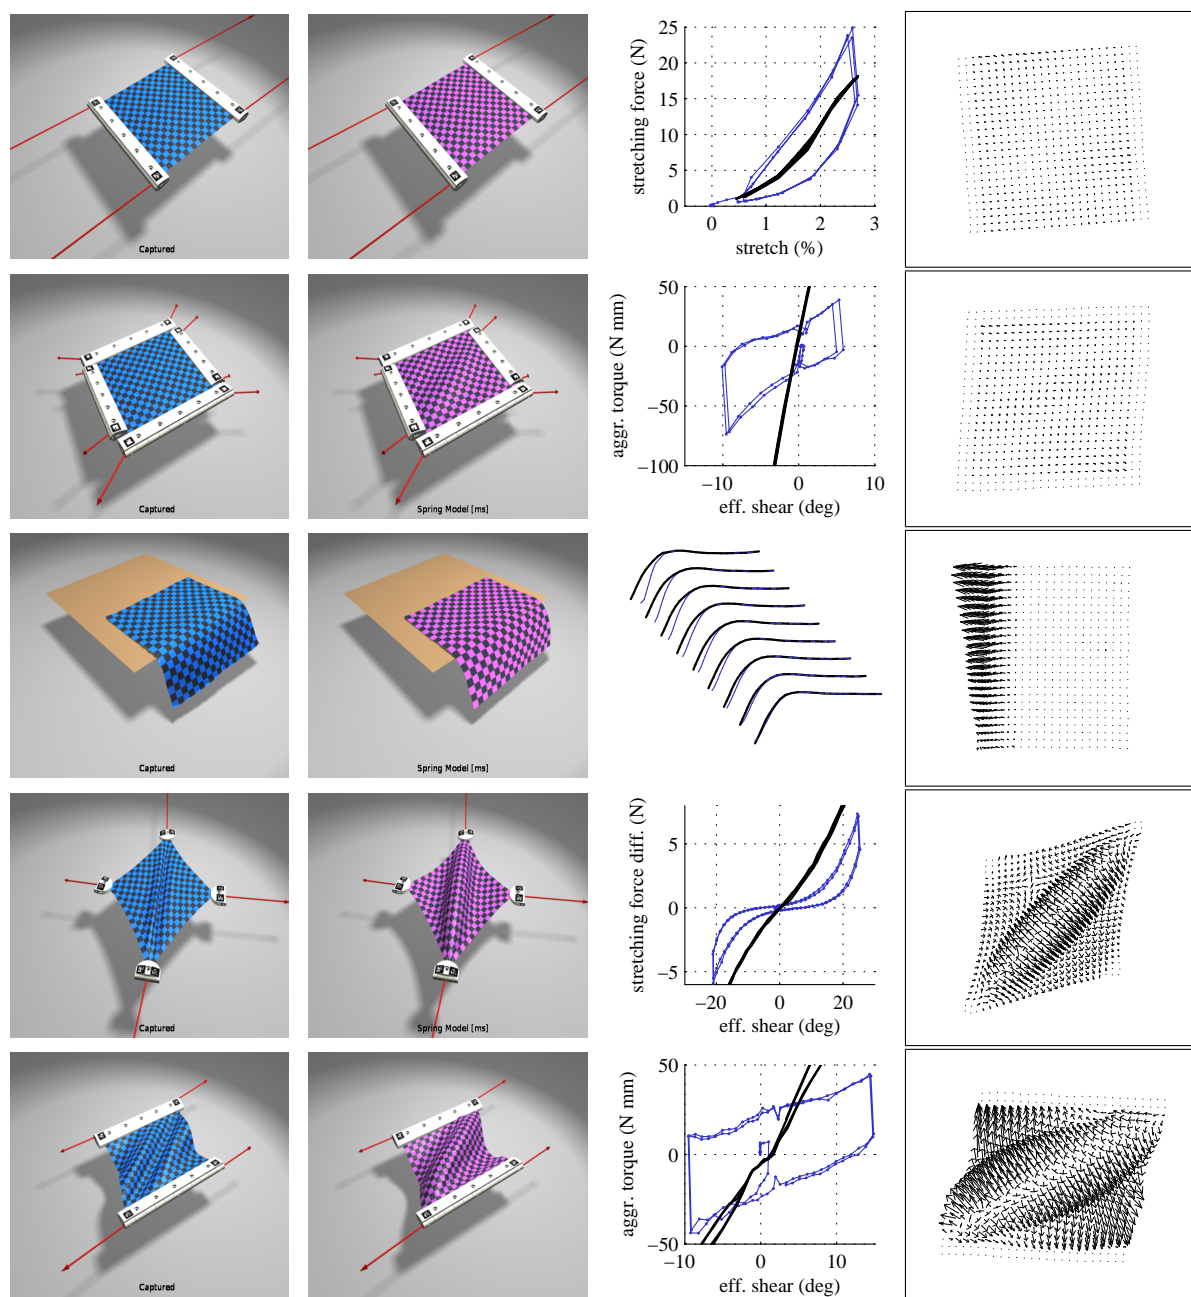


Fig. A.3 Fitting results for Sample #4, Spring model. Top to bottom; Stretch-X, Simple shear, Bend-X, Corner pull, Complex shear. Left to right: captured geometry, equilibrium of fitted model, force comparison (thin line: measurement; thick line: model), position residual (vertex position minus corresponding measured position, magnified 5x).

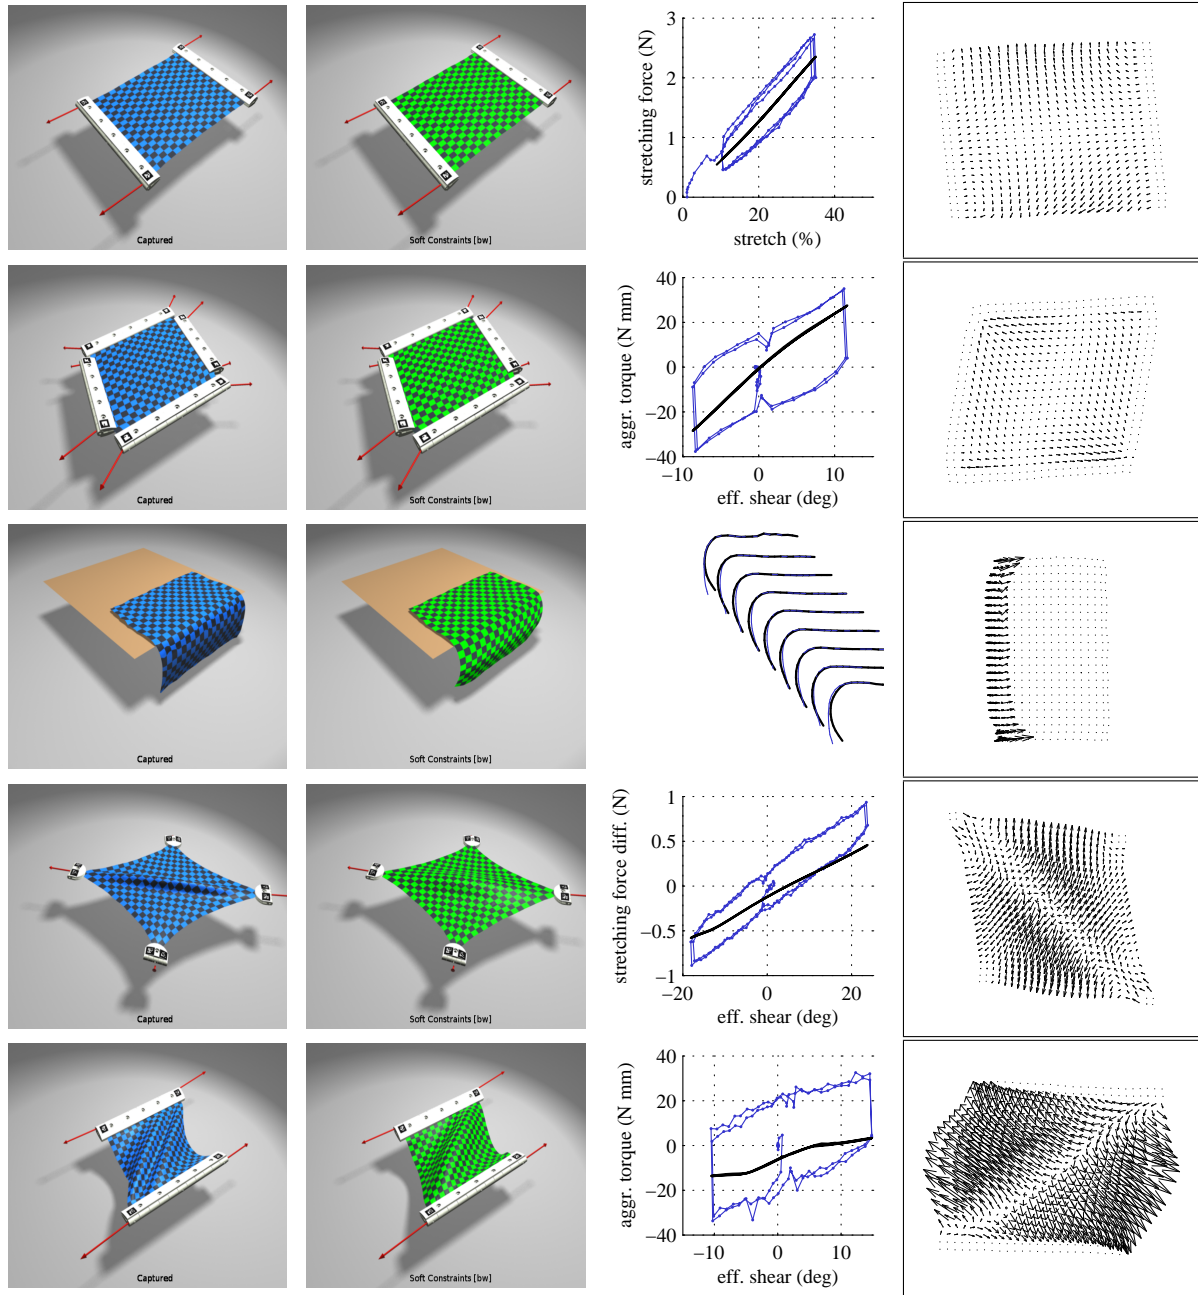


Fig. A.4 Fitting results for Sample #12, Soft Constraints model. Top to bottom; Stretch-X, Simple shear, Bend-X, Corner pull, Complex shear. Left to right: captured geometry, equilibrium of fitted model, force comparison (thin line: measurement; thick line: model), position residual (vertex position minus corresponding measured position, magnified 5x).

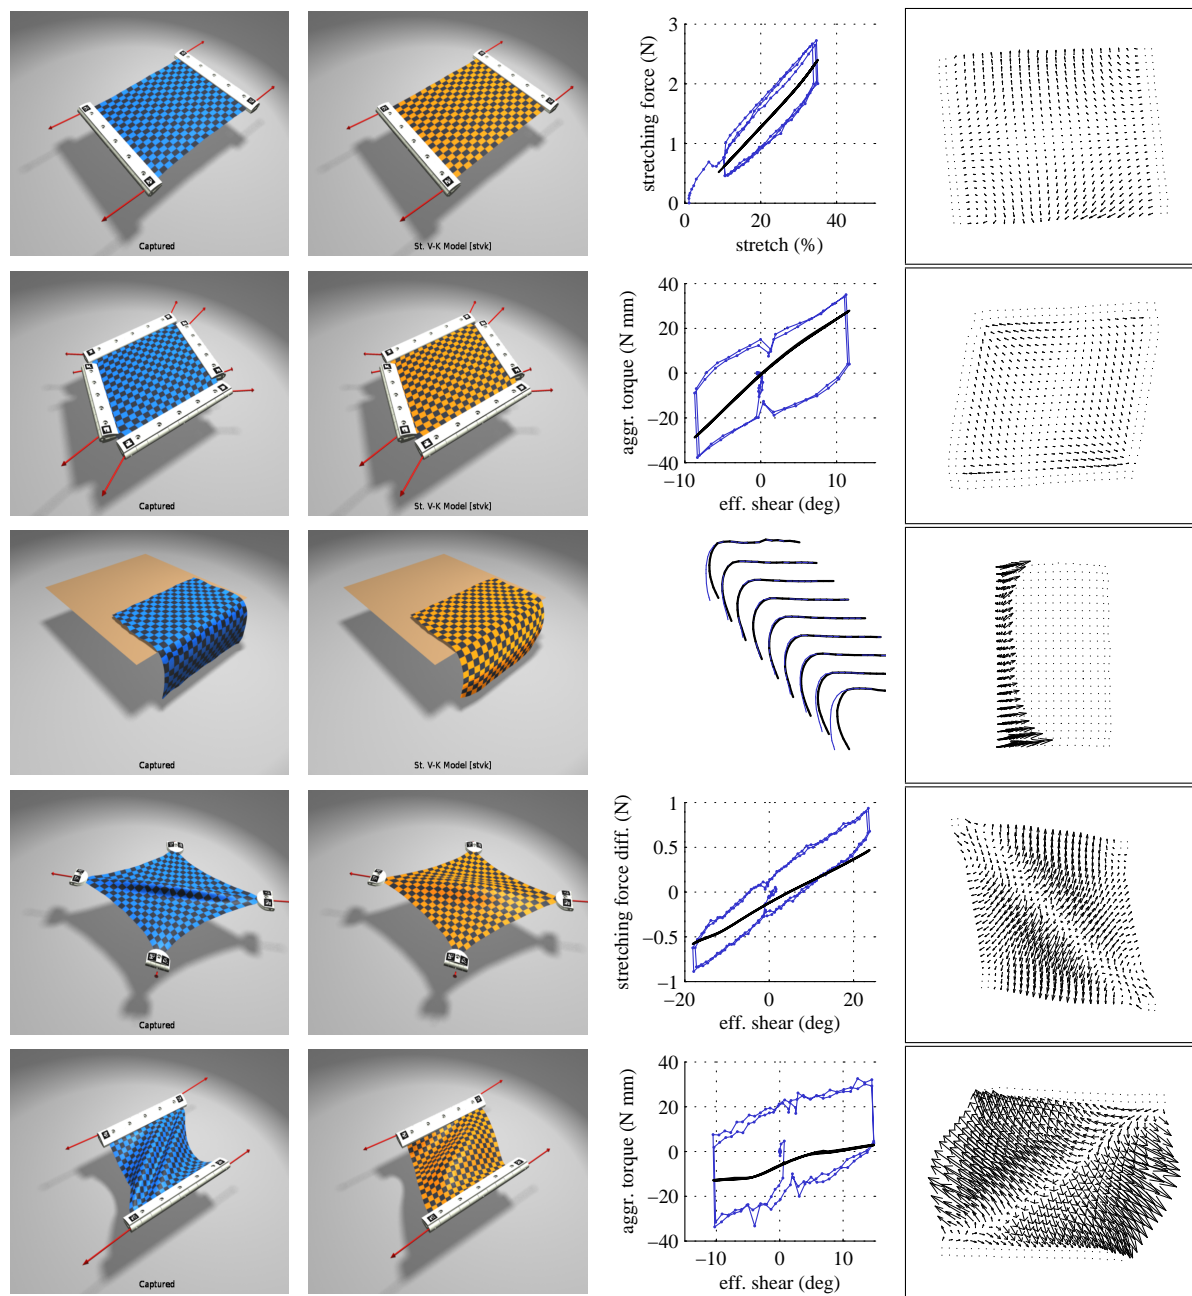


Fig. A.5 Fitting results for Sample #12, St. V-K model. Top to bottom; Stretch-X, Simple shear, Bend-X, Corner pull, Complex shear. Left to right: captured geometry, equilibrium of fitted model, force comparison (thin line: measurement; thick line: model), position residual (vertex position minus corresponding measured position, magnified 5x).

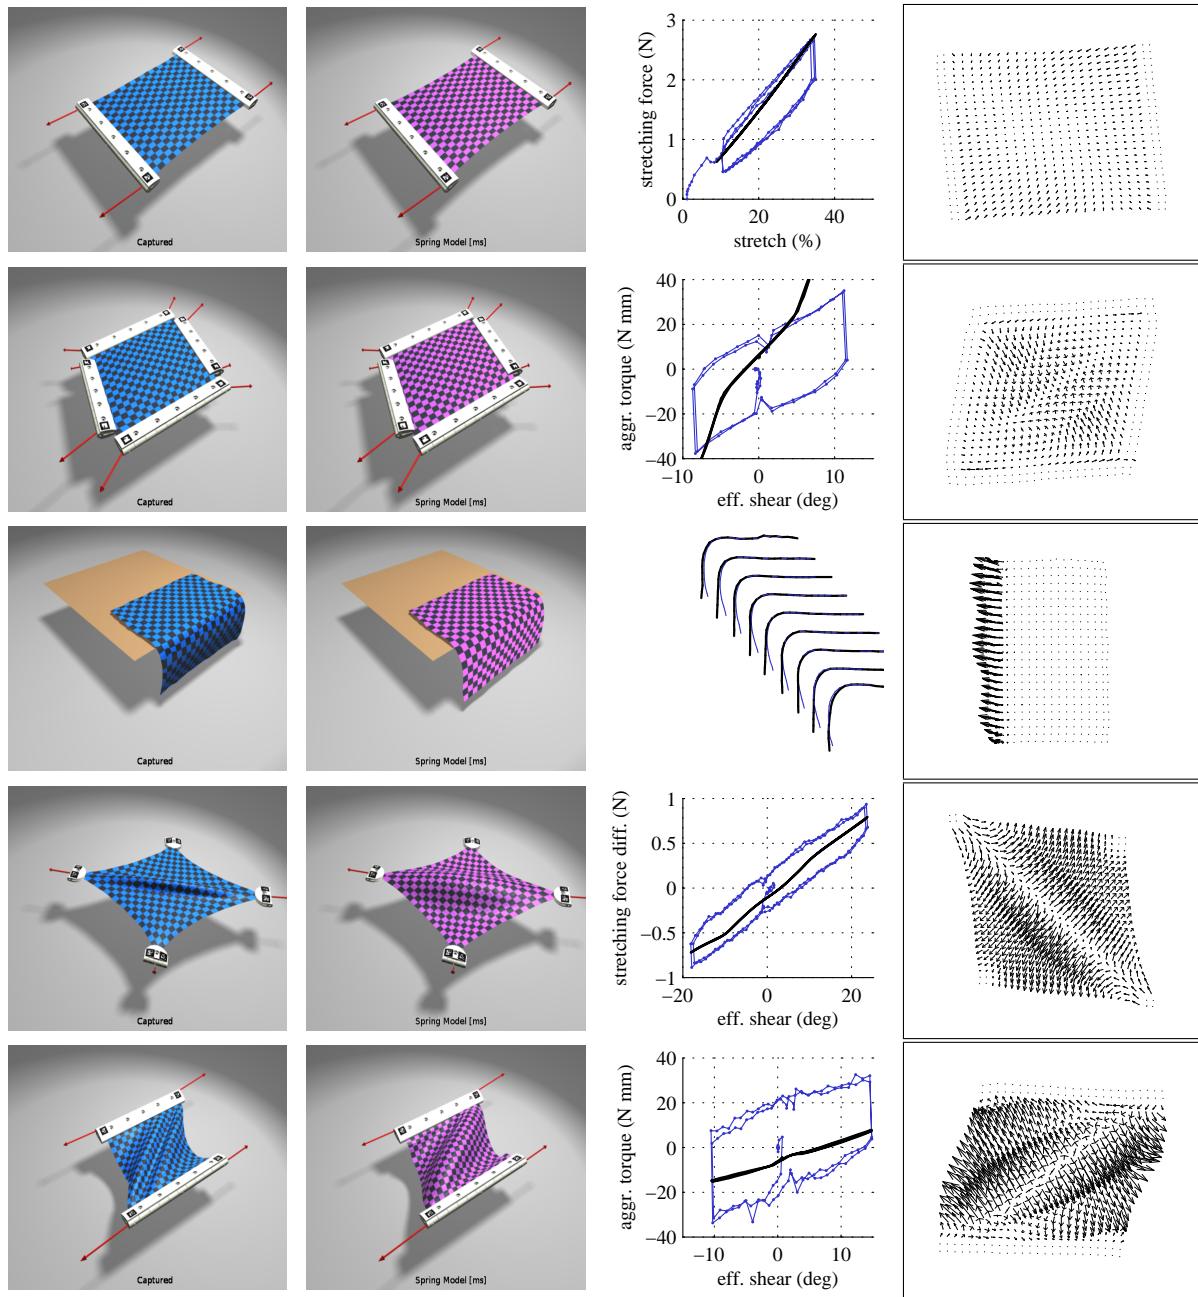


Fig. A.6 Fitting results for Sample #12, Spring model. Top to bottom; Stretch-X, Simple shear, Bend-X, Corner pull, Complex shear. Left to right: captured geometry, equilibrium of fitted model, force comparison (thin line: measurement; thick line: model), position residual (vertex position minus corresponding measured position, magnified 5x).

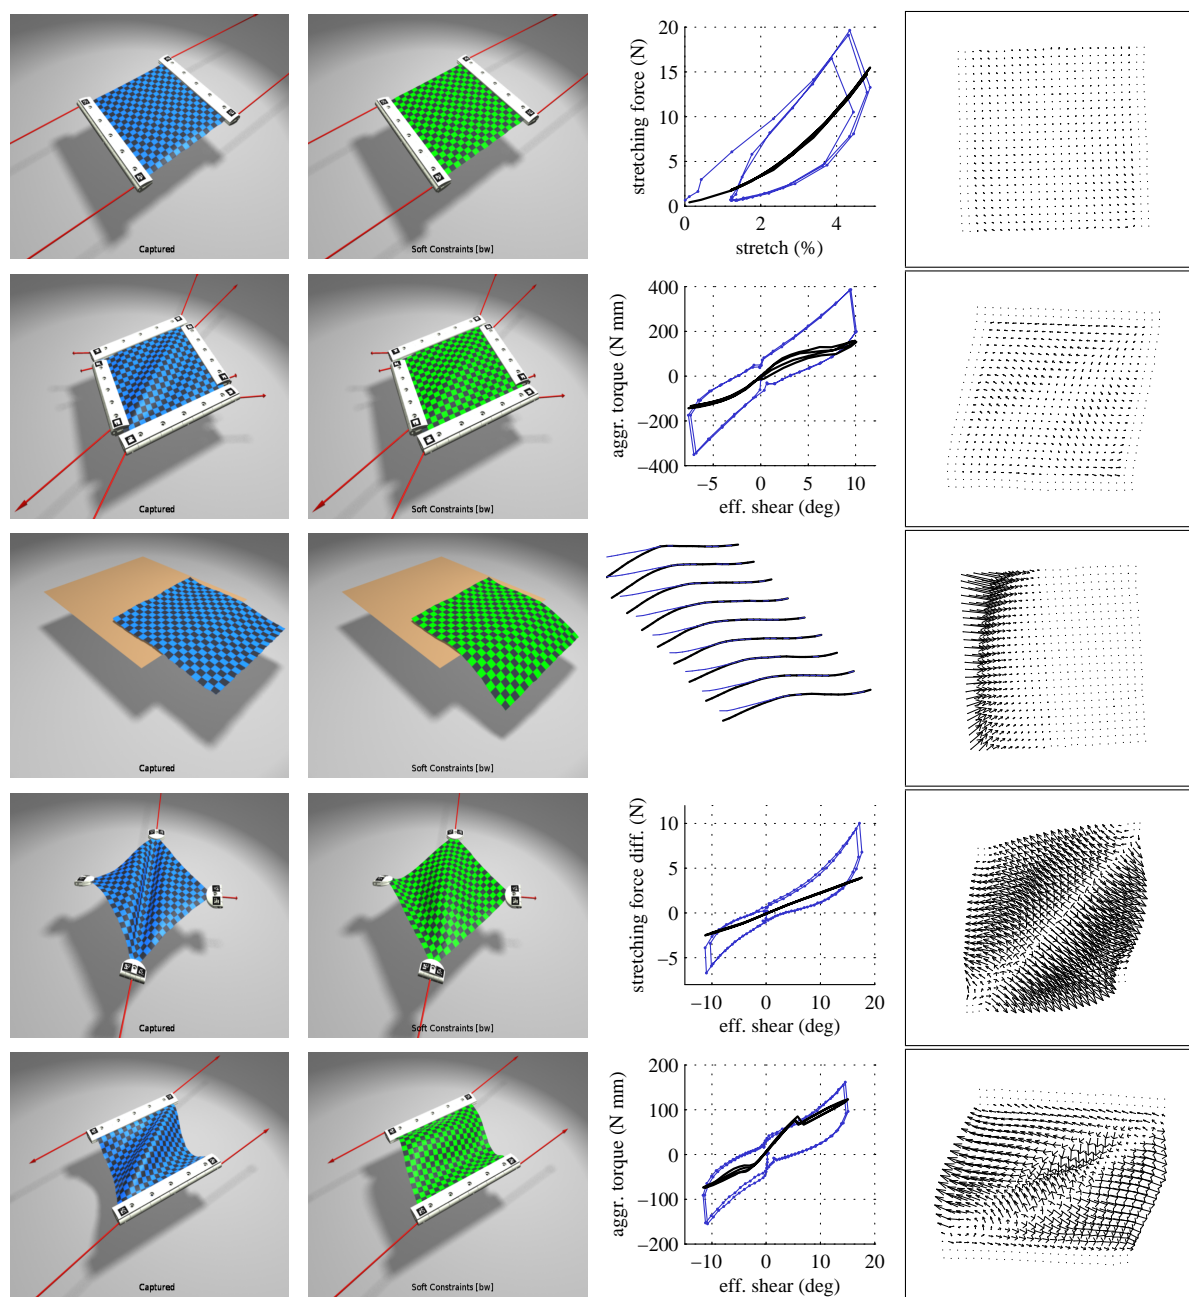


Fig. A.7 Fitting results for Sample #14, Soft Constraints model. Top to bottom; Stretch-X, Simple shear, Bend-X, Corner pull, Complex shear. Left to right: captured geometry, equilibrium of fitted model, force comparison (thin line: measurement; thick line: model), position residual (vertex position minus corresponding measured position, magnified 5x).

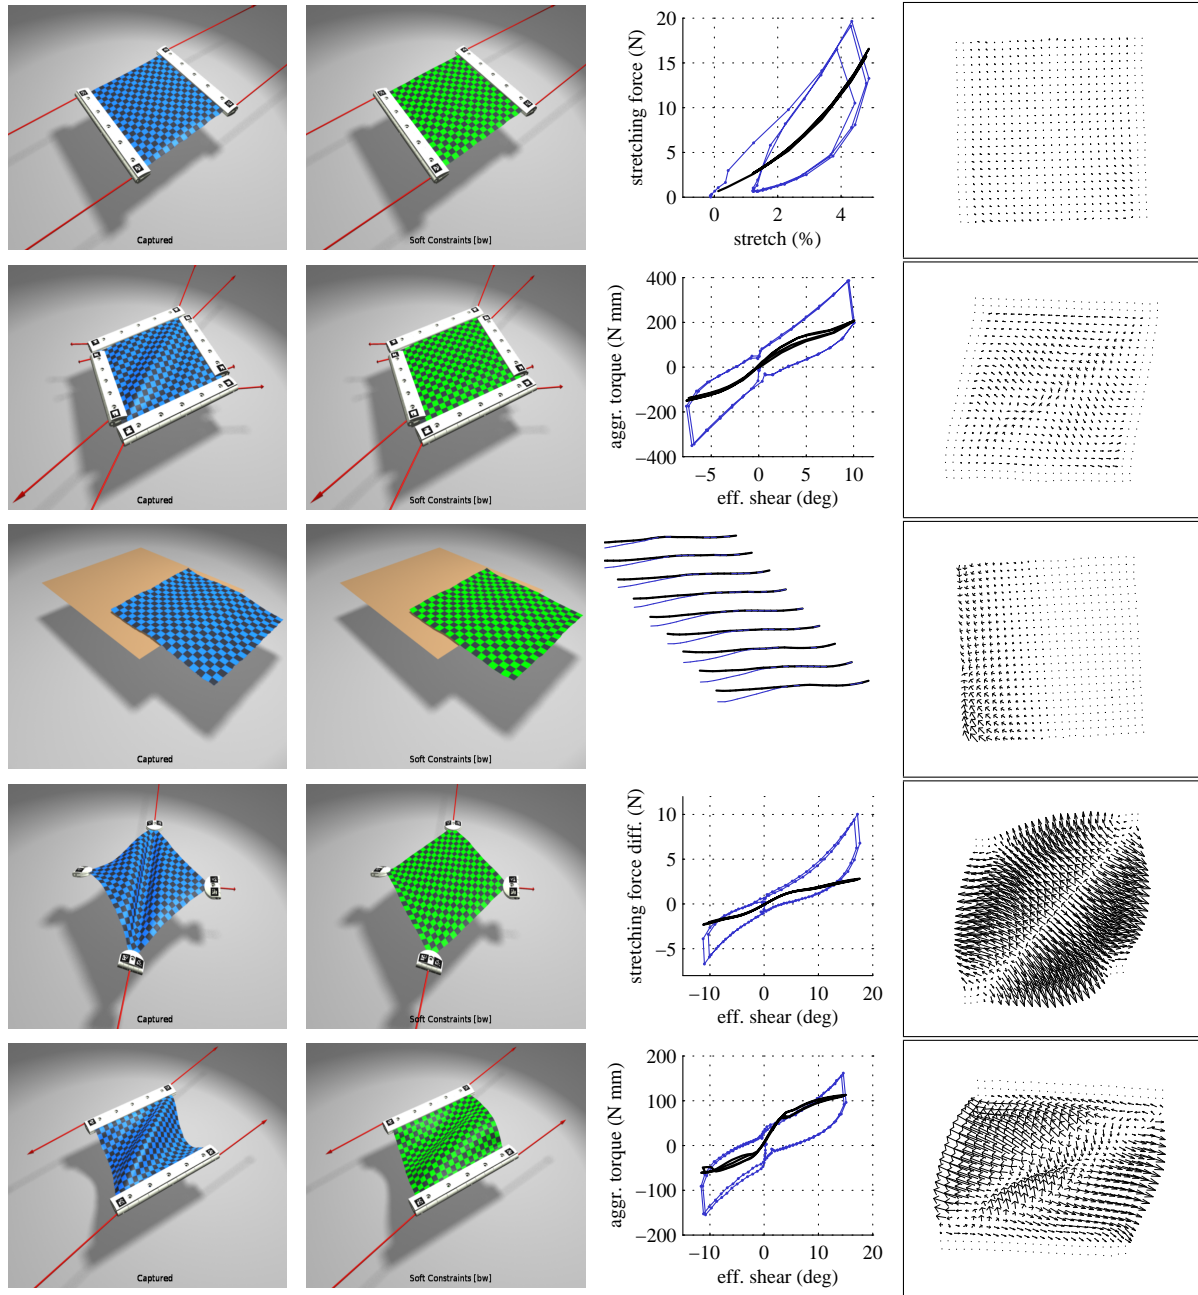


Fig. A.8 Fitting results for Sample #14, Soft Constraints model, isotropic. Top to bottom; Stretch-X, Simple shear, Bend-X, Corner pull, Complex shear. Left to right: captured geometry, equilibrium of fitted model, force comparison (thin line: measurement; thick line: model), position residual (vertex position minus corresponding measured position, magnified 5x).

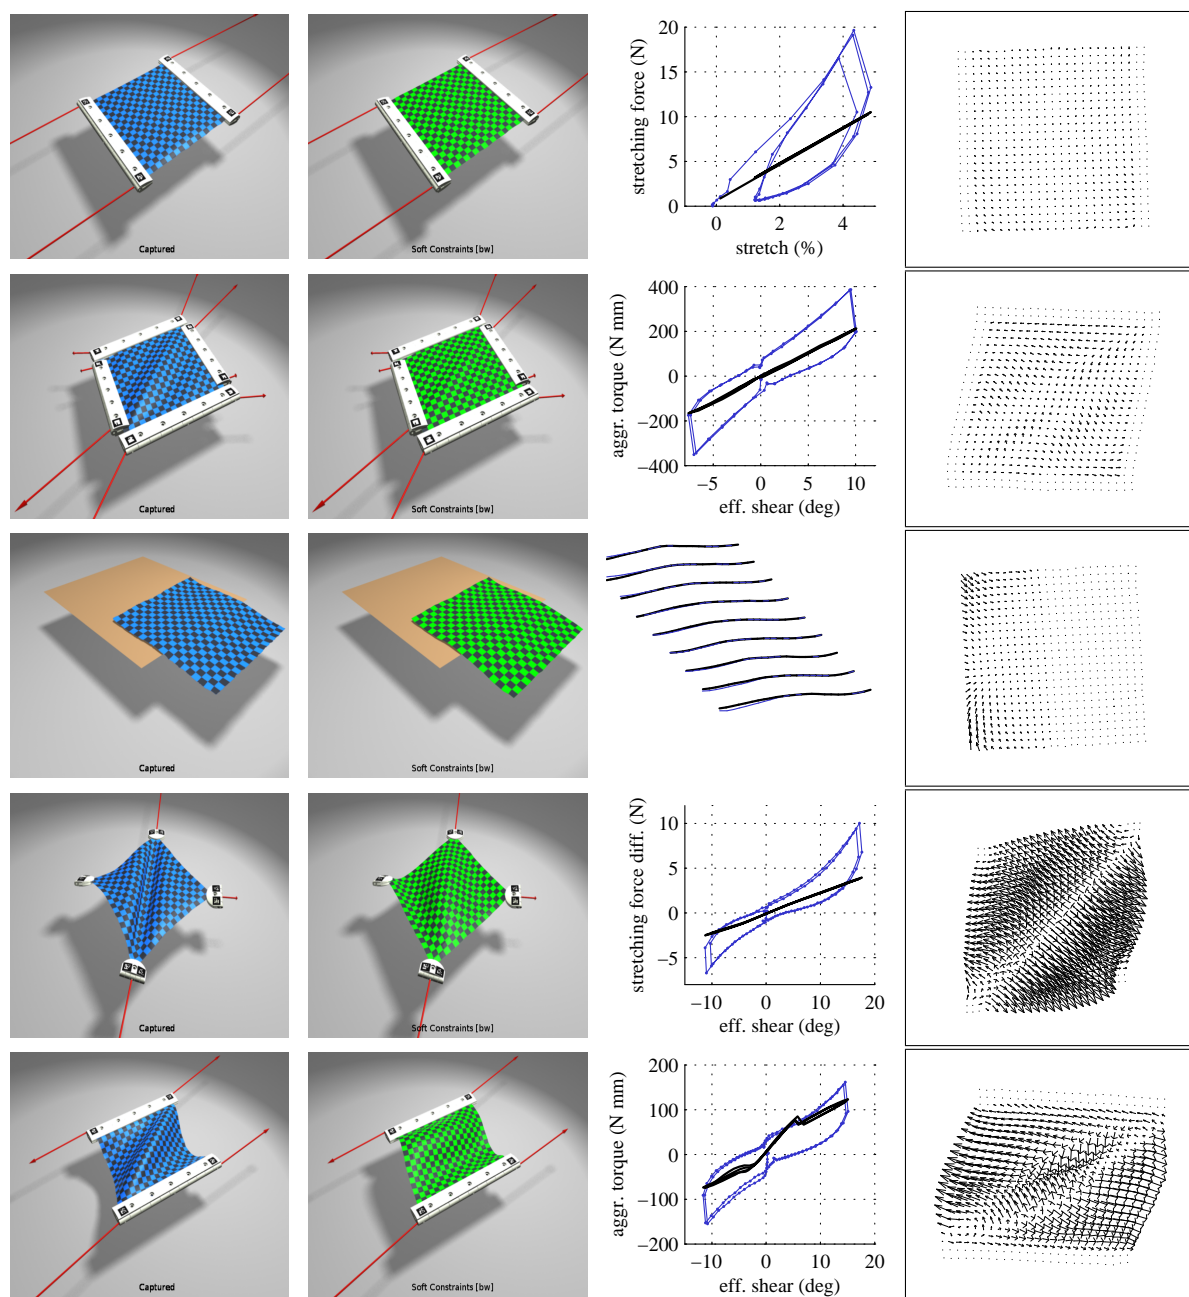


Fig. A.9 Fitting results for Sample #14, Soft Constraints model, linear. Top to bottom; Stretch-X, Simple shear, Bend-X, Corner pull, Complex shear. Left to right: captured geometry, equilibrium of fitted model, force comparison (thin line: measurement; thick line: model), position residual (vertex position minus corresponding measured position, magnified 5x).

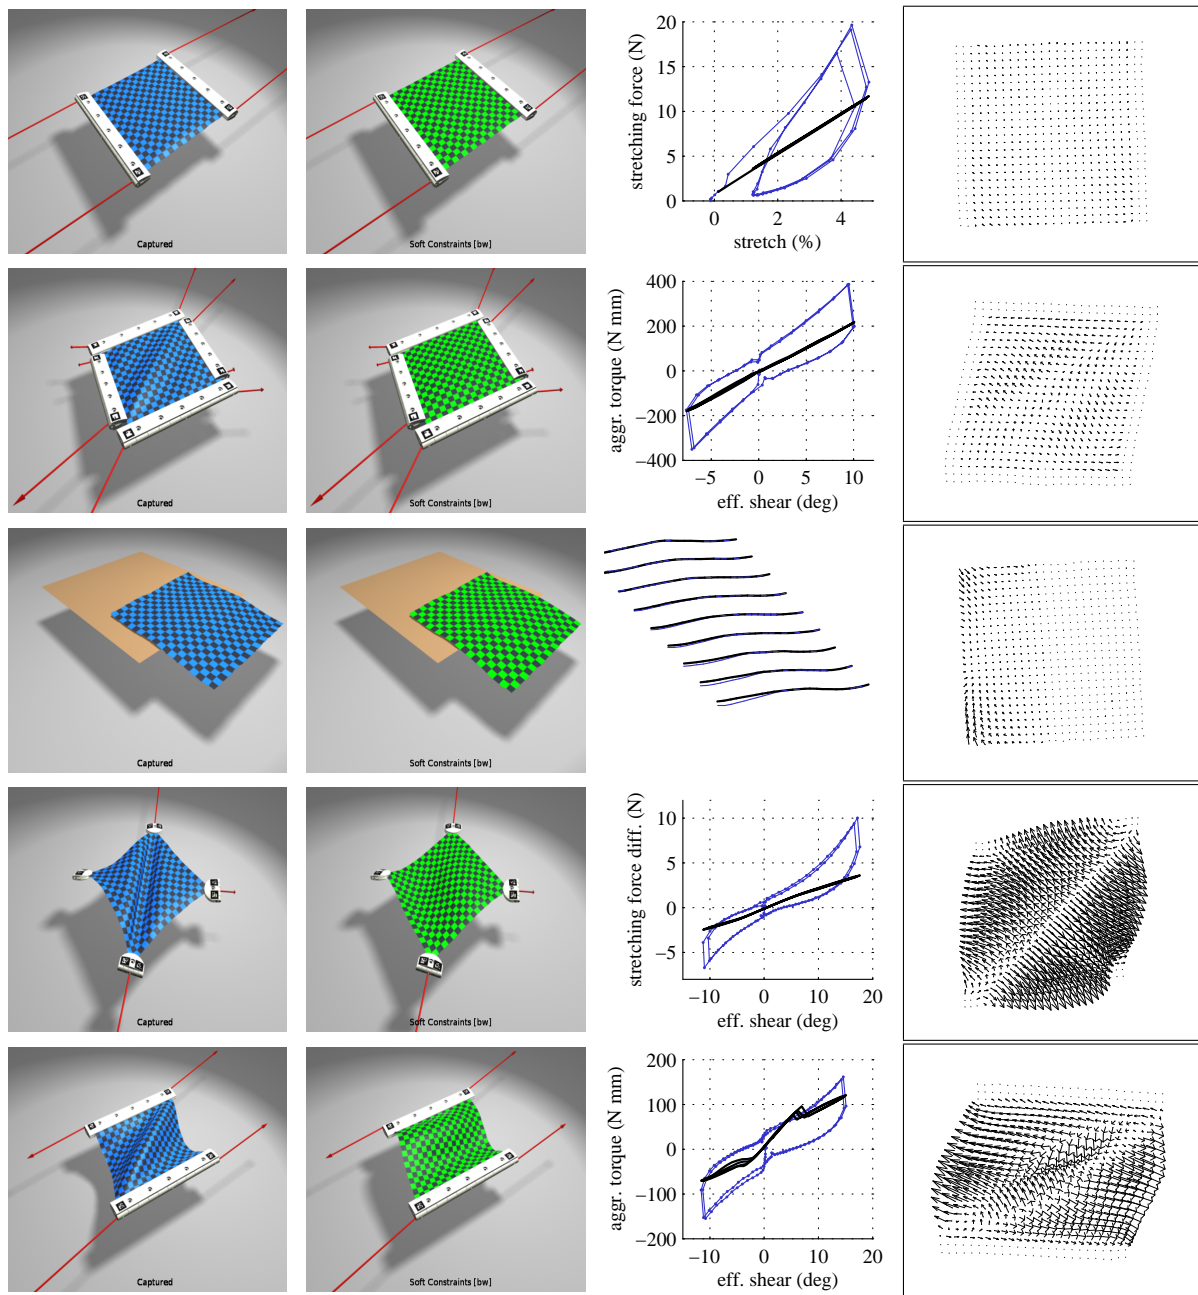


Fig. A.10 Fitting results for Sample #14, Soft Constraints model, linear isotropic. Top to bottom; Stretch-X, Simple shear, Bend-X, Corner pull, Complex shear. Left to right: captured geometry, equilibrium of fitted model, force comparison (thin line: measurement; thick line: model), position residual (vertex position minus corresponding measured position, magnified 5x).

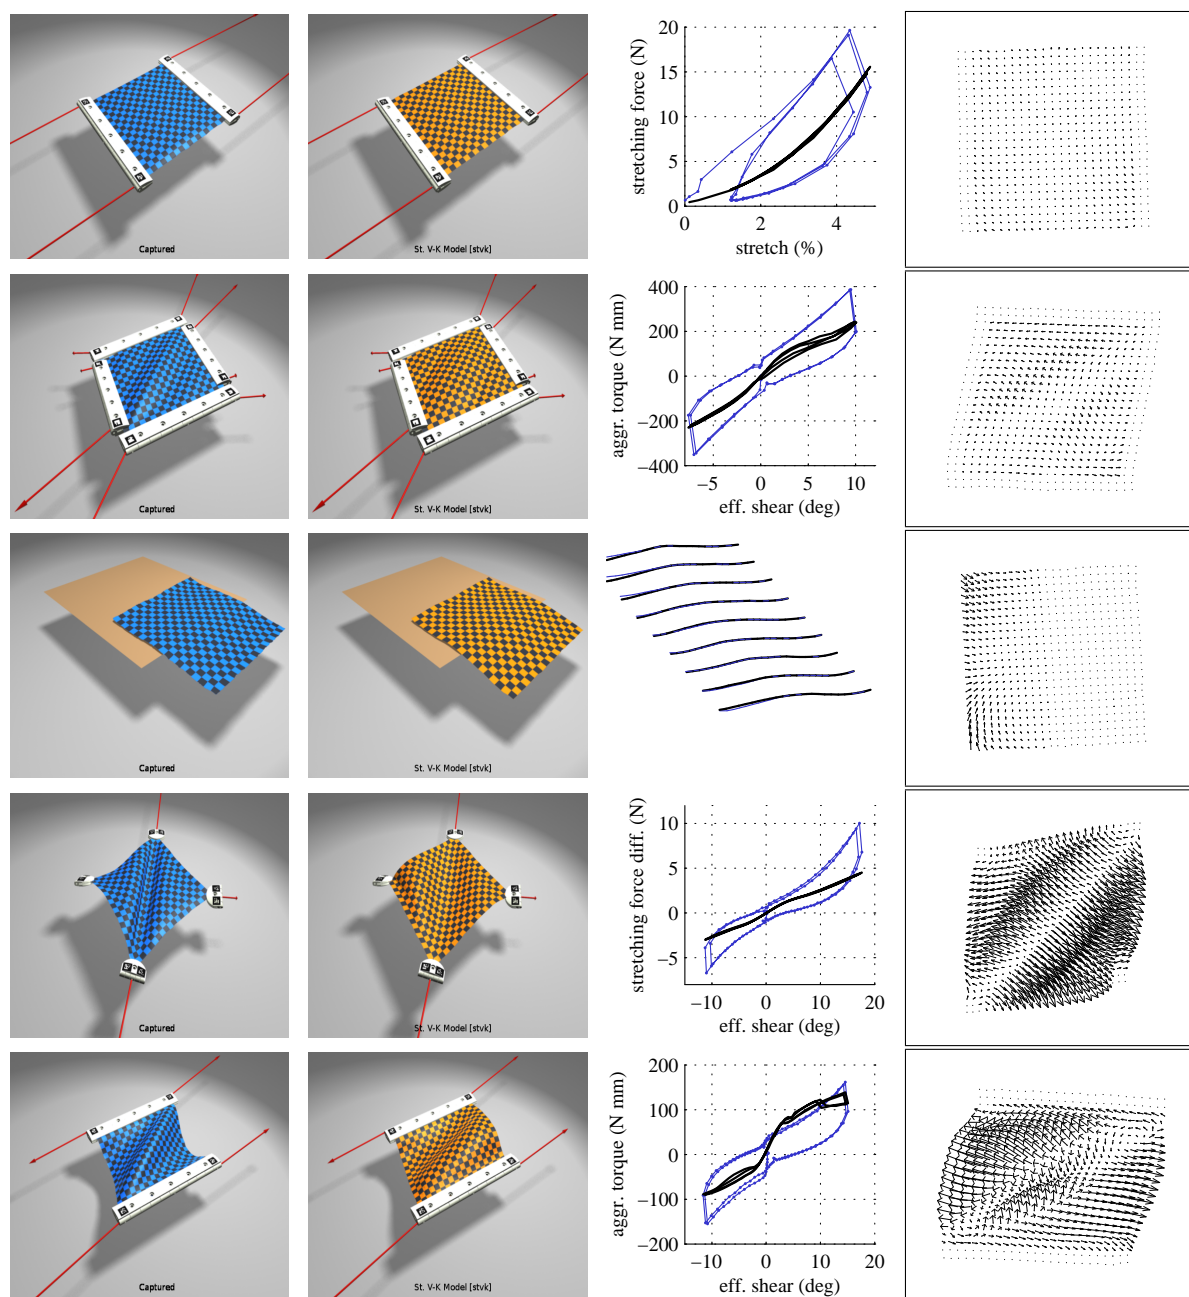


Fig. A.11 Fitting results for Sample #14, St. V-K model. Top to bottom; Stretch-X, Simple shear, Bend-X, Corner pull, Complex shear. Left to right: captured geometry, equilibrium of fitted model, force comparison (thin line: measurement; thick line: model), position residual (vertex position minus corresponding measured position, magnified 5x).

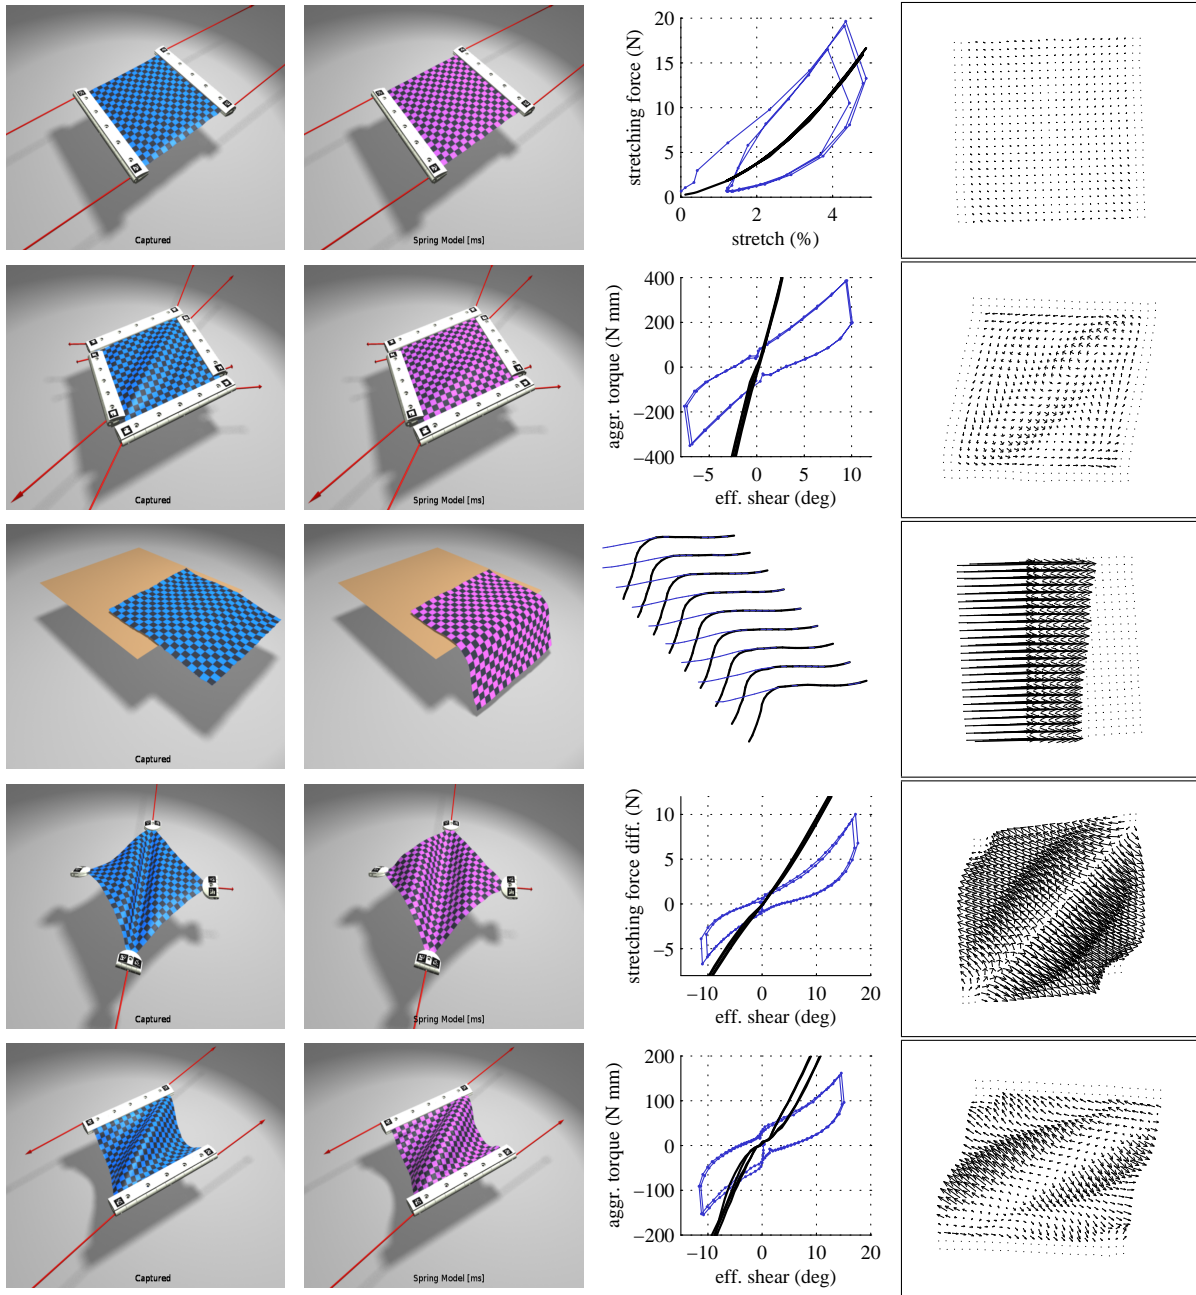


Fig. A.12 Fitting results for Sample #14, Spring model. Top to bottom; Stretch-X, Simple shear, Bend-X, Corner pull, Complex shear. Left to right: captured geometry, equilibrium of fitted model, force comparison (thin line: measurement; thick line: model), position residual (vertex position minus corresponding measured position, magnified 5x).

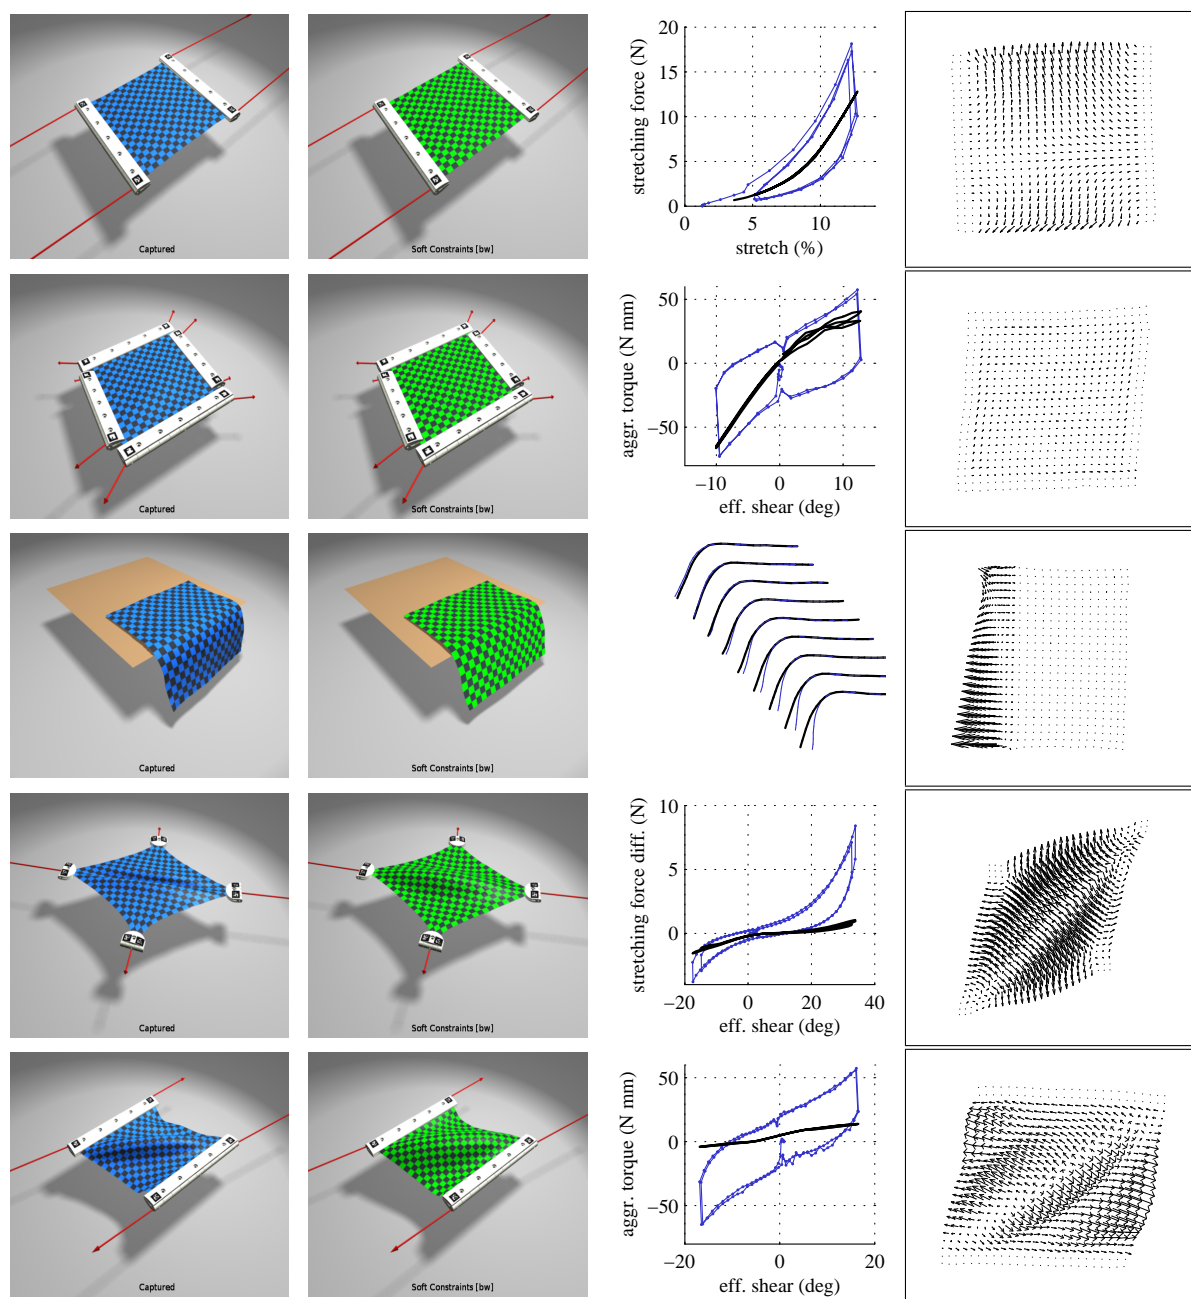


Fig. A.13 Fitting results for Sample #18, Soft Constraints model. Top to bottom; Stretch-X, Simple shear, Bend-X, Corner pull, Complex shear. Left to right: captured geometry, equilibrium of fitted model, force comparison (thin line: measurement; thick line: model), position residual (vertex position minus corresponding measured position, magnified 5x).

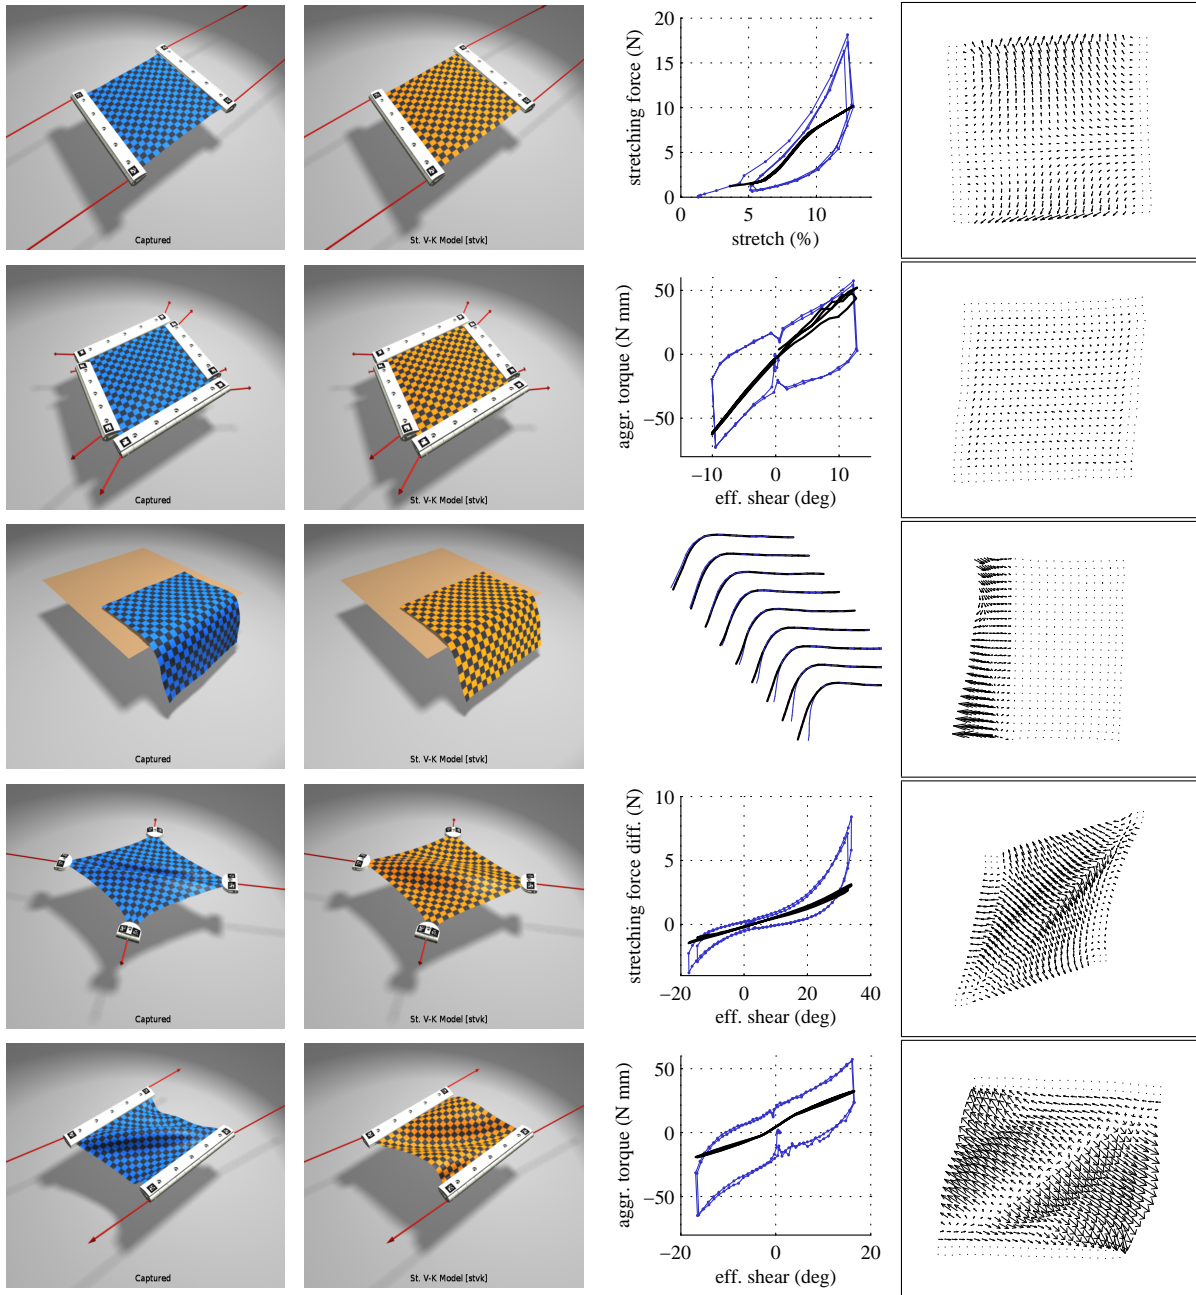


Fig. A.14 Fitting results for Sample #18, St. V-K model. Top to bottom; Stretch-X, Simple shear, Bend-X, Corner pull, Complex shear. Left to right: captured geometry, equilibrium of fitted model, force comparison (thin line: measurement; thick line: model), position residual (vertex position minus corresponding measured position, magnified 5x).

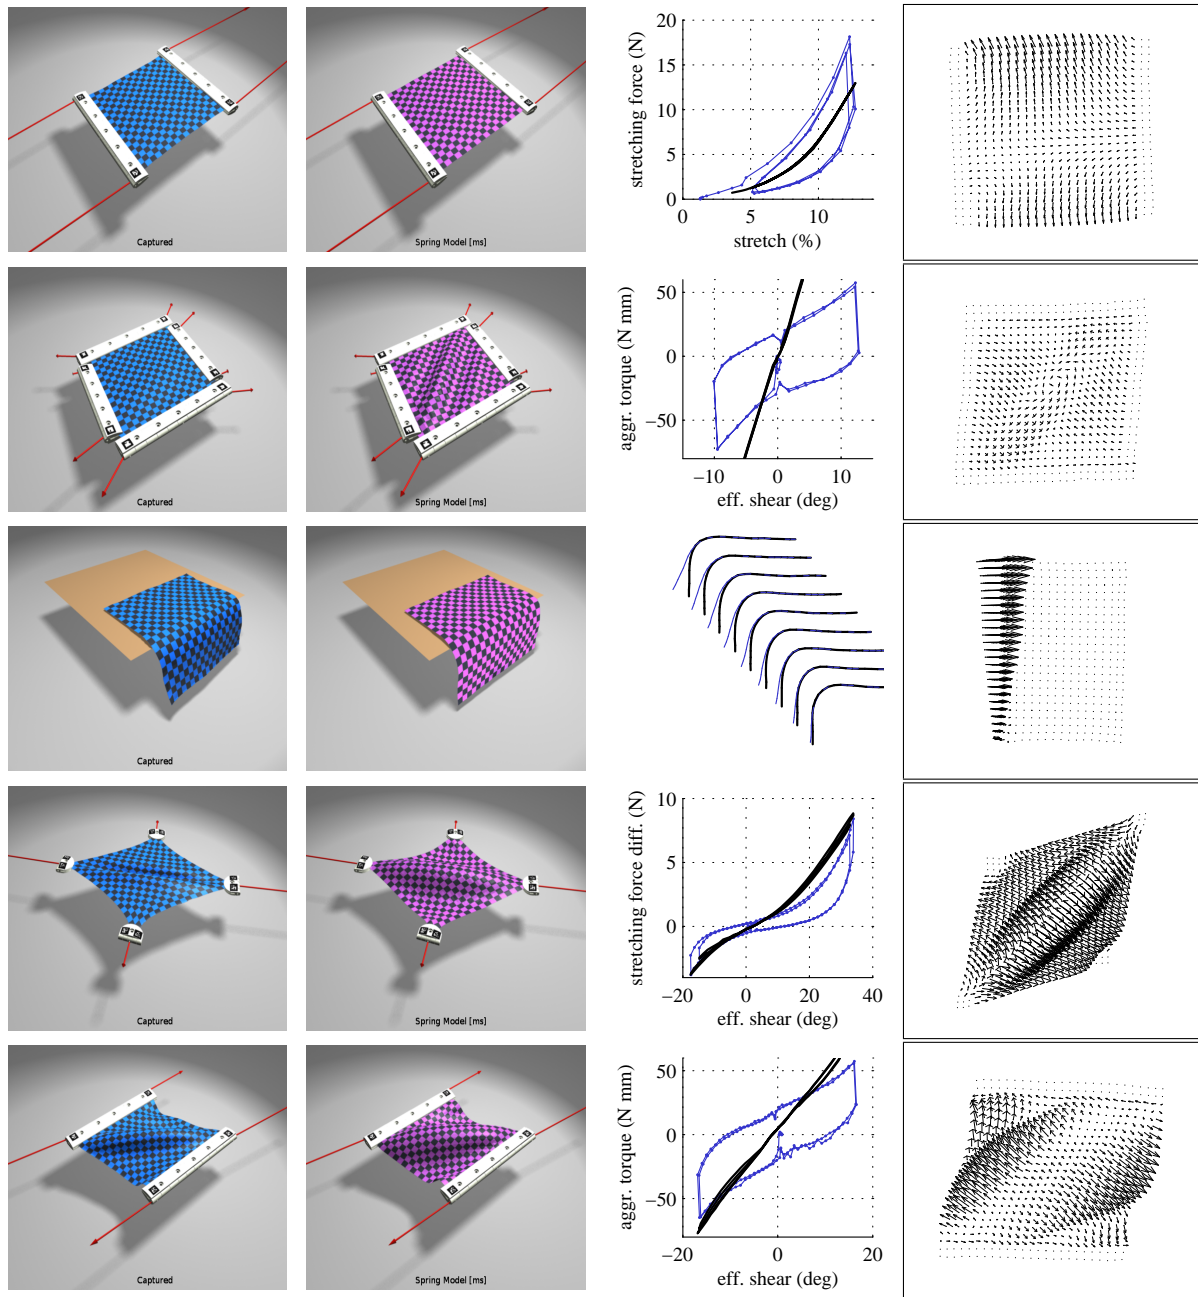


Fig. A.15 Fitting results for Sample #18, Spring model. Top to bottom; Stretch-X, Simple shear, Bend-X, Corner pull, Complex shear. Left to right: captured geometry, equilibrium of fitted model, force comparison (thin line: measurement; thick line: model), position residual (vertex position minus corresponding measured position, magnified 5x).

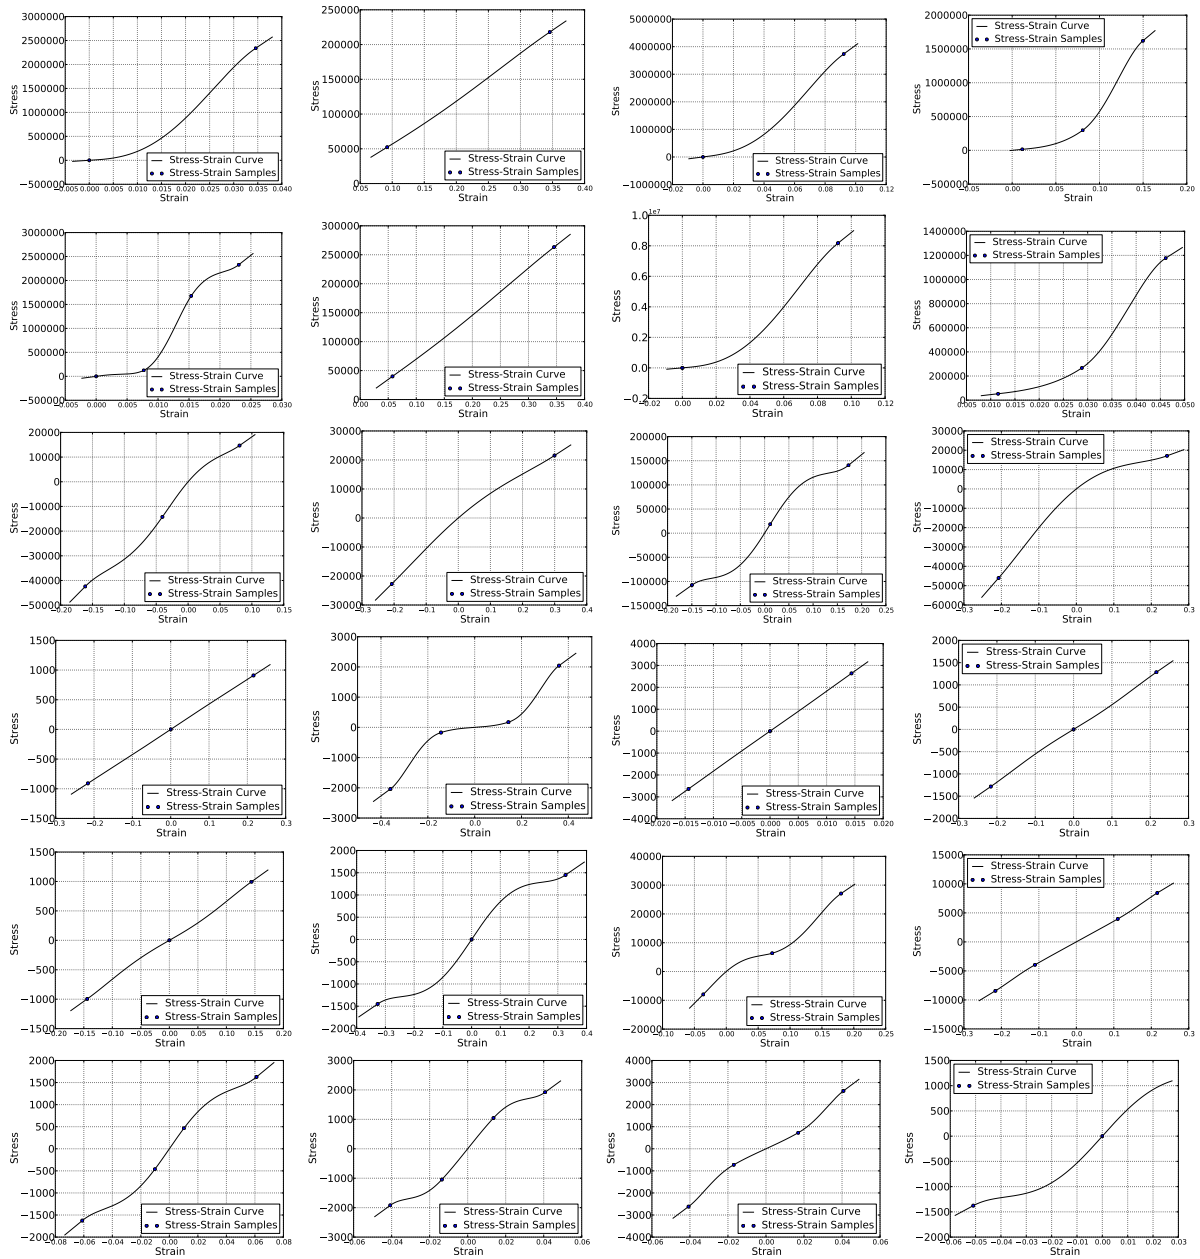


Fig. A.16 Stress-strain plots, Soft Constraints model. Cloth samples, from left to right: #4, #12, #14, #18. Deformation components, from top to bottom: warp-stretch, weft-stretch, shear, warp-bend, weft-bend, diagonal-bend.

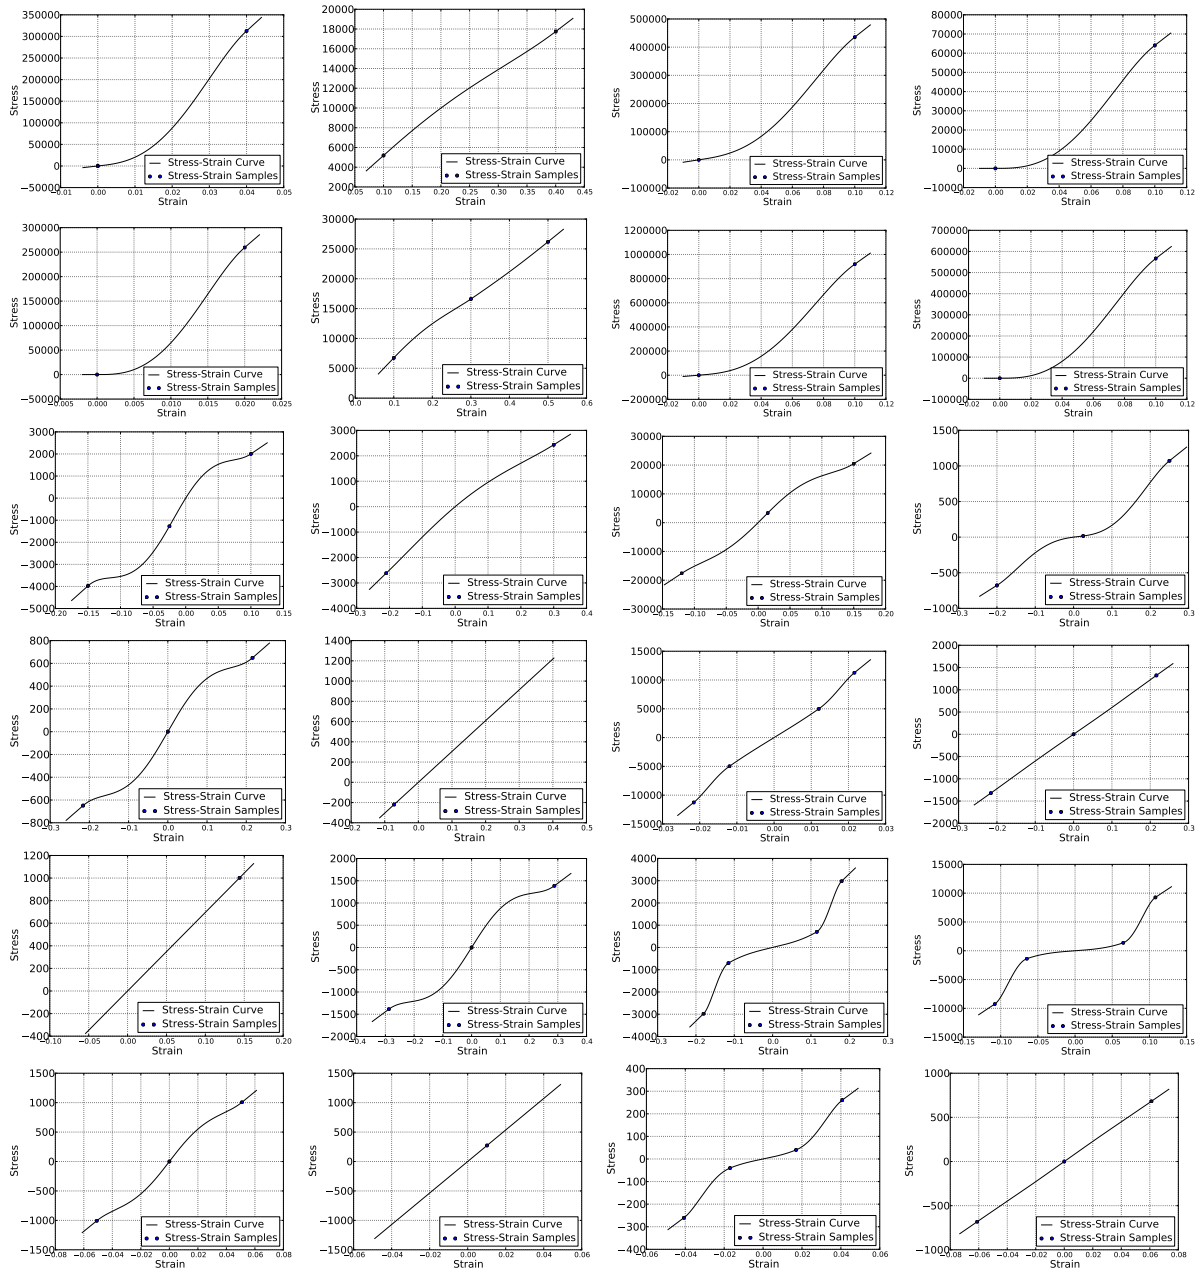


Fig. A.17 Stress-strain plots, St. V-K model. Cloth samples, from left to right: #4, #12, #14, #18. Deformation components, from top to bottom: warp-stretch, weft-stretch, shear, warp-bend, weft-bend, diagonal-bend.

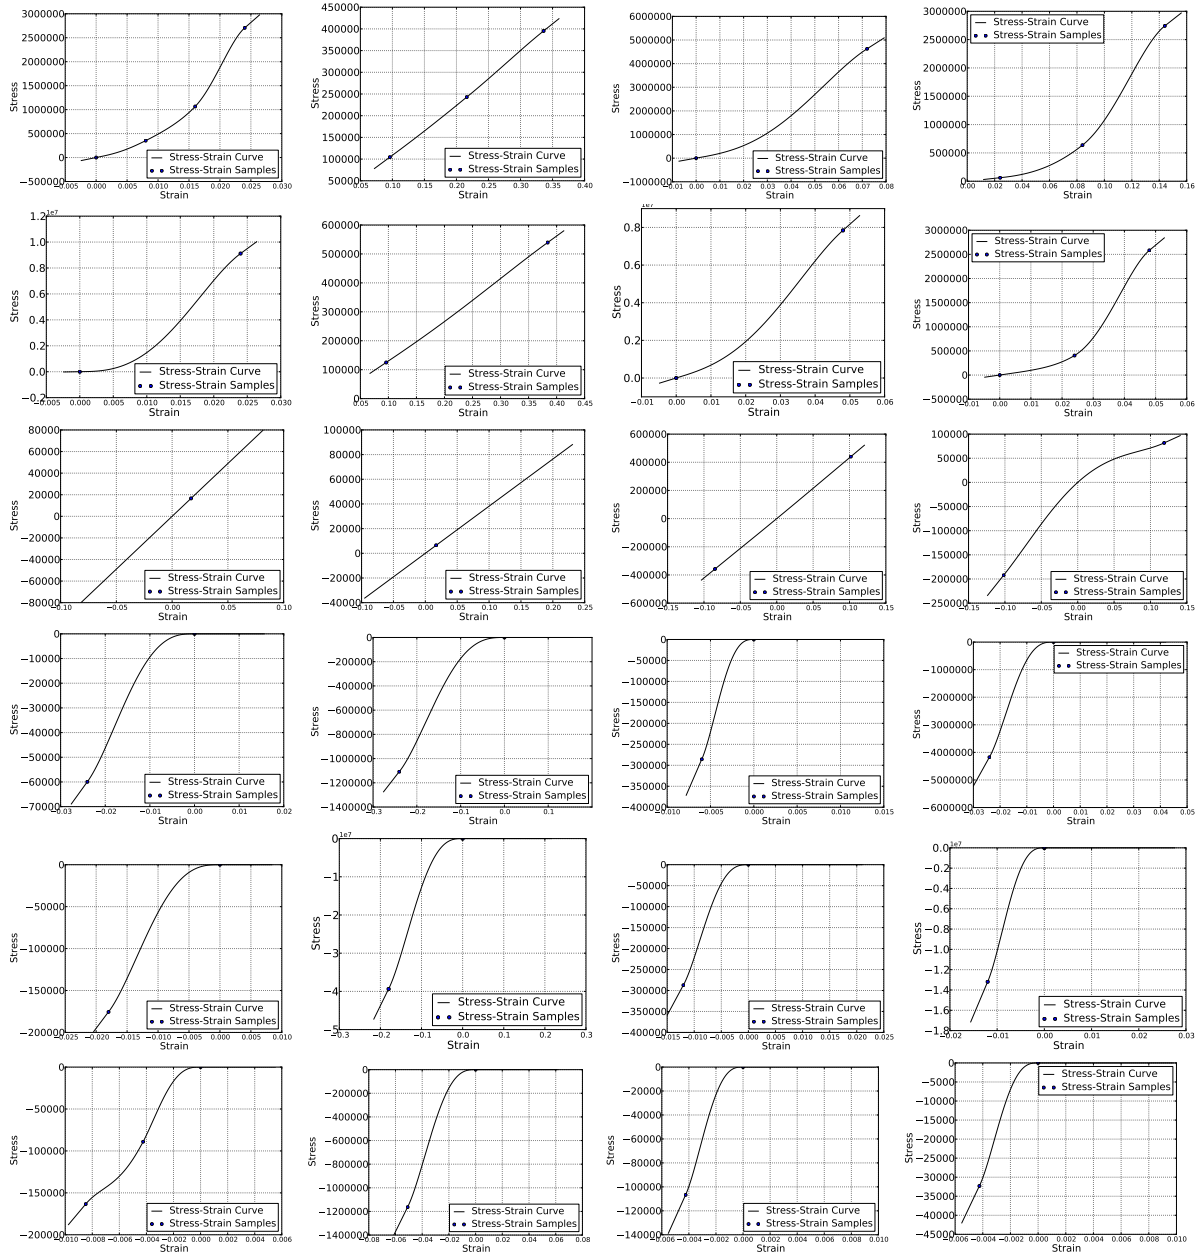


Fig. A.18 Stress-strain plots, Spring model. Cloth samples, from left to right: #4, #12, #14, #18. Deformation components, from top to bottom: warp-stretch, weft-stretch, shear, warp-bend, weft-bend, diagonal-bend.

Fit Variants	Cloth Model	Stretch-X	Stretch-Y	Shear	Bend-X	Bend-Y
Non-Linear Ortho.	Soft. Const.	5.546	6.394	0.33	16.10	33.14
	St.VK	5.558	6.323	0.331	19.51	33.22
	Spring	5.554	6.016	0.335	38.24	41.86
Linear Ortho.	Soft. Const.	6.708	7.859	0.336	25.88	33.94
	St.VK	6.656	7.829	0.337	28.42	34.59
	Spring	6.431	7.416	0.328	38.09	43.37
Non-Linear Iso.	Soft. Const.	7.438	7.438	0.329	39.53	39.53
	Spring	7.415	7.415	0.406	46.50	46.50
	St.VK	7.438	7.438	0.329	39.27	39.27
Linear Iso.	Soft. Const.	7.56	7.56	0.338	39.33	39.33
	Spring	7.466	7.466	0.333	45.91	45.91
	St.VK	7.543	7.543	0.338	39.59	39.59

Table A.1 Results of fitting to Sample-04 (cotton satin). Each entry shows position residual (mm, RMS over all free points) for bend scenarios, force residual (N, RMS over all clips) for Stretch scenarios, and clip parallel force residual (N, RMS over all clips) for Shear scenarios.

Fit Variants	Cloth Model	Stretch-X	Stretch-Y	Shear	Bend-X	Bend-Y
Non-Linear Ortho.	Soft. Const.	0.529	0.621	0.215	21.31	22.01
	St.VK	0.529	0.622	0.213	22.53	23.29
	Spring	0.529	0.612	0.345	32.01	43.39
Linear Ortho.	Soft. Const.	0.533	0.628	0.229	26.41	49.49
	St.VK	0.539	0.647	0.222	23.37	23.16
	Spring	0.531	0.619	0.237	31.54	43.69
Non-Linear Iso.	Soft. Const.	0.624	0.624	0.213	57.633	57.633
	Spring	0.606	0.606	0.364	39.54	39.54
	St.VK	0.621	0.621	0.213	58.921	58.921
Linear Iso.	Soft. Const.	0.637	0.637	0.225	57.565	57.565
	Spring	0.622	0.622	0.234	39.562	39.562
	St.VK	0.635	0.635	0.22	53.787	53.787

Table A.2 Results of fitting to Sample-12 (rayon/spandex knit). Each entry shows position residual (mm, RMS over all free points) for bend scenarios, force residual (N, RMS over all clips) for Stretch scenarios, and clip parallel force residual (N, RMS over all clips) for Shear scenarios.

Fit Variants	Cloth Model	Stretch-X	Stretch-Y	Shear	Bend-X	Bend-Y
Non-Linear Ortho.	Soft. Const.	4.508	5.883	1.837	16.15	20.84
	St.VK	4.498	6.04	1.857	18.74	70.38
	Spring	4.411	5.538	2.041	38.05	137.6
Linear Ortho.	Soft. Const.	5.165	6.163	2.286	11.88	84.91
	St.VK	5.118	6.168	2.303	24.11	64.63
	Spring	4.995	5.741	2.356	190.78	156.76
Non-Linear Iso.	Soft. Const.	5.614	1.707	124.812		
	Spring	5.947	2.015	149.943		
	St.VK	5.604	1.834	94.346		
Linear Iso.	Soft. Const.	5.811	2.306	124.497		
	Spring	5.947	2.352	158.902		
	St.VK	5.78	2.322	125.063		

Table A.3 Results of fitting to Sample-14 (cotton denim). Each entry shows position residual (mm, RMS over all free points) for bend scenarios, force residual (N, RMS over all clips) for Stretch scenarios, and clip parallel force residual (N, RMS over all clips) for Shear scenarios.

Fit Variants	Cloth Model	Stretch-X	Stretch-Y	Shear	Bend-X	Bend-Y
Non-Linear Ortho.	Soft. Const.	3.961	5.324	0.485	22.29	19.91
	St.VK	3.961	5.338	0.492	24.57	23.61
	Spring	3.936	5.298	0.489	62.73	26.64
Linear Ortho.	Soft. Const.	5.087	6.326	0.556	26.96	29.47
	St.VK	4.937	6.281	0.563	32.11	89.45
	Spring	5.046	6.217	0.563	30.69	33.56
Non-Linear Iso.	Soft. Const.	5.594	0.578	84.73		
	Spring	5.784	0.683	107.1		
	St.VK	6.974	0.627	88.55		
Linear Iso.	Soft. Const.	6.507	0.483	81.32		
	Spring	6.538	0.498	118.7		
	St.VK	7.699	0.492	81.55		

Table A.4 Results of fitting to Sample-18 (wool/cotton blend). Each entry shows position residual (mm, RMS over all free points) for bend scenarios, force residual (N, RMS over all clips) for Stretch scenarios, and clip parallel force residual (N, RMS over all clips) for Shear scenarios.

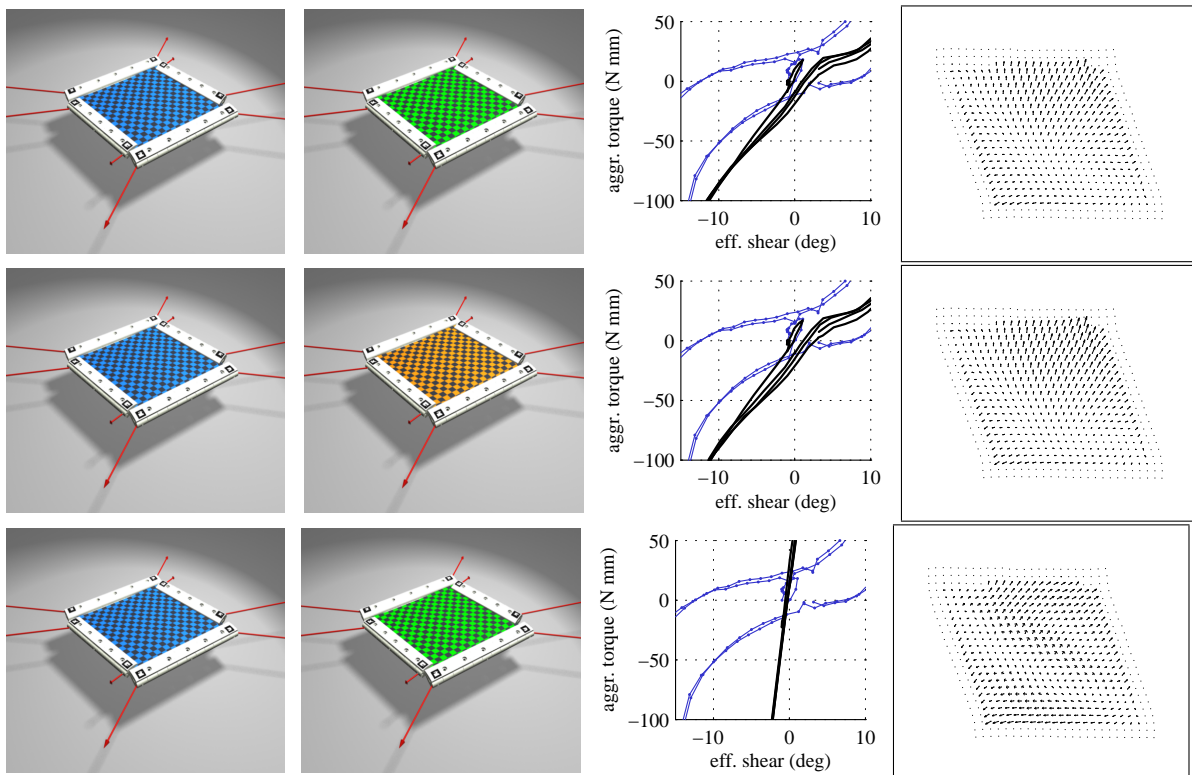


Fig. A.19 Fitting results for Sample #04.2. Top to bottom; Soft Constraints model, St. V-K model, Spring model. Left to right: captured geometry, equilibrium of fitted model, force comparison (thin line: measurement; thick line: model), position residual (vertex position minus corresponding measured position, magnified 5x).

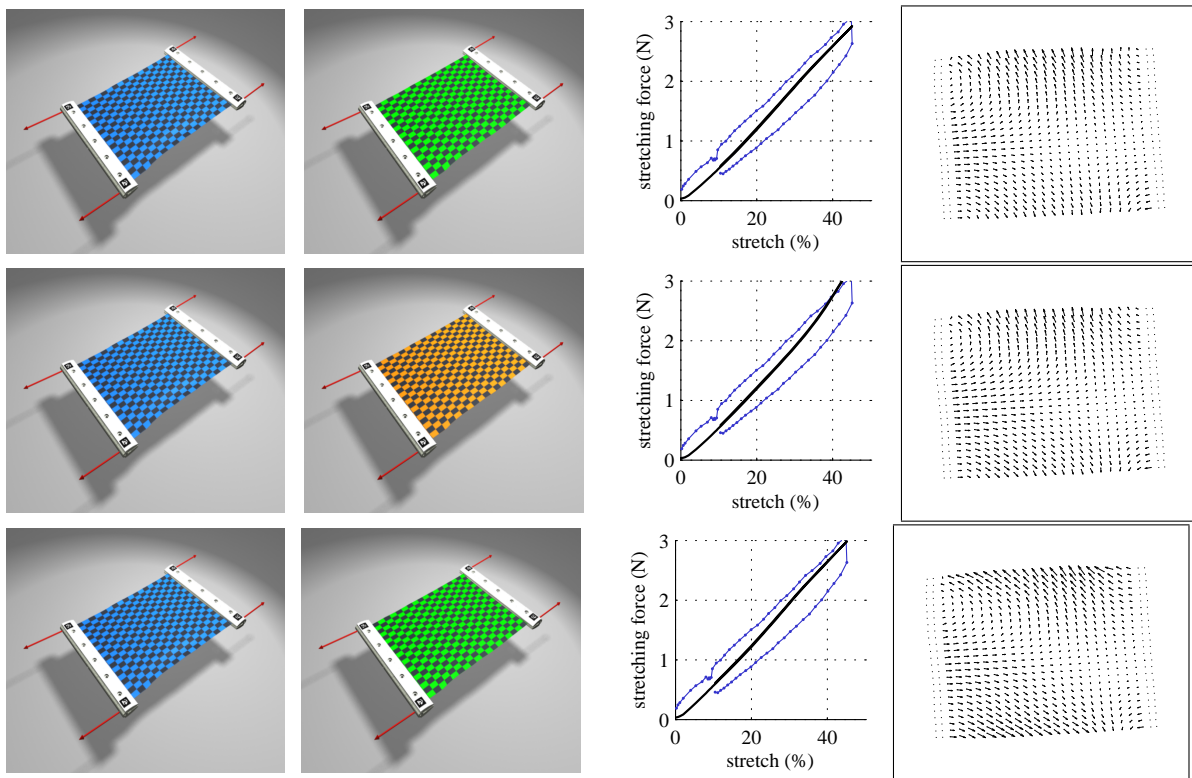


Fig. A.20 Fitting results for Sample #12.2. Top to bottom; Soft Constraints model, St. V-K model, Spring model. Left to right: captured geometry, equilibrium of fitted model, force comparison (thin line: measurement; thick line: model), position residual (vertex position minus corresponding measured position, magnified 5x).

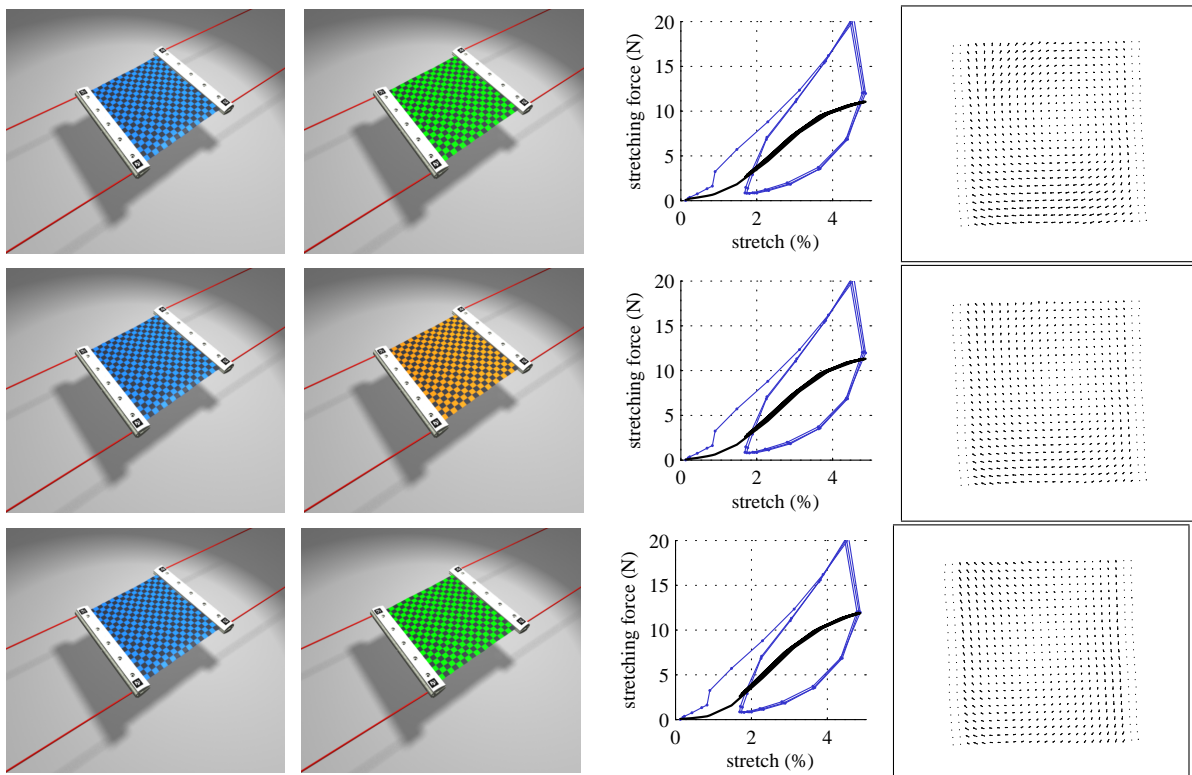


Fig. A.21 Fitting results for Sample #14.2. Top to bottom; Soft Constraints model, St. V-K model, Spring model. Left to right: captured geometry, equilibrium of fitted model, force comparison (thin line: measurement; thick line: model), position residual (vertex position minus corresponding measured position, magnified 5x).

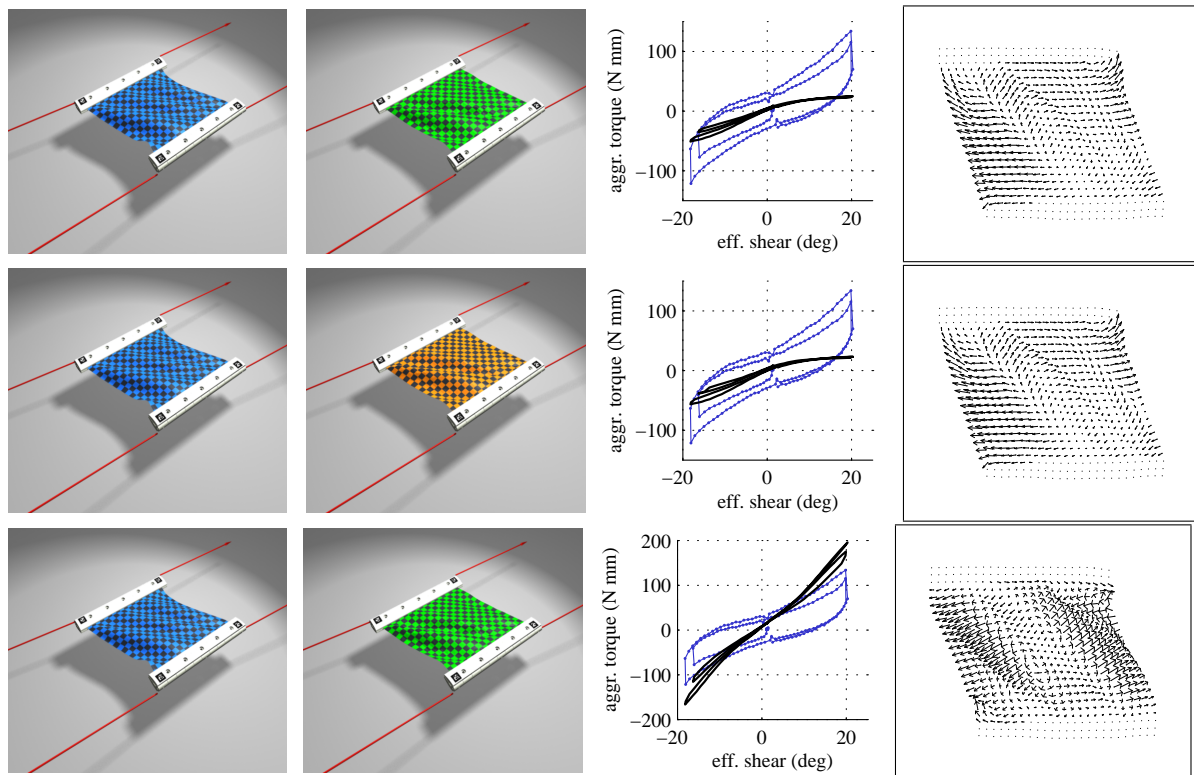


Fig. A.22 Fitting results for Sample #18.2. Top to bottom; Soft Constraints model, St. V-K model, Spring model. Left to right: captured geometry, equilibrium of fitted model, force comparison (thin line: measurement; thick line: model), position residual (vertex position minus corresponding measured position, magnified 5x).

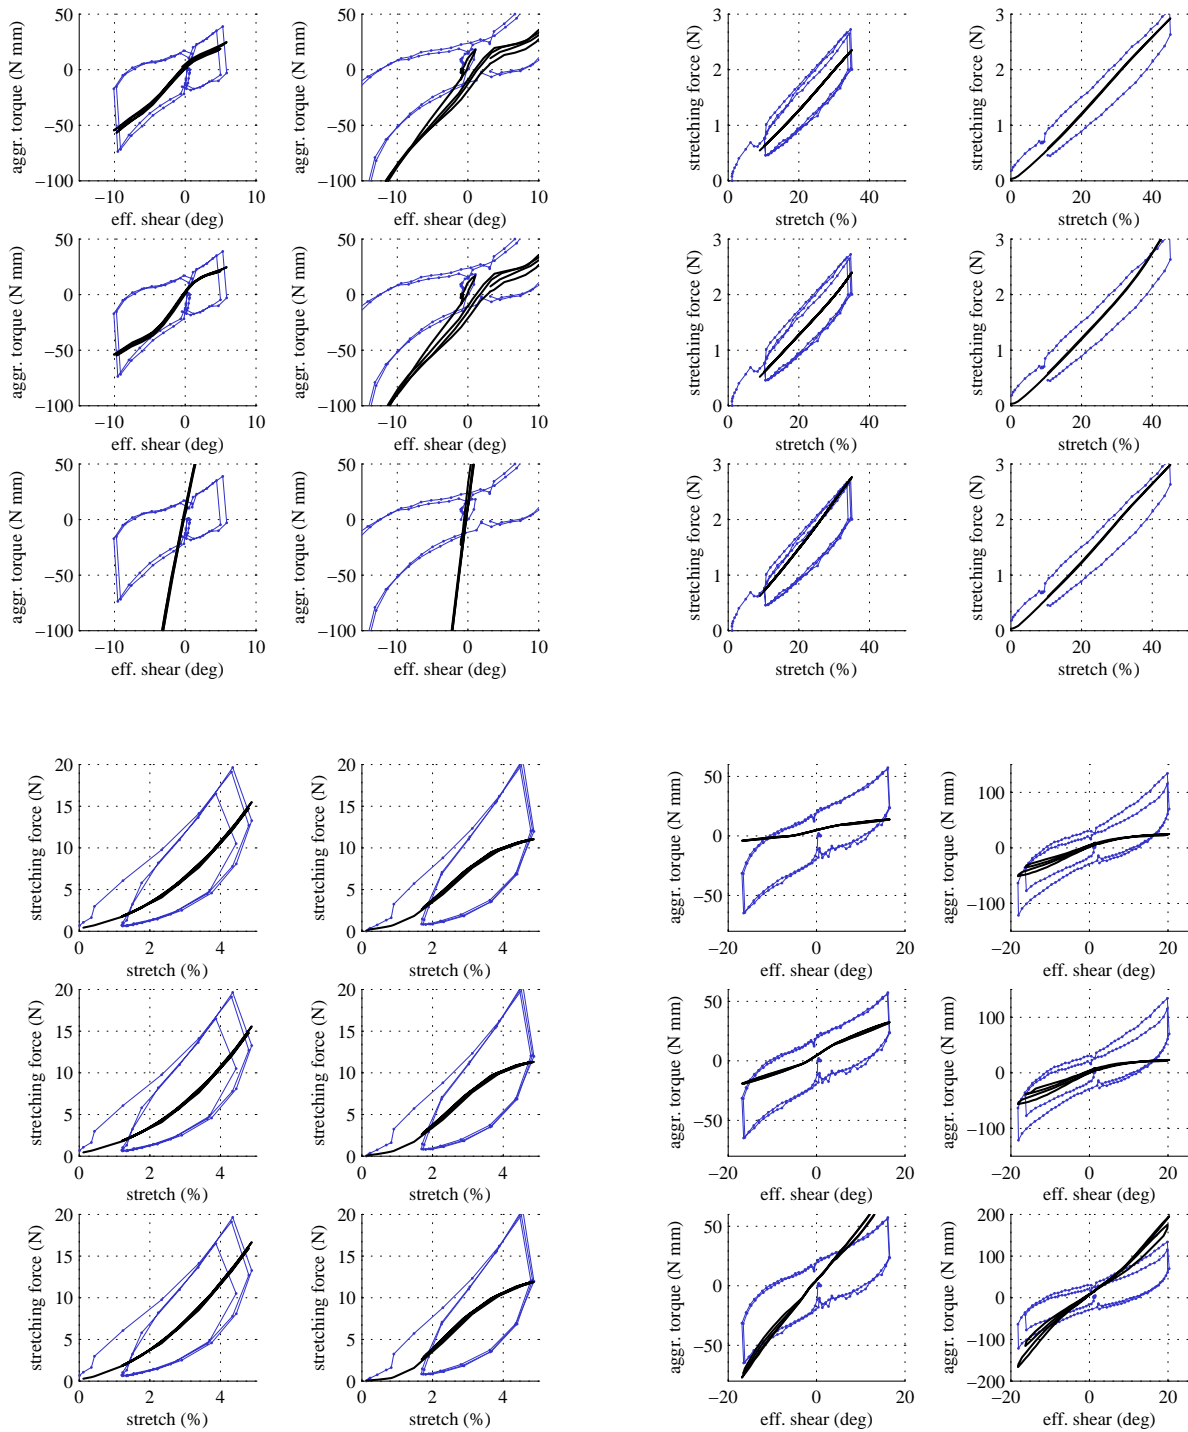


Fig. A.23 Comparison of evaluation results for different cloth samples for each fabric. Top-left: Samples #4 and #4.2, Shear scenario. Top-right: Samples #12 and #12.2, Stretch-X scenario. Bottom-left: Samples #14 and #14.2, Stretch-X scenario. Bottom-right: Samples #18 and #18.2, Shear-X scenario. Top to bottom on each set: Soft Constraints model, St. V-K model, Spring model. Each plot shows force comparison (thin line: measurement; thick line: model).