
Estimating the Material Properties of Fabric Through the Observation of Motion

by

Katherine L. Bouman

B.S.E., Electrical Engineering, University of Michigan, 2011

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Science
in Electrical Engineering and Computer Science
at the Massachusetts Institute of Technology

September 2013

© 2013 Massachusetts Institute of Technology
All Rights Reserved.

Signature of Author: _____

Department of Electrical Engineering and Computer Science
August 30, 2013

Certified by: _____

William T. Freeman
Professor of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by: _____

Leslie A. Kolodziejski
Professor of Electrical Engineering and Computer Science
Chair, Committee for Graduate Students

Estimating the Material Properties of Fabric Through the Observation of Motion

by Katherine L. Bouman

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of
Master of Science

Abstract

We present a framework for predicting the physical properties of moving deformable objects observed in video. We apply our framework to analyze videos of fabrics moving under various unknown wind forces, and recover two key material properties of the fabric: stiffness and mass. We extend features previously developed to compactly represent static image textures to describe video textures such as fabric motion. A discriminatively trained regression model is then used to predict the physical properties of fabric from these features. The success of our model is demonstrated on a new database of fabric videos with corresponding measured ground truth material properties that we have collected. We show that our predictions are well correlated with both measured material properties and human perception of material properties. Our contributions include: (a) a method for predicting the material properties of fabric from a video, (b) a database that can be used for training and testing algorithms for predicting fabric properties containing RGB and RGBD videos of real videos with associated material properties and rendered videos of simulated fabric with associated model parameters, and (c) a perceptual study of humans' ability to estimate the material properties of fabric from videos and images.

Thesis Supervisor: William T. Freeman

Title: Professor of Electrical Engineering and Computer Science

Acknowledgments

I would like to thank my advisor, Bill Freeman, for his invaluable guidance and patience in my work. His creative approach to solving both new and old problems is a truly rare gift that I hope to one day embody at least a fraction of.

I'd also like to thank Bei Xiao, Peter Battaglia, John Fisher and Wojciech Matusik for all of their help and advice on this work.

I would like to thank Edward Delp for taking me under his wing as a high school student and introducing me to image processing and research in general, Mimi Boutin for all of her encouragement and support, and Clayton Scott for introducing me to machine learning while I was at the University of Michigan. Additionally, I would like to thank my mentors Sergio Goma, Kalin Atanassov, and Vikas Ramachandra at Qualcomm, Karl Ni at MIT Lincoln Labs, and Ce Liu at Microsoft Research for rewarding summers as an intern.

I'd like to thank my friends for being there for me and making me laugh when I need it most. I especially want to thank everyone in the vision group for the helpful discussions and making coming to work so much fun: Andrew, Adrian, Ramesh, Phillip, Neal, Roger, Donglai, Bei, Christian, Kayhan, Yichang, and particularly my officemates of D460, Miki, Tianfan, and Hossein.

Most of all, I'd like to thank my father, Charles, my mother, Cristina, and my siblings, Amanda and Alexander for their love and support over the years. My parents' guidance and encouragement has been invaluable, and I will always appreciate how much they have sacrificed for both my siblings and my success.

This work was supported by an NSF Graduate Fellowship and the NSF grants "Analyzing Images Through Time" and "Reconstructive recognition: Uniting statistical scene understanding and physics-based visual reasoning".

Contents

Abstract	3
Acknowledgments	4
List of Figures	9
List of Tables	13
1 Introduction	15
1.0.1 Contributions	16
1.0.2 Thesis Outline	17
2 Background	19
2.0.3 Material Properties of Fabric	19
Cloth Models	19
Invasively Fitting Cloth Model Parameters for Real Fabric	20
Passively Fitting Cloth Model Parameters for Real Fabric	20
2.0.4 Texture	20
Parametric Texture Models	21
Temporal Textures For Material Classification	22
3 Approach	23
3.0.5 Material Localization	23
3.0.6 Motion Estimation	24
3.0.7 Statistical Features	25
Marginal Statistics	27
Autocorrelation	27
Magnitude Correlation	28
Phase Correlation	29
3.0.8 Model Learning	29
4 Database	31

4.0.9	Real Fabric Dataset	31
	RGB-Depth Videos	32
	Material Property Measurements	33
4.0.10	Simulated Fabric Dataset	33
5	Human Material Perception	37
5.0.11	Experimental Setup	37
5.0.12	Data Analysis and Discussion	39
6	Algorithm Results and Discussion	41
6.0.13	Implementation Details	41
6.0.14	Results	42
6.0.15	Discussion	42
7	Conclusion	47
	Bibliography	49

List of Figures

- 1.1 A sample of the fabrics in our collected database ranked according to stiffness predicted by our model. The top panel shows physical fabric samples hanging from a rod. The bottom panel shows a slice through time of the x-axis when the fabrics are blown by the same wind force. Bendable fabrics generally contain more high frequency motion than stiff fabrics. 15
- 1.2 Visual cues from a static image can often tell you a lot about the material in the scene. Image (a) has many small folds, which suggests it is a flexible material. Image (d) has no small folds and more large folds. This suggests image (d) it is a stiffer material. However, sometimes visual cues can be misleading, such as in Image (b). Although this image is silk its pattern and drape makes it seem like it has material properties closer to wool (such as in Image (c)). In these cases a video may help to disambiguate the material properties. Images were taken from the Nordstrom website. 16
- 2.1 Texture must be analyzed with respect to a frame of reference. Here we show three resolutions of the same image. Although they come from the same stochastic process, they are perceptually distinguishable textures. 21

3.1	Illustration of our framework for estimation of material properties through video. The input to our system is a video containing fabric (a) along with a mask of what pixels contain the material. The masked magnitude of motion is extracted from the video of moving fabric via optical flow (b). Statistics are computed from the masked magnitude of the motion (c). These statistics are computed on a decomposition of the motion into sub-bands associated with concentric spheres in the frequency domain. These statistics are used as features for our model. PCA is then used to reduce feature dimensionality (d). These features are fed into a regression model that predicts the material properties of the fabric in the input video. The regression model was trained using features extracted from videos of other fabric where ground truth was available (e). This model is used to estimate material properties of the fabric in the input video (f).	24
3.2	A slice of a video in space and space-time (a) along with the associated mask (b) and magnitude of motion in the video (c). Note that different parts of the fabric move at different speeds at an instant in time.	25
3.3	A horizontal space \times time slice of a sample of the fabrics in the collected databased ordered by increasing stiffness and mass. Note how the magnitude of frequencies in the motion of the fabric changes as the material properties change. Specifically, less high frequency motion is seen in the more stiff and heavier fabrics.	26
3.4	The local magnitude of a 4-scale decomposition of a video using the 3D complex multi-resolution pyramid. The top panel shows the magnitude of a frame of the decomposition. The bottom panel shows a slice through time of the x-axis. White indicates large magnitude. Note that low frequency motion is pronounced in the courser resolution sub-bands whereas higher frequency motion is pronounced in the finer resolution sub-bands.	27
3.5	For completeness we show the local phase of a 4-scale decomposition of a video using the 3D complex multi-resolution pyramid. The top panel shows the phase of a frame of the decomposition. The phase displayed ranges from 0 to 2π	28
4.1	Each fabric was exposed to three different intensities of wind forces of varying intensity. Here is an example of a horizontal space \times time slice of the same fabric exposed to the three different types of wind. Note that the motion of the cloth looks very different under the different wind conditions.	32
4.2	A dataset containing videos of the fabrics moving in response to 3 different wind forces are collected, including (a) RGB and (b) Depth information for 30 different fabrics	33

- 4.3 An example of videos taken from our dataset for simulated fabric. Each simulated fabric was exposed to three different types of wind forces of varying intensity. Here we show a frame from each video along with a horizontal space \times time slice of identically parameterized simulated fabric exposed to the three different types of wind. The rendered texture on the fabric is different even though they are parameterized identically so that algorithms focus on characterizing the motion rather than classifying a fabric's printed pattern. A pattern is rendered on the fabric so that accurate optical flow can be computed from the videos. 34
- 5.1 Experiment setup of pairwise comparisons of material properties (stiffness or mass) from video stimuli. Subjects were asked to compare material properties of the two fabrics on a 7 point scale. A similar setup was also used to compare the stiffness and mass of fabrics given image stimuli. 37
- 5.2 Comparisons of perceptual responses versus measurements of pairwise difference in the log of material properties. (a) Difference of the log of measured stiffness between two fabrics is plotted against the average perceptual difference for the same fabric pair using video stimuli. Red stars correspond to stimuli pairs that had the same wind force applied to both fabrics and blue circles represent stimuli pairs that contained fabrics exposed to different wind forces. (b) Difference of the log of measured stiffness between two fabrics is plotted against the average perceptual difference between two fabrics when using image stimuli. (c) and (d) contain similar scatter plots of the difference of the log of measured mass between two fabrics versus average perceptual difference between the same fabric pair for both video and image stimuli respectfully. 38
- 5.3 Perceptual differences for fabrics exposed to the same wind intensity and the perceptual difference measurements for fabrics exposed to different wind intensities. This suggests that humans are partially invariant to the forces in estimating the stiffness and mass of moving fabric from video. 39
- 5.4 Perceptual comparisons of cloth stiffness and mass were made between pairs of fabrics shown under the same wind conditions (horizontal axis) and under different wind conditions (vertical axis). Plots show the correlation between those perceived differences in stiffness (a) and mass (b). The strong correlation revealed in both plots shows that shows that the subjects were substantially invariant to the wind conditions when comparing the physical properties of different fabrics. 40
- 6.1 Fabrics omitted from results due to lack of texture or specularities that causes inaccurate optical flow estimates of the motion. Fabrics (a)-(e) were omitted due to their lack of texture. Fabric (f) was omitted since its pattern is so tight that it causes aliasing. Fabric (g) was omitted due to specularities. 41

- 6.2 Comparisons of model predictions for pairwise differences in the log of material properties (a) against perceptual differences for stiffness (b) against ground truth differences of the log of stiffness (c) against perceptual differences for mass (d) against ground truth differences of the log of mass. These plots show that our estimates of the material properties of the fabric in a video are well correlated with both the ground truth material property values and the perceptual estimates of material properties. Red stars (*) indicate differences for pairs exposed to the same force, and blue circles (o) indicate differences for fabric pairs under different forces. The correlation in these two cases are very similar, suggesting that our model is partially invariant to changes in force when estimating the intrinsic material properties of fabrics from videos. The range of values were scaled from 0 to 1 for display purposes. Plots (a) and (c) contain less data points since only 15 of the 30 fabrics were used in the perceptual study. 43
- 6.3 Comparison of pairwise differences predicted by our model (blue stars) with perceptual estimates (red line, with red shade representing a standard deviation). For both (a) stiffness and (b) mass, our estimated difference often falls within the perceptual score's standard deviation. . . . 44
- 6.4 Predicted comparisons of cloth stiffness and mass were made between pairs of fabrics shown under the same wind conditions (horizontal axis) and under different wind conditions (vertical axis) by computing the difference in predicted material properties from our model. Plots show the correlation between those predicted differences in stiffness (a) and mass (b). The strong correlation revealed in both plots shows that shows that our model is substantially invariant to the wind conditions when comparing the physical properties of different fabrics. 44
- 6.5 Fabrics in our collected database ranked (in row-major order) according to stiffness predicted by our model. The top image of each row shows the fabric samples hanging from a rod. The bottom image shows a slice through time of the x-axis when the fabrics are blown by the same wind force. The ground truth and predicted measurements of stiffness are shown below each image pair. 45
- 6.6 Fabrics in our collected database ranked (in row-major order) according to mass predicted by our model. The top image of each row shows the fabric samples hanging from a rod. The bottom image shows a slice through time of the x-axis when the fabrics are blown by the same wind force. The ground truth and predicted measurements of mass are shown below each image pair. 46

List of Tables

4.1	Table containing a summary of what is provided in the real and simulated fabric datasets in our database. This database will be made publicly available at the authors website (http://people.csail.mit.edu/klbouman).	31
4.2	Measured properties for the 30 fabrics in our database. Properties include, bending stiffness, area weight, mass, and also density (not shown). The average value of these properties across multiple samples is displayed in the table. An * in the description of a fabric indicates that only one measurement was made for this fabric. Most fabric's properties are averaged across two measurements made in orthogonal directions.	35
6.1	Percentage error and ensemble percentage error calculated for stiffness and mass. Percentage error is calculated by taking the average percentage difference between a predicted log measurement for each video segment and all log measurements for a specific fabric. Ensemble percentage error is calculated by finding the average percentage difference between the average predicted log measurement and average log measurement for each fabric.	42

Introduction



Figure 1.1. A sample of the fabrics in our collected database ranked according to stiffness predicted by our model. The top panel shows physical fabric samples hanging from a rod. The bottom panel shows a slice through time of the x-axis when the fabrics are blown by the same wind force. Bendable fabrics generally contain more high frequency motion than stiff fabrics.

AUTOMATIC scene understanding is a fundamental goal of computer vision research. Although the computer vision community has made great strides in the last couple decades towards achieving scene understanding from image or video data with work in object detection, 3D reconstruction, etc., there has been very little work on understanding the material properties of objects in a scene. For instance, is an object hard or soft, rough or smooth, flexible or rigid? Humans passively estimate the material properties of objects on a daily basis. Designing a system to estimate material properties of objects in a scene without invasive interactions is a difficult problem that is essential for scene understanding. Knowing the material properties of objects in a scene allows one to have a better understanding of how objects can and will interact with their environment. This can be very useful for applications such as robotics, online shopping, material classification, material editing, and predicting the object’s behavior under different applied forces.

In this thesis we focus on passively estimating the material properties of fabrics by observing a video of the fabric exposed to unknown wind forces. Our focus on fabric

is motivated by two reasons. First, a number of metrics exist to describe the intrinsic material properties of fabric. These metrics can be measured using setups such as the Kawabata system [8]. Second, fabric is intrinsically a 2D material, making most of its motion easily visible from video.

Humans use cues from the visual appearance of fabric to give them an understanding of the fabric’s material properties. Perhaps evidence of the importance of the visual appearance is in the way that, with just a simple, static sketch, an artist is able to convey the properties of a material by focusing on the way that the material drapes and folds. However, often static cues are not enough to disambiguate material properties of fabric (Refer to Figure 1.2). Cues from a video of moving fabric help humans make further judgements about the material properties of the fabric in a scene. We discuss this more in Chapter 5. For this reason, recently, a number of online-shopping sites have started showing short videos of models moving in their clothing. These videos show the customer how the clothes move and cling to the body. This information, along with a customer’s prior knowledge about the forces they expect to be exerted on the clothing, help a customer make judgments about the fabric’s material properties.



Figure 1.2. Visual cues from a static image can often tell you a lot about the material in the scene. Image (a) has many small folds, which suggests it is a flexible material. Image (d) has no small folds and more large folds. This suggests image (d) it is a stiffer material. However, sometimes visual cues can be misleading, such as in Image (b). Although this image is silk its pattern and drape makes it seem like it has material properties closer to wool (such as in Image (c)). In these cases a video may help to disambiguate the material properties. Images were taken from the Nordstrom website.

■ 1.0.1 Contributions

The motion of a fabric is determined by resistance to bending, stretching, shearing, external forces, aerodynamic effects, mass, friction, and collisions [2]. In this work we restrict our attention to recovering two of the most distinctive properties of fabric in natural scenes - the bending stiffness and mass. We aim to develop video features that correspond to successful estimations of the intrinsic material properties of fabric. We

believe this work is the first attempt to passively estimate material properties of fabric from video when the fabric is moving in a simple natural scene due to unknown forces.

Additionally, we have collected a database that can be used by researchers to further explore this problem. This database contains RGB and RGBD videos of real fabric with corresponding measured material properties as well as a dataset containing the mesh and rendered video of animated computer generated fabric. We also present a perceptual study of humans' ability to estimate the material properties of fabric from videos and images.

■ 1.0.2 Thesis Outline

This thesis is organized as follows. In Chapter 2 we provide a background of previous, applicable work. Chapter 3 presents our algorithm for predicting the material properties of fabric. Chapter 4 describes the database we have collected for training and testing of our algorithm. Chapter 5 describes a perceptual study on how well humans are able to estimate the material properties of fabric from video and image data. Chapter 6 contains results of our algorithm and a discussion of the results. We conclude with Chapter 7.

Background

RELATED previous work has focused on understanding static properties of materials, such as surface reflectance [18], material category [11], and roughness [4]. In contrast, we address the problem of passively estimating material properties of objects that are evident through dynamic motions. In this work we use simple videos of materials subjected to unknown forces in order to passively estimate the material properties of the objects.

In this chapter we review a number of papers from the graphics community that deal with modeling fabric and recovering the material properties of fabric through both invasive and passive methods. Our algorithm to passively estimate material properties from video makes use of features previously developed to compactly represent static image textures in order to describe video textures - such as is seen in fabric motion. Thus, in this chapter we also review papers related to parametric texture models.

■ 2.0.3 Material Properties of Fabric

Cloth Models

Most of the work on cloth animation in the graphics community has been on developing models that are both realistic looking and fast to render. Baraff and Witkins seminal paper describing a cloth model that uses stiff springs with implicit time integration to simulate fabric allowed physics-based cloth models to become feasible to render [1]. Since then, a multitude of models have been developed to animate fabric. However, although these models may produce reasonably realistic renderings of fabric, they are not necessarily able to reproduce the behavior of real fabric.

The work of Miguel et al. used data driven techniques to evaluate a number of common cloth models (spring-mass, soft constraints, and the StVK models) for the purposes of reproducing the behavior of real cloth [12]. A dedicated setup was used to measure the behavior of the cloth in response to a number of different forces and then fit the best set of parameters to the gathered data for each model. Their work showed that each model contains limitations in reproducing the behavior of real cloth - the most pronounced of errors coming from modeling the fabric as being elastic. However, the spring-mass cloth model remains a fairly simple model that does a comparably good job at mimicking the behavior of fabric.

Invasively Fitting Cloth Model Parameters for Real Fabric

In order to fit the behavior of real cloth, parameters of a model must be physically measured using expensive and time-intensive systems. These systems use invasive interactions to precisely measure a fabric's response to many different forces. These measurements can then be used to optimize the parameters of a cloth model.

The most well known setup used to measure these parameters is the Kawabata evaluation system [8]. The Kawabata system performs six tests to measure tensile strength, shear strength, bending stiffness, compression, surface friction, and roughness of the fabric. Since the development of the Kawabata system other systems have been developed to measure the properties of fabric by physically manipulating the fabric [20].

In this work we have created a database of thirty fabrics with their associated bending stiffness. The bending stiffness was measured using techniques described in [19]. The setup works by applying a load to a hanging fabric specimen in order to deform it. The specimen was then imaged and the bending stiffness of the fabric was determined through an iterative algorithm using a finite element model.

Passively Fitting Cloth Model Parameters for Real Fabric

Jojic and Huang attempted to estimate the cloth draping parameters from the 3D range data of a static scene of cloth draped over an object [5]. However, because the cloth was static, the dynamic properties of the cloth are not able to be estimated using this system. Additionally, the system needs very accurate 3D information in order to perform the inference.

Bhat et al. [2] presented a method to estimate the simulation parameters for cloth by using video rather than a still image; thus, they were able to accurately recover a cloth parameter that captured some of the dynamic properties of the cloth. Since they were using video of moving fabric as input into their system, they were also able to estimate properties, such as air drag and damping, that are not able to be measured using setups such as the Kawabata system. However, this system also contains many limitations. The system presented in [2] requires a controlled setup of structured light projected onto the fabric and only allows movement due to a known gravitational force. Such a system is inapplicable to our more general problem. Instead, we wish to estimate these parameters of the cloth in a more natural setting where the cloth is exposed to unknown forces.

■ 2.0.4 Texture

Texture is generally thought of as a sample from an underlying stochastic process - usually with some repetitive or periodic element. This stochastic process can have a varying degree of randomness. Although two images from the same stochastic process may not be identical, they may be pre-attentively indistinguishable to humans. Texture analysis focuses on extracting a number of features from two images that are equivalent if and only if the images come from the same texture process. In this case, the set




Figure 2.1. Texture must be analyzed with respect to a frame of reference. Here we show three resolutions of the same image. Although they come from the same stochastic process, they are perceptually distinguishable textures.

of statistical measurements are both sufficient and necessary to guarantee perceptual equivalence. It is even more desirable to extract features that have a monotonic relationship between measured and perceived similarity [14]. In this case, two textures that are perceptually similar would have more similar features than textures that are perceptually dissimilar.

Characterizing visual textures as a stationary random field described through a set of statistics was proposed by Julesz in [6, 15]. Julesz hypothesized that the N th-order joint empirical densities of the image pixels could be used to cluster textures into perceptually indistinguishable classes. More precisely, he predicted that there exists a set of k functions, ϕ_i for $i \in \{1, \dots, k\}$, such that for any two random fields, X and Y , if

$$\begin{aligned} \mathbb{E}[\phi_k(X)] &= \mathbb{E}[\phi_k(Y)] \\ \forall k &\Rightarrow \text{samples of } X \text{ and } Y \text{ are} \\ &\text{perceptually equivalent.} \end{aligned} \tag{2.1}$$

Julesz originally conjectured that only identical second order statistics were necessary to cluster pre-attentively perceptually indistinguishable textures [6]. However, he later disproved this conjecture and showed that textures with identical third and even fourth order statistics can be pre-attentively perceptually distinguishable [7].

Parametric Texture Models

Heeger and Bergen proposed a feature set to characterize a 2D visual texture by using the first order statistics of a chosen set of linear filter outputs [3]. A fundamental characteristic of texture is it cannot be analyzed without respect to a frame of reference. For instance, Figure 2.1 contains different resolutions of the same image. Although all images come from an overarching stochastic process, they are perceptually distinguishable. Heeger and Bergen used a multi resolution steerable pyramid to decompose each

image into different spatial resolution sub-bands. A histogram was then computed for each sub-band and used as the feature set characterizing the texture. This feature set was then used to synthesize new textures similar to the original texture. Although the technique worked well for stochastic homogenous textures, it was limited to synthesising images with only minimal structure.

In response, Portilla and Simoncelli developed a parametric statistical model for 2D visual textures in [15]. In this work, a complex steerable pyramid was used to decompose the image into different spatial resolution sub-bands. This allowed them to characterize not only the local magnitude, but also the phase of the texture. By designing a method to efficiently synthesize random image samples subject to a set of constraint functions, Portilla and Simoncelli were able to identify a minimal set of statistical constraints necessary for perceptual equivalence in 2D visual textures. This set of minimal constraints parameterizes a statistical model that can be fit to any 2D visual texture and then can be used as features for texture analysis. Their final feature set, in addition to containing first order statistics such as was proposed by Heeger and Bergen, included second order statistics.

Temporal Textures For Material Classification

Temporal textures are motions characterized as repetitive in both space and time. Classic examples include fire, steam, or falling snow. A majority of the motions that we encounter on a daily basis cannot be classified as temporal textures. For instance, a video of a person walking is periodic in time but spatially restrictive. Alternatively, a microscopic video of bacteria multiplying is repeated in space but not in time.

Algorithms have been used previously in the vision and graphics communities to synthesize or classify temporal textures. Rahman et al. used features extracted from videos to classify videos into a number temporal texture semantic categories including waving flags [16].

Approach

ONE goal of computer graphics is to create models of physical objects that can be used to synthesize realistic images and videos. In this work, we solve the inverse problem - derive a model and its parameters that fit the observed behavior of a moving deformable object in videos. A candidate solution is to use the same generative model to solve the inverse problem as is used in the forward rendering [1]. However, this would require us to first infer the geometry of the moving object at every instant in time before fitting an inverse model to the data. This intermediate inference step would both have a high computational cost and a large chance of introducing errors that an inverse model may be sensitive to. Thus, we look towards alternate methods to predict the material properties of a deformable object more robustly.

In this thesis, we use statistics characterizing temporal textures in order to predict the material properties of deformable objects. This differs from traditional temporal texture classification since we quantify a continuum of differences within the same class of textures rather than clustering different temporal textures into discrete classes.

We extend Portilla and Simoncelli's work to temporal textures for the application of inferring material properties to create a set of statistics characterizing perceptually indistinguishable motion. These statistics are then used as parameters to a regression model that estimates the material properties of a previously unseen object moving due to unknown forces. We apply this technique to videos of fabric moving under different wind forces. A flow diagram of our algorithm can be seen in Figure 3.1.

■ 3.0.5 Material Localization

We are interested in determining the material properties of a specific object. Thus, to begin, we must compute a mask that identifies what pixels in a video belong to the targeted object, and which do not. In this work we focus on identifying the material properties of fabric from our collected database. Videos in this database contain cloth with a green backdrop. This setup allows us to easily identify which regions are the targeted fabric, and which are not. We use simple color thresholding along with a 2D median filter to create a mask for each video in our database (Chapter 4). Figure 3.2b shows the masked region for a sample video.

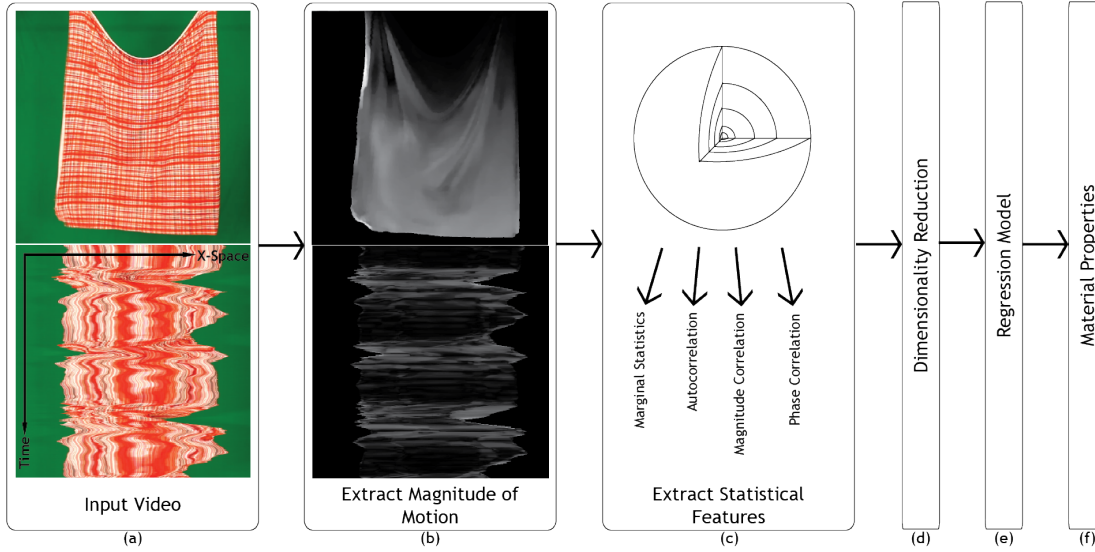


Figure 3.1. Illustration of our framework for estimation of material properties through video. The input to our system is a video containing fabric (a) along with a mask of what pixels contain the material. The masked magnitude of motion is extracted from the video of moving fabric via optical flow (b). Statistics are computed from the masked magnitude of the motion (c). These statistics are computed on a decomposition of the motion into sub-bands associated with concentric spheres in the frequency domain. These statistics are used as features for our model. PCA is then used to reduce feature dimensionality (d). These features are fed into a regression model that predicts the material properties of the fabric in the input video. The regression model was trained using features extracted from videos of other fabric where ground truth was available (e). This model is used to estimate material properties of the fabric in the input video (f).

■ 3.0.6 Motion Estimation

The 3D intensity values of a video contain information about both the appearance and motion in a scene. The printed texture of a material is less useful for the purpose of material property prediction since the pattern on a fabric does not, in general, affect the material properties of the underlying fabric. Therefore, we would like to separate the appearance of the printed texture from the motion signals in a video in order to just focus on characterizing the object’s motion.

Many methods can be used to extract the motion from video. Optical flow has been previously explored in the context of temporal textures for motion extraction [13, 16]. In our work we have also chosen to extract the motion using optical flow. Code provided with [10] was used to compute the optical flow vectors for each video of moving fabric. Once the optical flow has been computed, the magnitude of the flow is found and any region not containing the fabric is masked out by assigning it a flow magnitude of zero. Figure 3.2c shows the masked magnitude of flow for a sample video. Note that different parts of the fabric move at different speeds at an instant in time.




Figure 3.2. A slice of a video in space and space-time (a) along with the associated mask (b) and magnitude of motion in the video (c). Note that different parts of the fabric move at different speeds at an instant in time.

■ 3.0.7 Statistical Features

Once we have extracted the motion field from a video, our goal is to extract a set of features from the motion that are descriptive of the material properties. We hypothesize that if two videos contain perceptually indistinguishable motions under the influence of similar forces, then their material properties are similar. Thus, we would like to find a set of statistics such that two videos contain perceptually indistinguishable motions *if and only if* both videos' motion fields are drawn from random fields with matching statistics. In designing our feature set, we draw inspiration from Portilla and Simoncelli's constraints that were developed for synthesizing perceptually indistinguishable textures. Although these proposed statistics are not necessarily bidirectional, they empirically have been shown to produce similar statistics for perceptually indistinguishable textures.

First, we decompose our motion field using a 3D complex multi-resolution pyramid.

Similar to a 2D complex steerable pyramid, this pyramid uses a set of local filters to recursively decompose a video into sub-band videos at N_{sc} different spatiotemporal scales and N_{or} orientations; however, steerability does not hold in this representation. Each sub-band contains a local estimate of the magnitude and phase of the 3D signal around a pixel. In this thesis we have chosen to decompose the magnitude of our motion field into $N_{sc} = 4$ scales and $N_{or} = 1$ orientation. Figure 3.1c shows how the frequency domain is split up for our decomposition.

Decomposing the motion field in this way is desirable for our application because different material properties may be more pronounced in the features of the motion from different spatiotemporal scales. For instance, a fabric’s mass may have a larger effect on the motion in the higher spatiotemporal scales (i.e., at lower motion frequencies), whereas a fabric’s bending stiffness may have a larger effect on the motion in the lower spatiotemporal scales, (i.e., at higher motion frequencies). Figure 3.3 shows a space time slice of a sample of fabrics in our database ordered by increasing stiffness and mass. Note how the magnitude of frequencies in the motion of the fabric changes as the material properties change.

We define our statistics on coefficients of the decomposition of the magnitude of the motion field. Figures 3.4 and 3.5 show the local magnitude and phase for the decomposition of a video for the four sub-bands respectively.




Figure 3.3. A horizontal space \times time slice of a sample of the fabrics in the collected database ordered by increasing stiffness and mass. Note how the magnitude of frequencies in the motion of the fabric changes as the material properties change. Specifically, less high frequency motion is seen in the more stiff and heavier fabrics.

The following sections describe the statistical features we have computed from the coefficients of the decomposed motion field in order to characterize the motion of a




Figure 3.4. The local magnitude of a 4-scale decomposition of a video using the 3D complex multi-resolution pyramid. The top panel shows the magnitude of a frame of the decomposition. The bottom panel shows a slice through time of the x-axis. White indicates large magnitude. Note that low frequency motion is pronounced in the coarser resolution sub-bands whereas higher frequency motion is pronounced in the finer resolution sub-bands.

fabric.

Marginal Statistics

Statistics defined over the histogram of motion magnitudes in a video are a simple but very powerful feature to use in describing a motion field. Many texture analysis [3, 15, 21] and action recognition [17] algorithms have either used marginal statistics or histograms directly to characterize marginal distributions. We measure the variance, skew, kurtosis and range (minimum and maximum) of the motion magnitude. Additionally, the variance, skew, and kurtosis for each of the $N_{sc} = 4$ lowpass videos are computed from the complex 3D pyramid. The marginal statistics on the lowpass videos characterize the distribution of motion magnitudes at different spatiotemporal scales and provides information about how much the material moves at different frequency bands.

Autocorrelation

In order to capture the second order spatiotemporal distribution, or structure, in the motion field we include the autocorrelation of the spatiotemporal signal as a statistical

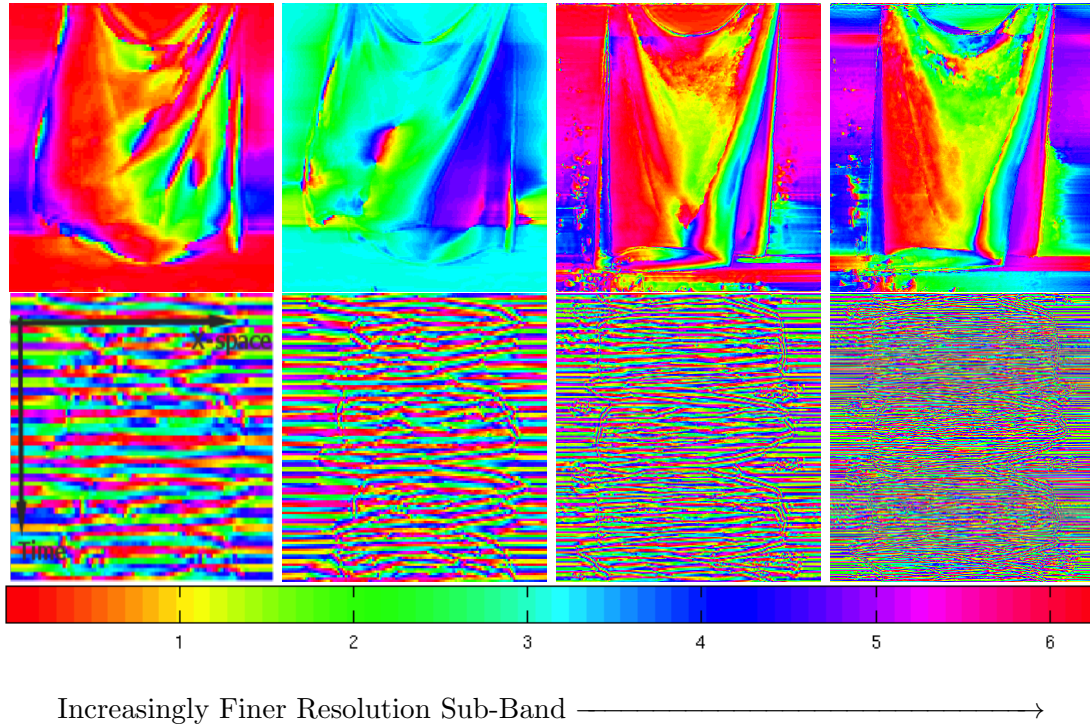


Figure 3.5. For completeness we show the local phase of a 4-scale decomposition of a video using the 3D complex multi-resolution pyramid. The top panel shows the phase of a frame of the decomposition. The phase displayed ranges from 0 to 2π .

feature. Julesz’s work in texture discrimination found that, although not always sufficient, second order statistics are often very important in guaranteeing pre-attentive perceptual equivalence of textures [6].

The circular autocorrelation for a 3D neighborhood of $N_l = 9$ pixels is computed for each of the $N_{sc} = 4$ lowpass videos. By using the same size neighborhood for the high and low spatiotemporal scales, the local autocorrelation captures higher spectral resolution in the lower spatiotemporal scales.

Magnitude Correlation

The correlation of the sub-band magnitudes of an image’s pyramid decomposition has been previously used to represent structures such as edges, bars, and corners in image textures [15]. Although bars and corners are rare in motion fields containing a single object, edges may occur due to occlusions. This is caused by the fabric moving at different speeds on either side of the occlusion. Thus, we include correlation of the decomposition’s neighboring sub-bands as a statistical feature of the motion field in a video. Capturing occlusions in space can be useful for identifying material properties such as stiffness; the less stiff a fabric is the more folds it generally contains. Before

computing the correlation of coefficients across scales we are required to upsample and interpolate the coarse scale sub-band to match the dimensions of the finer scale sub-band.

Phase Correlation

Local phase estimates of a signal indicate its gradient in a local region. In order to capture gradual gradient changes in the motion field, we compute the correlation across the local phases in the sub-bands of the video’s pyramid decomposition. Since the local phase changes twice as fast for each fine-scale sub-bands than it does for its neighboring coarse scale sub-band, we must first compensate for this by doubling the local phase of the coarse-scale coefficient before computing the cross-correlation. Additionally, similar to in magnitude correlation, in phase correlation we must upsample and interpolate the coarse scale sub-band to match the dimensions of the finer scale sub-band before computing the correlation.

■ 3.0.8 Model Learning

We aim to recover the underlying material properties from a video using the features described above. Specifically, we learn a function that maps the features to the log of ground truth stiffness and mass measurements (Section 4.0.9). We choose to work in the log domain since humans tend to be sensitive to the logarithm of material properties (Section 5.0.12), and the features we have chosen to use were initially developed for perceptual indistinguishability.

To force each feature type (eg. marginal statistics, autocorrelation, etc.) to contribute the same amount of variance in the feature vector, we normalize the feature vectors by subtracting the mean and forcing the variance of each feature to be proportional to the number of features in its feature type. Dimensionality of the whitened feature vectors is then reduced using PCA. Feature vectors were projected onto the eigenvectors which preserved 95% of the variance in the data. For feature vectors extracted from our fabric videos, this step reduces the size of the feature vector by an order of magnitude.

A simple linear regression model is used to map the resulting features to the ground truth material properties. We chose to use a linear regression rather than a more complex regression method to reveal the power in the features we have selected. To normalize for differences in sample sizes for different materials being analyzed, we add a weight to our regression model proportional to the number of samples containing the same material. Mathematically, we solve

$$W \odot Y = W \odot X\beta \tag{3.1}$$

for the weights β , given the dimensionality-reduced feature vectors X , log-domain ground truth measurements Y , and normalization weights W . Here, \odot denotes element-wise multiplication. This regression model is used in Chapter 6 to predict the material

properties of fabric.

Database

WE have compiled a database containing a video dataset of real fabrics and a video dataset of simulated fabrics along with their associated material properties. This database can be used for training and testing predictions of material properties. This database will be made publicly available on the author’s website¹. A summary of what is provided in this database can be seen in Table 4.1

■ 4.0.9 Real Fabric Dataset

We have compiled RGBD videos of 30 different types of real fabric along with measurements of their material properties. The fabrics span a variety of stiffness and area weights. Example categories include cotton, velvet, spandex, felt, silk, upholstery, wool, denim, and vinyl. Table 4.2 contains images and measurements of the fabrics in our

¹<http://people.csail.mit.edu/klbouman>

Real Fabric Dataset	Simulated Fabric Dataset
30 different fabrics cut to the same size	Script to render video & mesh of parameterized simulated fabric in response to a normal force or a parameterized wind force
Measured bending stiffness, density area weight, and mass for each fabric	
RGB video for fabric exposed to 3 wind intensities (858 × 850 pixels, 30 fps)	25 samples of simulated fabric parameterized by varying bending stiffness & mass
Kinect RGB video for fabric exposed to 3 wind intensities (640 × 480, 30 fps)	Rendered RGB videos of sample fabric exposed to 3 wind intensities (960 × 540 pixels, 24 fps)
Kinect Depth video for fabric exposed to 3 wind intensities (640 × 480, 30 fps)	Mesh of sample simulated fabric exposed to 3 wind intensities

Table 4.1. Table containing a summary of what is provided in the real and simulated fabric datasets in our database. This database will be made publicly available at the authors website (<http://people.csail.mit.edu/klbouman>).

database.

RGB-Depth Videos

Videos (858×850 pixel resolution) were recorded for all fabrics. Fabrics were hung from a bar and exposed to three different intensities of wind from an oscillating fan. The two-minute videos capture the fabrics moving in response to the wind force. Figure 4.1 shows a space time slice of the same fabric moving under the three wind forces. Note that the motion of the cloth looks very different under the different wind conditions.




Figure 4.1. Each fabric was exposed to three different intensities of wind forces of varying intensity. Here is an example of a horizontal space \times time slice of the same fabric exposed to the three different types of wind. Note that the motion of the cloth looks very different under the different wind conditions.

A green screen was placed behind the fabrics to allow for easy segmentation of the fabric from the background and for different complex backgrounds to be inserted behind the moving fabric. This could be used to create more challenging videos to process in the future.

RGBD Kinect videos of the scene were also recorded, providing a lower resolution RGB image along with a corresponding depth image at every frame. We have not used this data in this thesis, however this information could be used in the future to obtain motion along the depth dimension. Refer to Figure 4.2 for a sample RGB and depth frame.

All fabrics were cut to approximately 107×135 cm, and steamed to remove wrinkles. Cutting the fabrics to the same size reduces the problem of predicting the fabric surface weight to predicting the mass. Additionally, it removes any uncertainties due to scale that would confuse human observers or an algorithm. For instance, in a breeze a normal window curtain might move in a qualitatively different way than a curtain from a dollhouse even when cut from the same piece of fabric.




Figure 4.2. A dataset containing videos of the fabrics moving in response to 3 different wind forces are collected, including (a) RGB and (b) Depth information for 30 different fabrics

Material Property Measurements

The area weight, density, mass, and bending stiffness were measured for each fabric. Specimens of the fabric were sent to the Lowell Advanced Composite Materials and Textile Research Laboratory² in order to be tested for stiffness, density, and area weight. The stiffness of each fabric was measured by the methods presented in [19]. Briefly, each specimen was clamped in a fixture, a load was applied to it through a pulley system, and the magnitude that it deformed was recorded; two orthogonal directions for each fabric were tested in most fabrics. A finite element model (FEM) was fit to this data and yielded a bending stiffness estimate for the fabric. Across our fabrics the range of measured stiffness values was greater than an order of magnitude. Table 4.2 contains measurements of bending stiffness, area weight, and mass for each of the fabrics in our database.

■ 4.0.10 Simulated Fabric Dataset

A dataset of videos containing simulated fabric reacting to a variety of different forces was created. A simulated database allows us to have a collection of videos where we have precise measurements of the fabric parameters. Additionally, exact 3D motion can be recovered from this data by returning a mesh of the node locations of the simulated fabric over time. This dataset was not used in the work presented in this thesis.

Simulations of the fabric were rendered using the open source simulation software Blender. Blender uses a mass-spring system along with a number of parameters to define the behavior of a fabric mesh under a force field. These parameters include mass, bending stiffness, structural stiffness, spring damping, air damping, and fabric size. The simulated fabric was placed in an environment that parallels the data gathered for real fabric - the fabric mesh was hung and exposed to a wind force. Figure 4.3 shows renderings of the same fabric mesh exposed to the different force environments. A different

²<http://m-5.uml.edu/acmtrl/index.htm>




Figure 4.3. An example of videos taken from our dataset for simulated fabric. Each simulated fabric was exposed to three different types of wind forces of varying intensity. Here we show a frame from each video along with a horizontal space \times time slice of identically parameterized simulated fabric exposed to the three different types of wind. The rendered texture on the fabric is different even though they are parameterized identically so that algorithms focus on characterizing the motion rather than classifying a fabric’s printed pattern. A pattern is rendered on the fabric so that accurate optical flow can be computed from the videos.

image was rendered on each fabric mesh so that algorithms focus on characterizing the motion rather than classifying a fabric’s printed pattern.

In this work the mass and bending stiffnesses parameters of Blender’s model were varied. Video renderings of the behavior of the fabric mesh were collected. Additionally, the 3D positions of the mesh through time have been collected so that optical flow can be verified or initially avoided during algorithm development.

In addition to releasing a small dataset of videos that we have rendered, we have also created a script that allows a user to render fabric and return a mesh with any parameter set under two different types of environments: a wind force of parameterized intensity and the fabric falling over a sphere (normal force). This feature gives the user precise control over the fabric parameters that they would like the rendered fabric to have.

					
Description	Lycra	Faux Fur	Silk	Silk	Cotton
Bending Stiffness (lbf-in ²)	0.0250	0.0512	0.0075	0.0110	0.0100
Area Weight (oz/yd ²)	7.0600	13.5400	1.9750	2.6350	2.9400
Mass (g)	357	795	115	151	150
					
Description	Wool	Taffeta	Linen	Corduroy	Cotton
Bending Stiffness (lbf-in ²)	0.0375	0.0165	0.0225	0.0185	0.0225
Area Weight (oz/yd ²)	11.3400	3.3750	4.4750	4.5150	4.1045
Mass (g)	618	197	235	229	217
					
Description	Velvet	Fleece	Denim*	Upholstery	Upholstery
Bending Stiffness (lbf-in ²)	0.0200	0.0205	0.0250	0.0400	0.0750
Area Weight (oz/yd ²)	3.9250	5.6300	8.7500	9.8100	14.1600
Mass (g)	220	320	445	520	757
					
Description	Pleather*	Minky*	Damask Upholstry	Flannel Backed Vinyl*	Upholstry
Bending Stiffness (lbf-in ²)	0.1000	0.0600	0.0800	0.0350	0.0675
Area Weight (oz/yd ²)	16.6600	7.5100	14.6650	5.1400	10.2570
Mass (g)	854	393	782	247	546
					
Description	Outdoor Polyester*	Silk	Wool*	Canvas*	Nylon Rip Stop
Bending Stiffness (lbf-in ²)	0.0550	0.0300	0.0500	0.1000	0.0165
Area Weight (oz/yd ²)	6.1100	4.0935	9.4600	11.5900	1.7261
Mass (g)	332	222	510	621	93
					
Description	Terry Knit*	Lycra*	Laminated Cotton*	Lycra*	Upholstry
Bending Stiffness (lbf-in ²)	0.0330	0.0550	0.0500	0.0120	0.0875
Area Weight (oz/yd ²)	4.9850	6.6400	5.3200	5.2100	11.3900
Mass (g)	258	340	277	262	598

Table 4.2. Measured properties for the 30 fabrics in our database. Properties include, bending stiffness, area weight, mass, and also density (not shown). The average value of these properties across multiple samples is displayed in the table. An * in the description of a fabric indicates that only one measurement was made for this fabric. Most fabric's properties are averaged across two measurements made in orthogonal directions.

Human Material Perception

HUMANS naturally use their accumulated knowledge of mechanics to judge the material properties of objects in the world. We aim to understand material perception from a purely visual perspective. Psychophysical experiments were designed to measure how well observers are able to estimate the mass and stiffness of fabrics when observing video or image stimuli.

■ 5.0.11 Experimental Setup

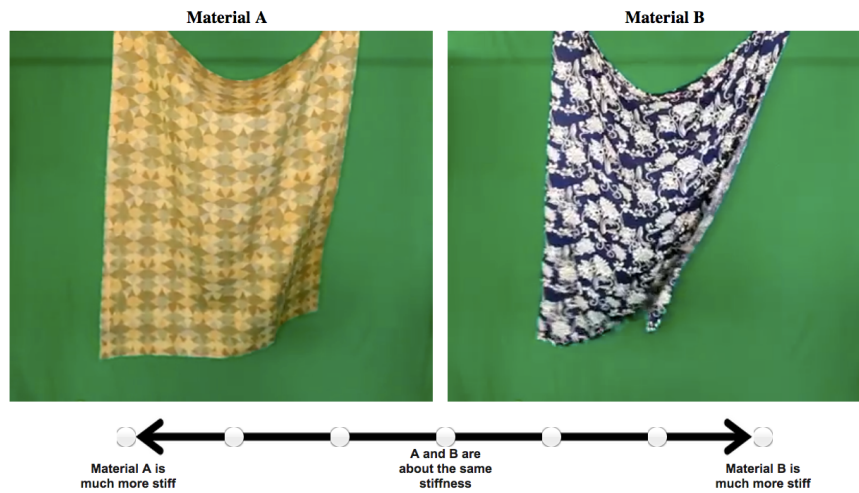


Figure 5.1. Experiment setup of pairwise comparisons of material properties (stiffness or mass) from video stimuli. Subjects were asked to compare material properties of the two fabrics on a 7 point scale. A similar setup was also used to compare the stiffness and mass of fabrics given image stimuli.

Stimuli included videos of 15 common fabrics exposed to 3 different intensities of wind. Specifically, we selected the fabrics corresponding to the first three rows of Table 4.2 as stimuli. The video data acquisition process is described in Section 4.0.9. A paired comparison method was used to measure perceived differences in the stiffness and mass between the fabrics in two videos [9]. The observer was shown two videos of different fabric stimuli moving by either the same or a different wind force and then was

asked to report which fabric was stiffer, the fabric in video A or B, by indicating their answer on a 7-point scale provided underneath the videos (Figure 5.1). Similarly, in a second experiment the observer was asked to report which video contained the heavier fabric on a 7-point scale.

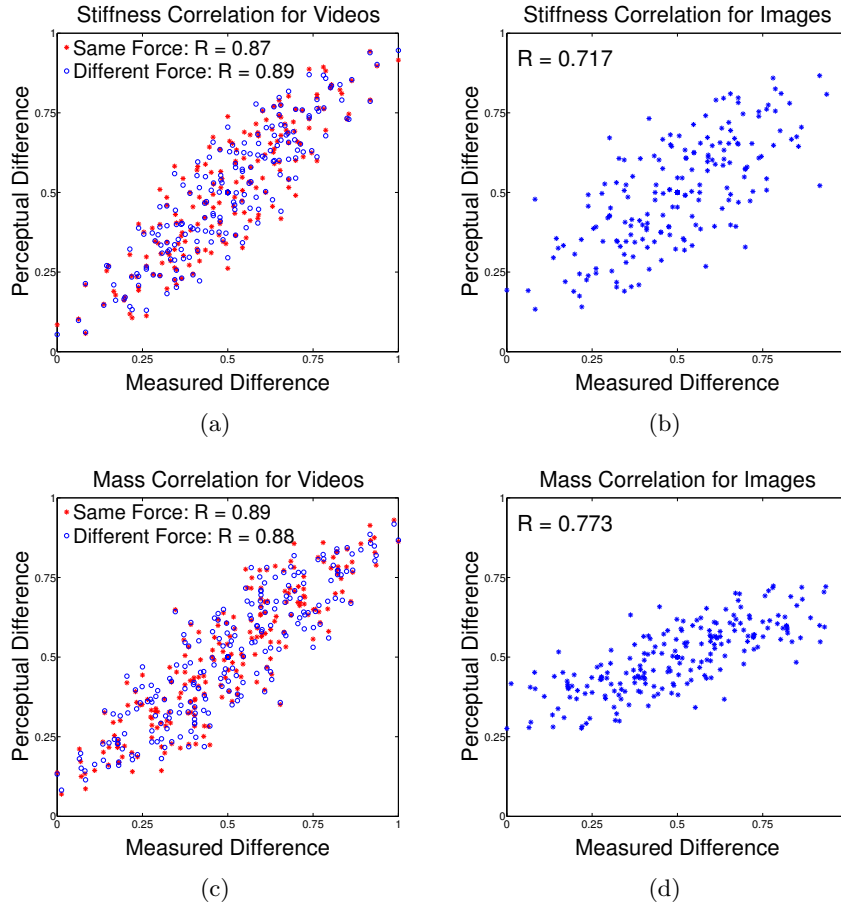


Figure 5.2. Comparisons of perceptual responses versus measurements of pairwise difference in the log of material properties. (a) Difference of the log of measured stiffness between two fabrics is plotted against the average perceptual difference for the same fabric pair using video stimuli. Red stars correspond to stimuli pairs that had the same wind force applied to both fabrics and blue circles represent stimuli pairs that contained fabrics exposed to different wind forces. (b) Difference of the log of measured stiffness between two fabrics is plotted against the average perceptual difference between two fabrics when using image stimuli. (c) and (d) contain similar scatter plots of the difference of the log of measured mass between two fabrics versus average perceptual difference between the same fabric pair for both video and image stimuli respectfully.

These experiments were conducted using Amazon Mechanical Turk. In order to maximize high quality responses, subjects were tested periodically throughout the experiment by asking them questions that they had previously been given the answer

to. Additionally, subjects were required to watch each pair of videos for 10 seconds before being allowed to respond. A total of 100 workers from Mechanical Turk (> 95% approval rate in Amazon's system) completed each experiment. Each subject answered an average of 250 questions.

A similar experimental setup was used to test the perception of mass and stiffness of 15 draped fabrics from a single still image. A sample image can be seen in Figure 4.2a. A total of 40 workers from Mechanical Turk (> 95% approval rate in Amazon's system) completed each experiment. Each subject answered 135 questions.

■ 5.0.12 Data Analysis and Discussion

Pairwise comparison scores from each experiment were scaled to take on values between 0 and 1. A score of 0.5 indicates that the two fabrics have the same stiffness/mass whereas a value of 0 or 1 indicates that material A or B is much stiffer/heavier than the other, respectively. A single comparative stiffness and mass score was computed for each pair of videos by averaging human responses.

In accordance with Weber's Law, we found that human responses to the pairwise comparison tasks were correlated with the difference of the log measured values (stiffness and mass). Figure 5.2 compares the difference of log measurements in the fabric stiffness and mass with the perceptual difference of the same material property. In these plots, the range of differences in the log of measured values were scaled to the range [0, 1] in order to match the range of perceptual scores.

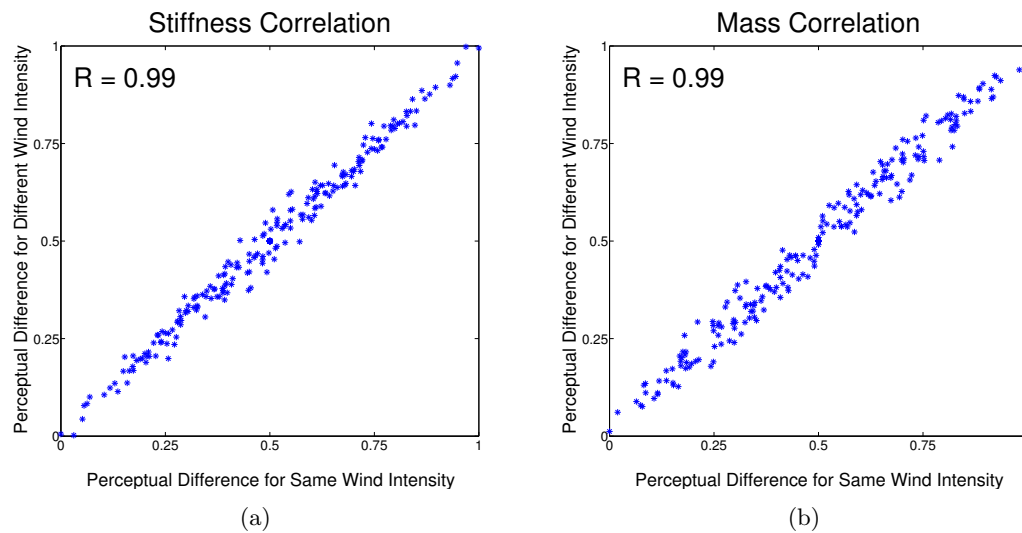


Figure 5.3. Perceptual differences for fabrics exposed to the same wind intensity and the perceptual difference measurements for fabrics exposed to different wind intensities. This suggests that humans are partially invariant to the forces in estimating the stiffness and mass of moving fabric from video.

Figures 5.2a and 5.2b compare the measured differences in stiffness and mass against

humans' perception of the same material properties when subjects observed the fabrics moving under wind forces. Figures 5.2b and 5.2d compare these differences when subjects observed only a static image of the draped fabrics. We find that humans' perception of stiffness and mass are very well correlated with ground truth measurements of these properties when the subjects observed the fabric moving. However, when the subjects were asked to judge the material properties from a single image, the correlation decreased substantially. These observations suggests that human observers use motion cues in videos to estimate material properties.

In Figures 5.2a and 5.2c we separate the cases where the same wind force intensity was applied (red star) to both fabrics and when different force intensities were applied (blue circle) to the two fabrics. The correlation in these two cases are very similar, suggesting that human observers are partially invariant to changes in force when estimating the intrinsic material properties of fabrics from videos. Figure 5.3 shows this point more clearly. For every pair of fabrics, the average perceptual difference is compared for when the fabrics were exposed to the same intensity of wind versus different intensities of wind. These plots show that the perceptual scores are highly correlated. Thus, subjects judged fabric's relative material properties similarly when the fabrics were in the same and different environments.




Figure 5.4. Perceptual comparisons of cloth stiffness and mass were made between pairs of fabrics shown under the same wind conditions (horizontal axis) and under different wind conditions (vertical axis). Plots show the correlation between those perceived differences in stiffness (a) and mass (b). The strong correlation revealed in both plots shows that the subjects were substantially invariant to the wind conditions when comparing the physical properties of different fabrics.

To evaluate human consistency, for each experiment, we randomly split responses into two independent halves and quantified how well pairwise scores measured from the first half of responses matched pairwise scores measured from the second half of responses (Figure 5.4). Humans showed a high consistency in their responses to the different stimuli.

Algorithm Results and Discussion

IN this thesis, our goal was to develop a set of features that enable successful estimation of the intrinsic material properties of a fabric in the presence of unknown forces. In Section 3.0.7 we introduced candidate features for use in predicting material properties. In this section, we evaluate the power of these features in predicting intrinsic material properties. Specifically, we fit a linear regression model to ground truth measurements and compare predictions of properties for previously unseen fabrics to both ground truth and perceptual estimates.

■ 6.0.13 Implementation Details

Twenty-three of the 30 fabrics in our database were selected for training and testing of our model. Seven of the fabrics, which either lacked texture or caused specularities in the videos, were removed from our results due to inaccurate optical flow estimates of the motion. Figure 6.1 contains images of the omitted fabrics.




Figure 6.1. Fabrics omitted from results due to lack of texture or specularities that causes inaccurate optical flow estimates of the motion. Fabrics (a)-(e) were omitted due to their lack of texture. Fabric (f) was omitted since its pattern is so tight that it causes aliasing. Fabric (g) was omitted due to specularities.

For each video we extracted multiple non-overlapping video segments, each 512 frames long. A single feature vector was computed for each segment. The linear regression model described in Section 3.0.8 was then used to learn a mapping from the feature vectors to the log of ground truth measurements. In the cases where a single fabric contains multiple ground truth measurements, we map each feature vector corresponding to that fabric to each of the collected measurements. We used a leave-one-out method for training the model and predicting the material properties of the fabric in each video

	Percentage Error	Ensemble Percentage Error
Stiffness	18.8%	17.2%
Mass	16.7%	12.6%

Table 6.1. Percentage error and ensemble percentage error calculated for stiffness and mass. Percentage error is calculated by taking the average percentage difference between a predicted log measurement for each video segment and all log measurements for a specific fabric. Ensemble percentage error is calculated by finding the average percentage difference between the average predicted log measurement and average log measurement for each fabric.

segment. More specifically, when making a prediction using a feature vector associated with a fabric, all feature vectors extracted from video segments corresponding to the same fabric were removed from the training set.

■ 6.0.14 Results

The predicted stiffness values from our regression model are shown in Figure 6.5 along with the ground truth measurements. The images of each fabric are sorted in row-major-order according to their predicted stiffness value from our model. The predicted stiffness value displayed was found by averaging the predicted stiffness values from all video segments containing the same fabric. Similarly, Figure 6.6 displays the same results for the mass of the fabrics.

We compare predicted measurements of stiffness and mass to ground truth measurements (Section 4.0.9) and perceptual measurements (Chapter 5) by computing the Pearson product-moment correlation coefficient or R value. Correlation of our predicted measurements for pairwise differences are plotted against perceptual differences and ground truth differences for both the stiffness and mass of fabrics in Figure 6.2. Red stars (*) indicate difference values for stimuli pairs that were exposed to the same force and blue circles (o) indicate difference values for stimuli pairs that were exposed to different forces. Percentage error for stiffness and mass for our results can be seen in Table 6.1.

■ 6.0.15 Discussion

Figure 6.2 shows that our estimates of the material properties of the fabric in a video are well correlated with both the ground truth material property values and the perceptual estimates of material properties. Individual predictions for each fabric are compared to ground truth measurements in Figures 6.5 and 6.6. Although our model is not able to recover specific values of material properties, our model is able to estimate a general trend of increasing stiffness or mass.

Figures 6.3a and 6.3b show how our model’s predictions of material properties for a pair of stimuli compare to the distribution of human responses for the same stimuli pair. Blue stars (*) indicate our model’s average prediction of the log of the material properties when the range of our estimated differences are scaled to lay in the same

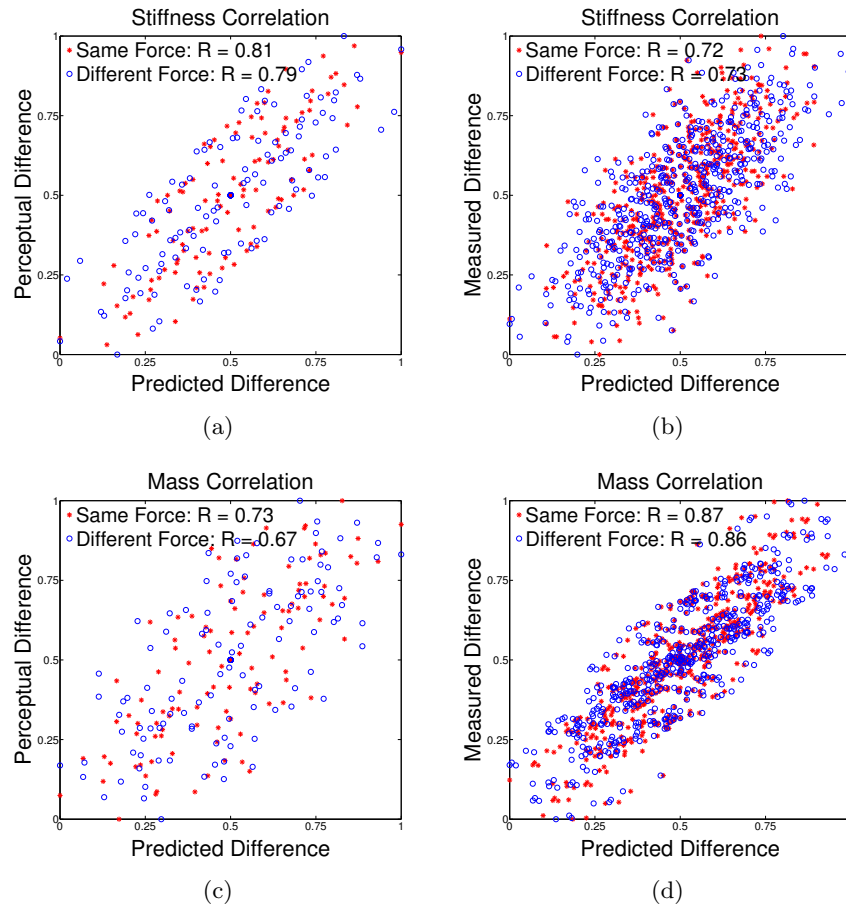


Figure 6.2. Comparisons of model predictions for pairwise differences in the log of material properties (a) against perceptual differences for stiffness (b) against ground truth differences of the log of stiffness (c) against perceptual differences for mass (d) against ground truth differences of the log of mass. These plots show that our estimates of the material properties of the fabric in a video are well correlated with both the ground truth material property values and the perceptual estimates of material properties. Red stars (*) indicate differences for pairs exposed to the same force, and blue circles (o) indicate differences for fabric pairs under different forces. The correlation in these two cases are very similar, suggesting that our model is partially invariant to changes in force when estimating the intrinsic material properties of fabrics from videos. The range of values were scaled from 0 to 1 for display purposes. Plots (a) and (c) contain less data points since only 15 of the 30 fabrics were used in the perceptual study.

range as the perceptual scores. The red region defines one standard deviation around the mean response (red line) of human observers for each pair or stimuli. These plots show that although humans on average are able to make accurate predictions about relative material properties, there is a large variance in their responses.

Figures 6.4a and 6.4b show the correlation of predicted difference measurements for fabrics exposed to the same wind intensity and the predicted difference measurements

for fabrics exposed to different wind intensities. The correlation for both stiffness and mass is very high. This shows that our model is somewhat invariant to the forces in recovering material properties of fabric.

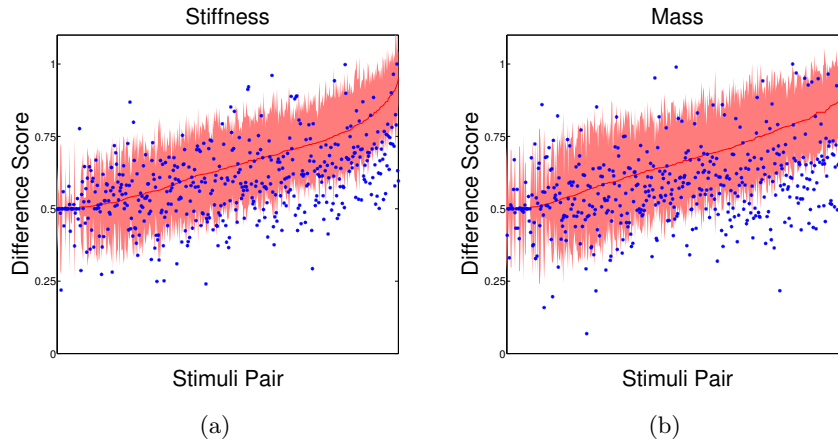


Figure 6.3. Comparison of pairwise differences predicted by our model (blue stars) with perceptual estimates (red line, with red shade representing a standard deviation). For both (a) stiffness and (b) mass, our estimated difference often falls within the perceptual score's standard deviation.

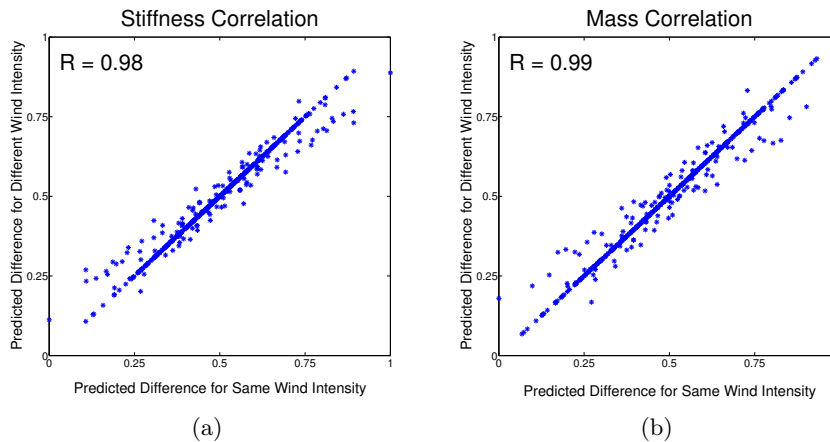


Figure 6.4. Predicted comparisons of cloth stiffness and mass were made between pairs of fabrics shown under the same wind conditions (horizontal axis) and under different wind conditions (vertical axis) by computing the difference in predicted material properties from our model. Plots show the correlation between those predicted differences in stiffness (a) and mass (b). The strong correlation revealed in both plots shows that our model is substantially invariant to the wind conditions when comparing the physical properties of different fabrics.

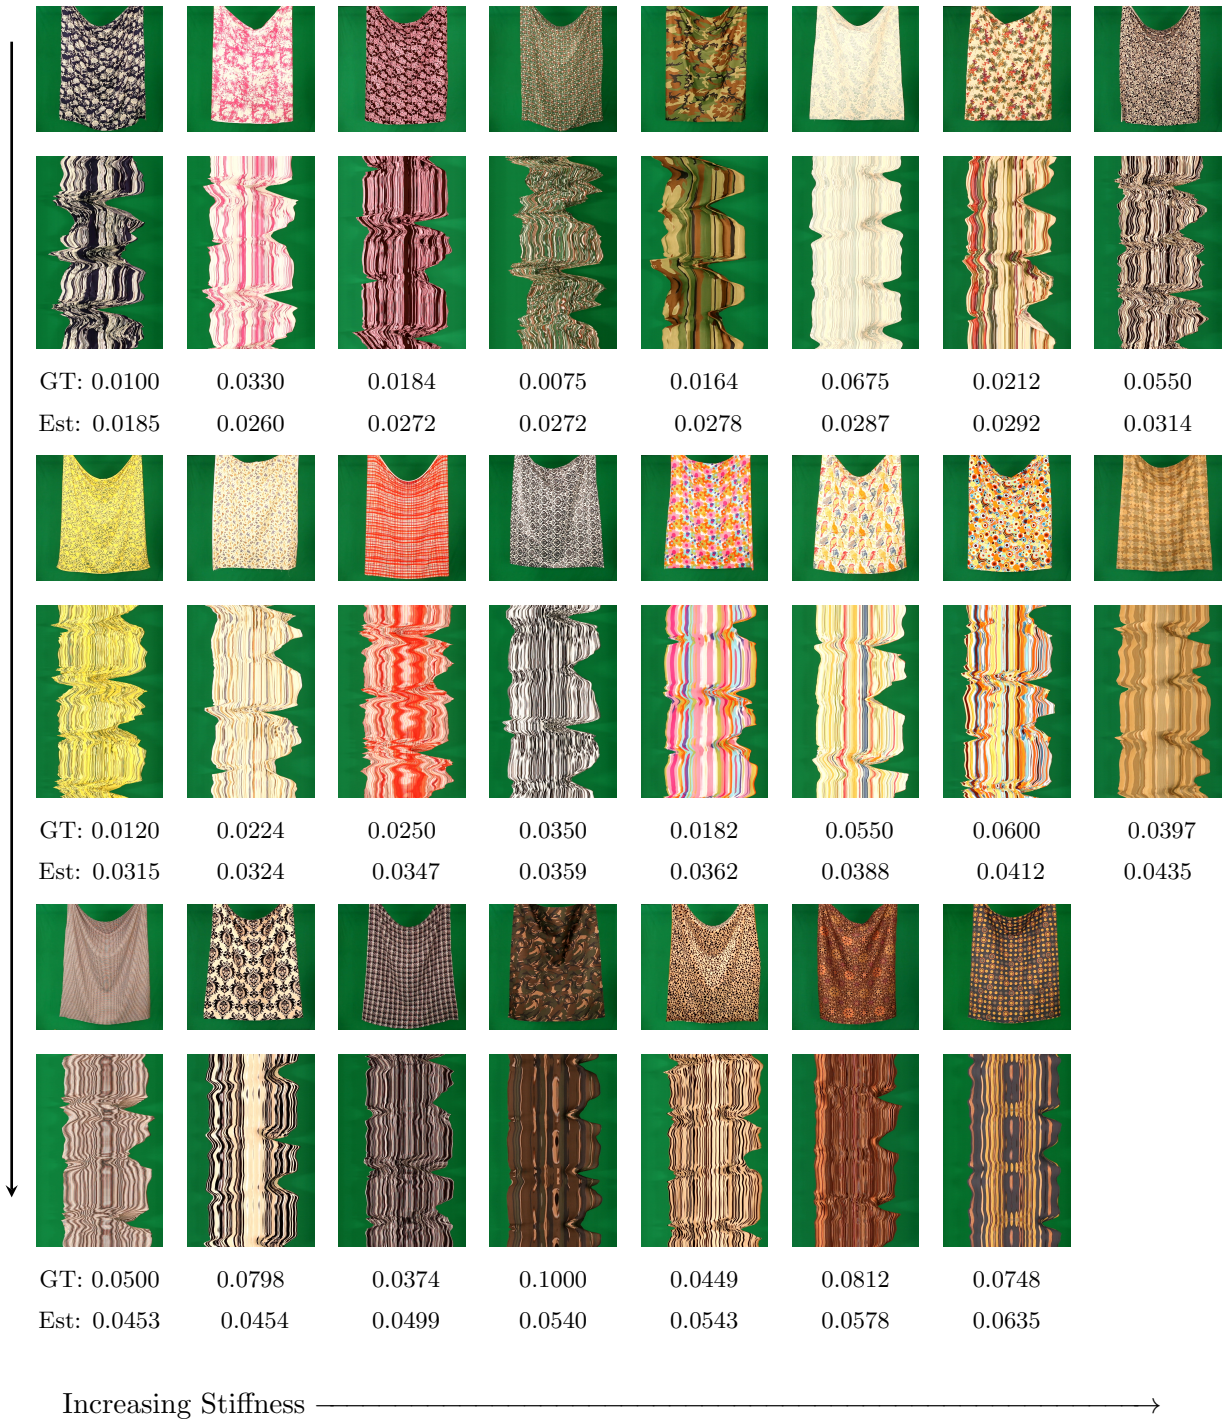
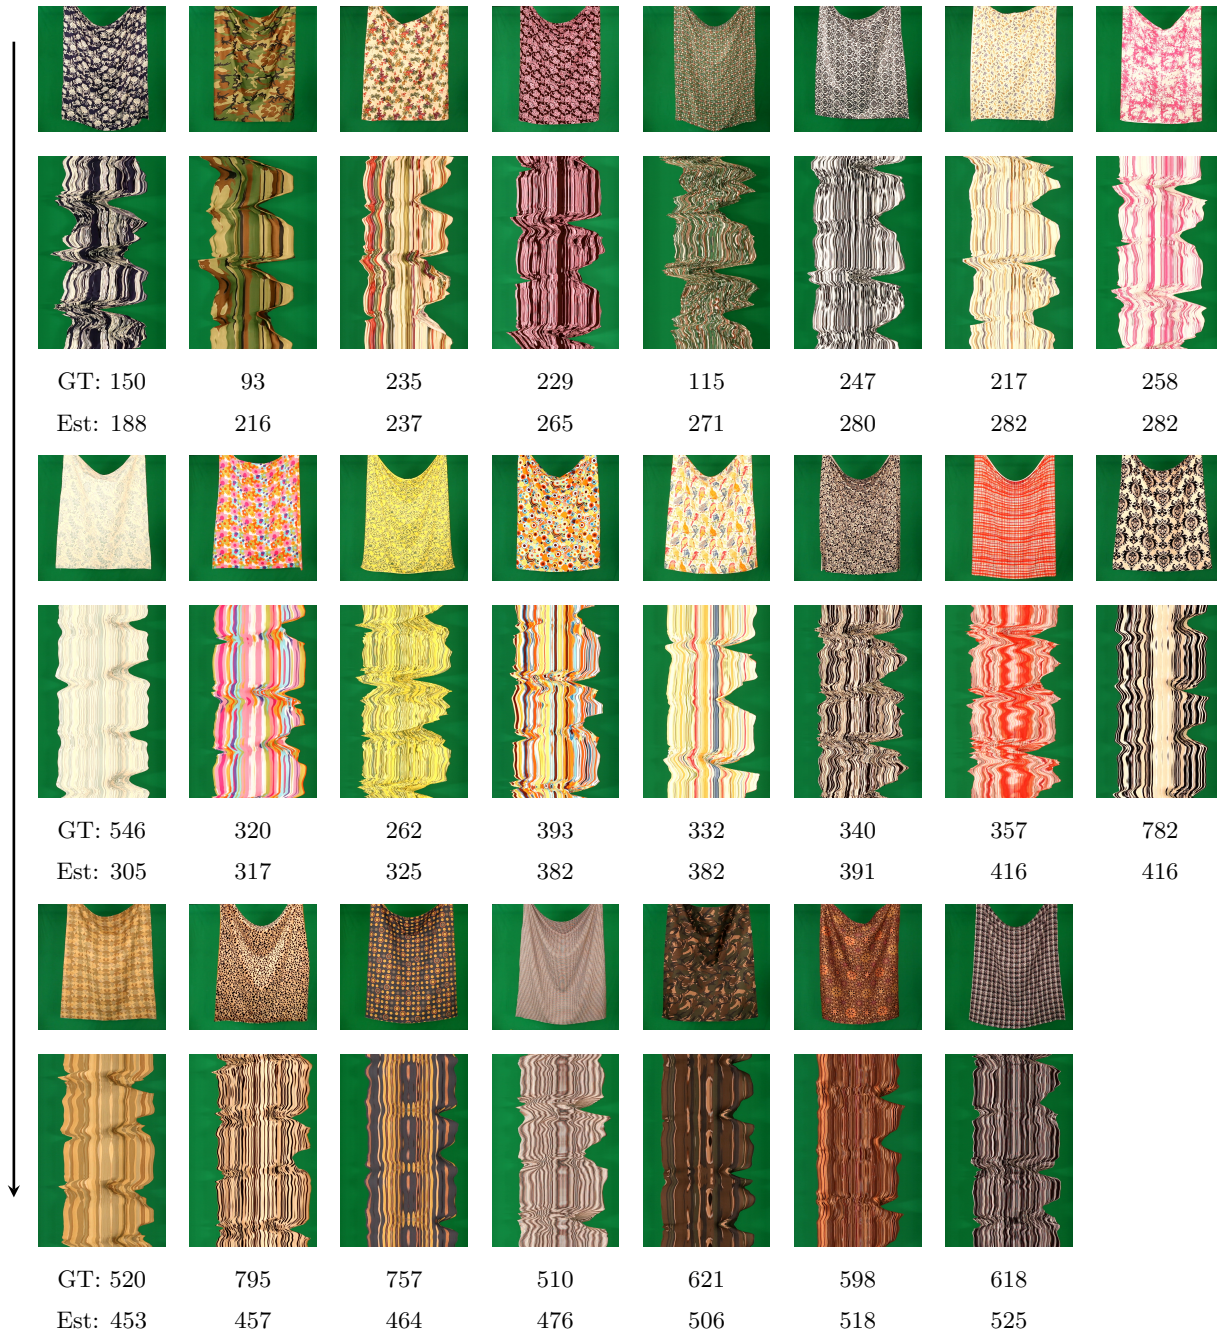


Figure 6.5. Fabrics in our collected database ranked (in row-major order) according to stiffness predicted by our model. The top image of each row shows the fabric samples hanging from a rod. The bottom image shows a slice through time of the x-axis when the fabrics are blown by the same wind force. The ground truth and predicted measurements of stiffness are shown below each image pair.



Increasing Mass \longrightarrow

Figure 6.6. Fabrics in our collected database ranked (in row-major order) according to mass predicted by our model. The top image of each row shows the fabric samples hanging from a rod. The bottom image shows a slice through time of the x-axis when the fabrics are blown by the same wind force. The ground truth and predicted measurements of mass are shown below each image pair.

Conclusion

WE have developed an approach for estimating the material properties of a fabric from video data through the use of features that capture spatiotemporal statistics in a video’s motion field. We tested our method on RGB videos from a new dataset on dynamic fabric movement and ground truth material parameters that we constructed. Our method recovers estimates of the stiffness and mass of fabrics that are well correlated with those of the log of ground truth measurements and estimates made by human observers who were asked to make analogous judgments. Both our method and the humans were able to partially discount the intensity of applied forces when forming judgments about material properties. We believe our dataset and algorithmic framework represent the first attempt to passively estimate the material properties from observing the motions of deformable objects. More generally, our work suggests that many physical systems with complex mechanics may generate image data that encodes the underlying physical parameters in a way that is extractable by efficient discriminative methods.

Bibliography

- [1] David Baraff and Andrew Witkin. *Large Steps in Cloth Simulation*. 1998.
- [2] Kiran Bhat, Christopher Twigg, Jessica Hodgins, Pradeep Khosla, Zoran Popovi, and Steven Seitz. Estimating cloth simulation parameters from video. pages 37–51. Eurographics Association, 2003. ISBN 1-58113-659-5. URL <http://portal.acm.org/citation.cfm?id=846276.846282>.
- [3] David J. Heeger and James R. Bergen. Pyramid-based texture analysis/synthesis. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, SIGGRAPH '95, page 229238, New York, NY, USA, 1995. ACM. ISBN 0-89791-701-4. doi: 10.1145/218380.218446. URL <http://doi.acm.org/10.1145/218380.218446>.
- [4] Yun-xian Ho, Michael S. Landy, and Laurence T. Maloney. How direction of illumination affects visually perceived surface roughness. *Journal of Vision*, 6: 634648, 2006.
- [5] Nebojsa Jojic and Thomas S. Huang. Estimating cloth draping parameters from range data. In *In International Workshop on Synthetic-Natural Hybrid Coding and 3-D Imaging*, page 7376, 1997.
- [6] B. Julesz. Visual pattern discrimination. *IRE Transactions on Information Theory*, 8(2):84–92, 1962. ISSN 0096-1000. doi: 10.1109/TIT.1962.1057698.
- [7] Bela Julesz. Textons, the elements of texture perception, and their interactions. *Nature*, 290(5802):91–97, March 1981. ISSN $\{\text{footerJournalISSN}\}$. doi: 10.1038/290091a0. URL <http://www.nature.com/nature/journal/v290/n5802/abs/290091a0.html>.
- [8] S. Kawabata and Masako Niwa. Fabric performance in clothing and clothing manufacture. *Journal of the Textile Institute*, 80(1):19–50, January 1989. ISSN 0040-5000, 1754-2340. doi: 10.1080/00405008908659184. URL <http://www.tandfonline.com/doi/pdf/10.1080/00405008908659184#.Uh1mNmS2Mts>.

- [9] Frdric B. Leloup, Michael R. Pointer, Philip Dutr, and Peter Hanselaer. Geometry of illumination, luminance contrast, and gloss perception. *Journal of the Optical Society of America A*, 27(9):2046–2054, September 2010. doi: 10.1364/JOSAA.27.002046. URL <http://josaa.osa.org/abstract.cfm?URI=josaa-27-9-2046>.
- [10] Ce Liu. *Beyond pixels : exploring new representations and applications for motion analysis*. Thesis, Massachusetts Institute of Technology, 2009. URL <http://dspace.mit.edu/handle/1721.1/53293>. Thesis (Ph. D.)—Massachusetts Institute of Technology, Dept. of Electrical Engineering and Computer Science, 2009.
- [11] Ce Liu, Lavanya Sharan, Edward Adelson, and Ruth Rosenholtz. Exploring features in a bayesian framework for material recognition. 2010.
- [12] E. Miguel, D. Bradley, B. Thomaszewski, B. Bickel, W. Matusik, M. A. Otaduy, and S. Marschner. Data-driven estimation of cloth simulation models. *Computer Graphics Forum*, 31(2pt2):519528, 2012. ISSN 1467-8659. doi: 10.1111/j.1467-8659.2012.03031.x. URL <http://onlinelibrary.wiley.com/doi/10.1111/j.1467-8659.2012.03031.x/abstract>.
- [13] Randal C. Nelson and Ramprasad Polana. Qualitative recognition of motion using temporal texture. *CVGIP: Image Understanding*, 56(1):78–89, July 1992. ISSN 1049-9660. doi: 10.1016/1049-9660(92)90087-J. URL <http://www.sciencedirect.com/science/article/pii/104996609290087J>.
- [14] Thrasyvoulos N. Pappas. The rough side of texture: texture analysis through the lens of HVEI. pages 86510P–86510P, March 2013. doi: 10.1117/12.2012991. URL <http://dx.doi.org/10.1117/12.2012991>.
- [15] Javier Portilla and Eero P. Simoncelli. A parametric texture model based on joint statistics of complex wavelet coefficients. *INTERNATIONAL JOURNAL OF COMPUTER VISION*, 40(1):4971, 2000.
- [16] A. Rahman and M. Murshed. A robust optical flow estimation algorithm for temporal textures. In *International Conference on Information Technology: Coding and Computing, 2005. ITCC 2005*, volume 2, pages 72–76 Vol. 2, April . doi: 10.1109/ITCC.2005.31.
- [17] C. Schuldt, I. Laptev, and B. Caputo. Recognizing human actions: a local SVM approach. In *Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004*, volume 3, pages 32 – 36 Vol.3, August 2004. doi: 10.1109/ICPR.2004.1334462.
- [18] Lavanya Sharan, Yuanzhen Li, Isamu Motoyoshi, Shin’ya Nishida, and Edward H Adelson. Image statistics for surface reflectance perception. *Journal of the Optical*

- Society of America. A, Optics, image science, and vision*, 25(4):846–865, April 2008. ISSN 1084-7529. PMID: 18382484.
- [19] Dimitri Soteropoulos, Konstantine Fetfatsidis, James A. Sherwood, and Joanna Langworthy. Digital method of analyzing the bending stiffness of NonCrimp fabrics. *AIP Conference Proceedings*, 1353(1):913–917, May 2011. ISSN 0094243X. doi: doi:10.1063/1.3589632. URL http://proceedings.aip.org/resource/2/apcpcs/1353/1/913_1.
- [20] Huamin Wang, James F. O’Brien, and Ravi Ramamoorthi. Data-driven elastic models for cloth: modeling and measurement. In *ACM SIGGRAPH 2011 papers*, SIGGRAPH ’11, page 71:171:12, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0943-1. doi: 10.1145/1964921.1964966. URL <http://doi.acm.org/10.1145/1964921.1964966>.
- [21] Song Chun Zhu, Yingnian Wu, and David Mumford. Filters, random fields and maximum entropy (FRAME) towards a unified theory for texture modeling. *INTERNATIONAL JOURNAL OF COMPUTER VISION*, 27(2):120, 1998.