Estimating the Material Properties of Fabric from Video

Katherine L. Bouman Bei Xiao Peter Battaglia William T. Freeman

Massachusetts Institute of Technology

{klbouman, beixiao, pbatt, billf}@mit.edu

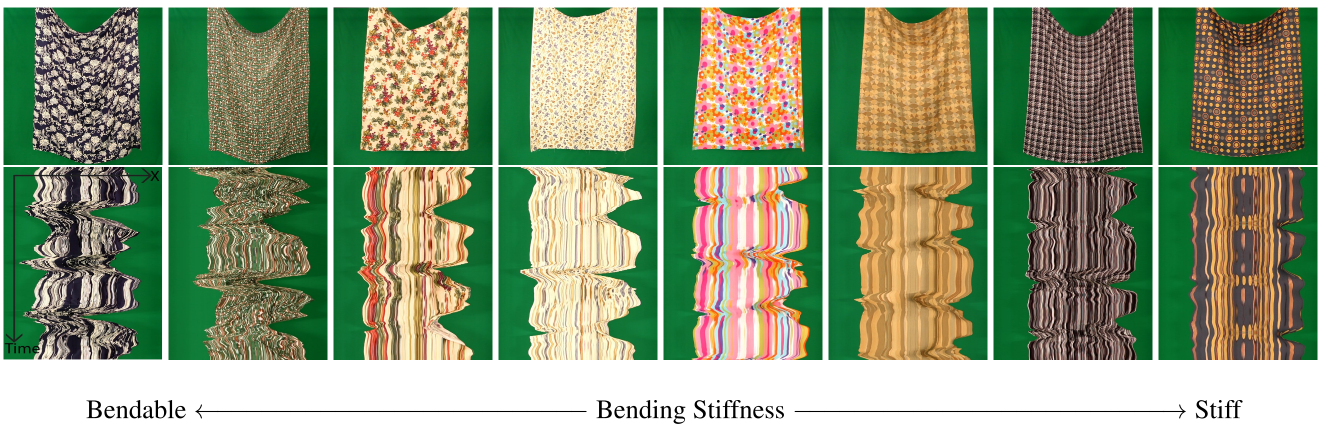

Figure 1. A sample of the fabrics in our collected database ranked according to stiffness predicted by our model. The top panel shows physical fabric samples hanging from a rod. The bottom panel shows a horizontal space $\times$ time slice of a video when the fabrics are blown by the same wind intensity. Bendable fabrics generally contain more high frequency motion than stiff fabrics.

Abstract

Passively estimating the intrinsic material properties of deformable objects moving in a natural environment is essential for scene understanding. We present a framework to automatically analyze videos of fabrics moving under various unknown wind forces, and recover two key material properties of the fabric: stiffness and area weight. 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, publicly available database of fabric videos with corresponding measured ground truth material properties. We show that our predictions are well correlated with ground truth measurements of stiffness and density for the fabrics. Our contributions include: (a) a database that can be used for training and testing algorithms for passively predicting fabric properties from video, (b) an algorithm for predicting the material properties of fabric from a video, and (c) a perceptual study of humans' ability to estimate the material properties of fabric from videos and images.

1. Introduction

Automatic scene understanding is a fundamental goal of computer vision. Although the computer vision community has made great strides in the last couple of decades towards achieving this goal, with work in object detection, 3D reconstruction, etc., there has been very little work on understanding the intrinsic 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 these properties from a video is a difficult problem that is essential for automatic scene understanding.

Knowing the material properties of objects in a scene allows one to have a better understanding of how objects will interact with their environment. Additionally, it can be very useful for many applications such as robotics, online shopping, material classification, material editing, and predicting objects' behavior under different applied forces. For instance, imagine a system that is able to automatically segment a video of a complex natural scene into different materials and estimate their intrinsic properties. Such a system would yield powerful meta-data that could be easily integrated into many video search applications.

In this paper, we focus on developing an algorithm to passively estimate the material properties of fabric from a video of the fabric moving due to unknown wind forces. Two considerations motivate starting with fabric. 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 [7]. Second, fabric is intrinsically a 2D material, making most of its motion easily visible from video.

The motion of a fabric is determined by its density, its resistance to bending, stretching, and shearing, external forces, aerodynamic effects, friction, and collisions [1]. In this work we restricted our attention to recovering two of the most distinctive properties of fabric in natural scenes - the bending stiffness and area weight (weight/area). We aimed to develop spatiotemporal visual features that capture these intrinsic material properties of fabric. To the best of our knowledge, our 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.

The remainder of this paper is structured as follows. In Section 2 we provide a background of previous applicable work. Section 3 describes the database we have collected for training and testing of our algorithm. Section 4 describes a perceptual study testing how well humans are able to estimate the material properties of fabric from video and image data. Section 5 presents our algorithm for predicting the material properties of fabric. Section 6 contains results of our algorithm and a discussion of the results.

2. Background

Previous work has focused on understanding static properties of materials, such as surface reflectance [14], material category [10], roughness [4], and surface gloss [2]. In contrast, we address the problem of passively estimating material properties of deformable objects in a natural scene that are visually evident through dynamic motions.

The material properties of fabric can be measured using expensive and time-intensive systems. These systems precisely measure a fabric's response to many different, controlled forces. The most well known setup used to measure these parameters is the Kawabata evaluation system [7]. Since the development of the Kawabata system, other systems have been developed to directly measure the properties of fabric [11, 17]. Although these systems produce accurate measurements of a fabric's material properties, they are undesirable for applications in which we are not able to directly manipulate a physical specimen of the fabric.

Jojic and Huang attempted to estimate a fabric's material parameters from 3D data of a static scene containing the fabric draped over an object [5]. However, because the fabric was static, the system was not able to estimate properties evident from dynamic motion. Additionally, the system needed very accurate 3D data in order to perform the inference. Bhat et al. presented a method to estimate the material properties of fabric from video data [1]. However, this system has several limitations as well; the system 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 the problem we focus on. Instead, we wish to estimate material properties in a more natural setting, when the fabric is exposed to unknown forces.

3. Database

In order to study this problem, we collected a database containing videos of moving fabrics along with their associated material properties. This database has been made publicly available online[1]. Thirty fabrics were collected for the database. The fabrics span a variety of stiffness and densities. Example categories include cotton, velvet, spandex, felt, silk, upholstery, wool, denim, and vinyl.

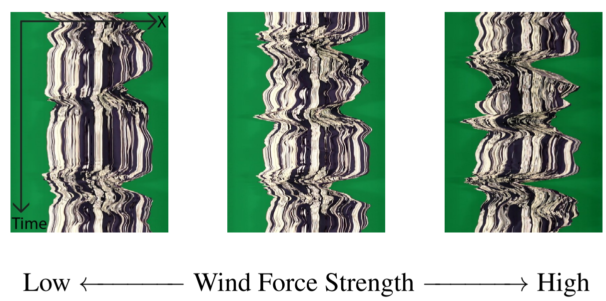

Figure 2. An example of a horizontal space × time slice of the same fabric exposed to the three different strengths of wind from an oscillating fan. Note that the motion of the fabric appears very different under the different wind strengths.

Ground Truth Measurements In order to obtain ground truth material property measurements for the fabrics, we sent specimens of each fabric to the Lowell Advanced Composite Materials and Textile Research Laboratory[2] to have their stiffness (lbf-in^2), area weight (oz/yd^2), and density (lb/in^3) measured [16]. Since the material properties of fabric often varies depending on the direction of measurement, for many of the fabrics we made measurements of the properties in two orthogonal directions.

In this work, we have focused on recovering two of these properties from video - the stiffness and area weight of fabrics. For convenience, in the remainder of this paper we refer to area weight as density.

[1]http://people.csail.mit.edu/klbouman
[2]http://m-5.uml.edu/acmtrl/index.htm

Videos Videos (859×851 pixel resolution) were recorded for all fabrics. Fabrics were hung from a bar and exposed to three different strengths of wind from an oscillating fan positioned to the right of the fabrics. The two-minute videos capture the fabrics moving in response to the wind force. Figure 2 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 strengths.

RGB-D Kinect videos (640×480 pixel resolution) 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 our work thus far, however this information could be used in the future to obtain motion along the depth dimension.

All fabrics were cut to approximately 107×135 cm, and steamed to remove wrinkles. Cutting the fabrics to the same size removes any uncertainties due to scale that would confuse human observers or an algorithm. For instance, a life-size window curtain may move in a qualitatively different way than a curtain from a dollhouse, even when cut from the same piece of fabric.

4. Human Material Perception

In order to design our own algorithm we first looked to humans for inspiration on what features may be important. While visual cues from a static image can sometimes reveal a lot about the materials in a scene, they can often be misleading. In these cases, a video may help to disambiguate the material properties.

To verify that humans use motion cues to passively estimate material properties, we designed a psychophysical experiment to understand material perception from a purely visual perspective. The experiment was designed to measure how well subjects are able to estimate the relative stiffness and density of fabrics when observing video or image stimuli. These experiments were conducted using Amazon's Mechanical Turk. Results of this study have been made publicly available online.

4.1. Experimental Setup

Video Stimuli Stimuli included the videos of 30 common fabrics exposed to 3 different strengths of wind from our database (Section 3). A paired comparison method was used to measure perceived differences in the stiffness and density between the fabrics in two videos [8]. Specifically, a subject was shown two videos of 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 answering on a 7-point scale provided underneath the videos (Figure 3). This *pairwise score*, which takes a value in $\{-3, -2, -1, 0, 1, 2, 3\}$, indicates which fabric the subject believed was stiffer, and the degree of stiffness difference between two fabrics. Similarly, in a second experiment, a subject was asked to report a pairwise score indicating the relative weight of the fabric. Since fabrics in the videos were cut to approximately the same size, the task of predicting a fabric's area weight reduces to predicting its weight in this experiment.

A total of 100 workers from Mechanical Turk ($> 95\%$ approval rate in Amazon's system) completed each experiment by answering 100 questions. To prevent biases from seeing different wind combinations across video pairs, a particular subject always saw the same combination of wind strengths between the two videos.

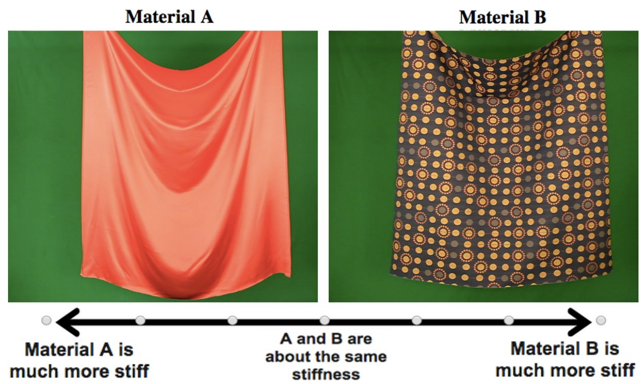

Figure 3. Experimental setup of pairwise comparisons of material properties (stiffness or density) from image stimuli. Subjects were asked to compare material properties of the two fabrics on a 7 point scale. A similar setup was also used to compare the stiffness and density of fabrics given video stimuli.

Image Stimuli A similar experimental setup was used to test the perception of density and stiffness of the same 30 draped fabrics from a single still image. A total of 25 workers from Mechanical Turk ($> 95\%$ approval rate in Amazon's system) completed each experiment. Each subject answered 100 questions.

Participants In order to maximize high quality responses, subjects were required to watch each pair of videos for 15 seconds and each pair of images for 2 seconds before responding. Additionally, subjects were tested periodically throughout the experiment by answering questions that they had been given the answer to previously. Subjects who did not respond to over 80% of these questions correctly were removed from our analysis.

4.2. Data Analysis and Discussion

Given N pairwise scores, a single perceptual score for each of the K fabrics was found by solving a linear regression problem. For each question $n \in \{1, ..., N\}$ a stimuli pair containing fabrics $i_n, j_n \in \{1, ..., K\}$ was observed and assigned a pairwise score $b(n) \in \{-3, 2, 1, 0, 1, 2, 3\}$. To obtain a single score for each fabric we solve $Ax = b$ for x, where $A(n, i_n) = -1$ and $A(n, j_n) = +1$ for all n.

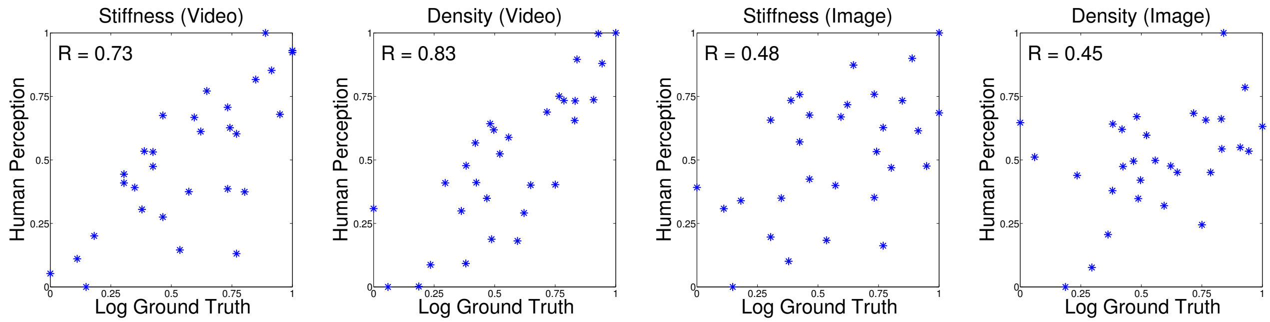

(a) (b) (c) (d)

Figure 4. Comparisons of ground truth material properties with human predictions when subjects observed video (a,b) and image (c,d) stimuli. Each star in the plots represents a single fabric. The Pearson product-moment correlation coefficient (R-value) is shown for each comparison. These plots suggest that human observers use motion cues in videos to estimate material properties. Results are shown scaled to values in the range of 0 to 1.

In accordance with Weber's Law, we found that human responses were well correlated with the log of ground truth stiffness and density measurements when they were asked to make judgments from videos of the fabric. However, the responses were much less correlated with ground truth measurements when the subjects were asked to make judgments only from still images of the fabric. Figure 4 compares the log of ground truth stiffness and density of the fabrics with the perceptual score of the same material property for these experiments. These plots suggest that human observers use motion cues in videos to estimate material properties.

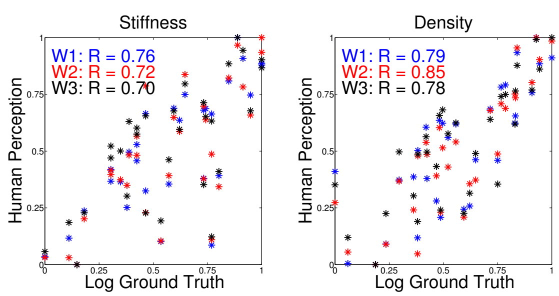

(a) (b)

Figure 5. Comparison of ground truth stiffness (a) and density (b) versus perceptual scores computed from responses where subjects observed fabrics exposed to the same wind strength. Results are colored according to the wind strength applied (indicated by W1, W2, and W3). This suggests that humans were somewhat invariant to the strength of wind when making predictions about material properties of fabric.

Next, we evaluated the effect of wind force on subjects' responses in estimating material properties of fabric. To do so, for each pair of fabrics, we measured how the *difference* in wind strength affected the fabrics' pairwise score. Table 1 shows the average change in pairwise score for every increase in wind strength difference. We find that while a wind's strength has a small effect on human perception of stiffness and density, relative judgements of the material properties are largely unchanged under different force

Stiffness	Density
-4.4% $\pm$ 5.8%	-4.7% $\pm$ 5.3%

Table 1. The average sensitivity of humans to the strength of a wind force in estimating material properties of fabric. The average percentage change (and standard deviation) of a pairwise score for every wind strength increase applied to the fabric in the second stimuli (Material B). This value indicates that subjects on average judged fabric moving with an increased force as 4.4% less stiff and 4.7% less heavy than they would have with a weaker force.

environments. Figure 5 illustrates how subjects' responses correlated with ground truth material properties in varying wind conditions for pairs of fabric moving under the *same* wind strength.

5. Approach

A 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. 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 work, we use statistics characterizing temporal textures in order to predict the material properties of fabric.

A flow diagram of our algorithm can be seen in Figure 6. The input to our system is a video of a previously unseen fabric moving (Figure 6a) along with a mask of which pixels contain the fabric in every frame; the output is a value indicating its stiffness or density.

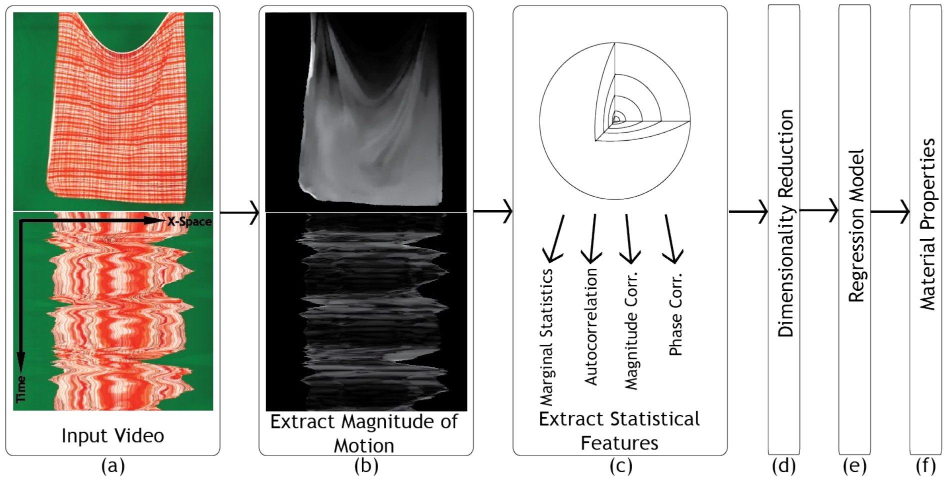

Figure 6. Illustration of our framework for estimation of material properties of fabric from video. The input to our system is a video containing fabric (a) along with a mask of what pixels contain the material. The masked magnitude of motion is extracted from the video of moving fabric via optical flow (b). Features are computed from the masked magnitude of the motion (c). These features are computed on a decomposition of the motion into sub-bands associated with concentric spheres in the frequency domain. 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 the stiffness or density of the fabric in the input video (f).

5.1. Motion Estimation

The intensity values of a video contain information about both the appearance and motion in a scene. The printed pattern on a fabric is less useful for the purpose of material property prediction since the pattern does not, in general, affect the material properties of the underlying fabric. Therefore, since we would like to focus on characterizing the fabric's motion, we separate the appearance of the printed pattern from the motion field in a video by computing the magnitude of the optical flow [9]. Any region in the video not containing the fabric is masked out by assigning it a flow magnitude of zero. Figure 6b shows the masked magnitude of flow for a sample video. Note that different parts of the fabric move at different speeds, even at a single instant in time.

5.2. Statistical Features

Once we have extracted the motion's magnitude from a video, our goal is to extract a set of features from the motion field that are descriptive of the material properties. Motivated by humans' ability to passively estimate the relative material properties of fabric, we would like to find a set of features that have a monotonic relationship between their computed values and the perceived similarity of the motion fields.

In designing our feature set, we draw inspiration from Portilla and Simoncelli's constraints that were developed for synthesizing perceptually indistinguishable 2D visual textures [12]. Portilla and Simoncelli developed a compact, parametric statistical model that could then be used for texture analysis. We extend Portilla and Simoncelli's work to 3D temporal textures for the application of inferring material properties.

Pyramid Decomposition 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 [12, 15]. Each sub-band contains a local estimate of the magnitude and phase of the 3D signal around a pixel. We have chosen to decompose the magnitude of our motion field into $N_{sc} = 4$ scales and $N_{or} = 1$ orientation. Figure 6c shows how the frequency domain is split up for our decomposition. Features are computed from the sub-bands of the multi-resolution complex pyramid decomposition.

Decomposing the motion field in this way is desirable for this application because different material properties may be more pronounced in the motion from different spatiotemporal scales. For instance, a fabric's density may have a larger effect on the low frequency motion, whereas a fabric's bending stiffness may have a larger effect on the high frequency motion.

The following sections describe features computed from the coefficients of the decomposed motion field in order to characterize the motion of a fabric. Implementation details for the special case of these features in 2D can be found in [12].

5.2.1 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, 12, 18] and action recognition [13] algorithms have either used marginal statistics or histograms directly to characterize marginal distributions. We measure the mean, skew, kurtosis and range (minimum and maximum) of the motion magnitude. Additionally, the mean, skew, and kurtosis for each of the $N_{sc} = 4$ lowpass videos are computed from the complex 3D pyramid. The marginal statistics of the lowpass videos characterize the distribution of motion magnitudes at different spatiotemporal scales.

5.2.2 Autocorrelation

Julsez'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]. 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 feature.

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.

5.2.3 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 [12]. 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 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.

5.2.4 Phase Correlation

Local phase estimates of a signal indicate its gradient in a local region [12]. In order to capture gradual changes in the motion field, we compute the correlation across the local phases in the neighboring sub-bands of the video's pyramid decomposition.

5.3. 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 density measurements. Motivated by Weber's Law, we choose to work in the log domain since humans tend to be sensitive to the logarithm of material properties and the features we have chosen to use were initially developed for perceptual indistinguishability.

We first standardize each feature by subtracting the mean and dividing by the the standard deviation. We would like each feature-type (e.g., marginal statistics, autocorrelation, etc.) to contribute the same amount of variance to the feature vector. Thus, we force the variance of each feature to be proportional to the number of features in its feature-type. We do this by dividing each feature by the square root of the number of elements in its feature-type. Dimensionality of the standardized feature vectors is then reduced using PCA. Feature vectors are projected onto the top M eigenvectors, E_m, that preserve 95% of the variance in the data.

A simple linear regression model is used to map the resulting features to the ground truth material properties. We chose to use linear regression rather than a more complex regression method to more directly reveal the power in the selected features. 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 $W \odot Y = W \odot X\beta$, 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.

5.4. Implementation Details

Twenty-three of the 30 fabrics in our database were selected for training and testing of our model. Fabrics were removed that either lacked texture or caused specularities in the videos since they produced inaccurate optical flow estimates of the motion.

Videos were first cropped to 832×832 pixels. Then, for each video we extracted two non-overlaping video segments, each 512 frames long. A single feature vector was computed for each segment. The linear regression model described in Section 5.3 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 contained multiple ground truth measurements, we mapped 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 segment. More specifically, when making a prediction using a feature vector associated with a fabric, all feature vectors extracted from video segments correspond-

Error	Stiffness	Density
Motion	17.2%	13.8%
Intensity	34.7%	28.9%

Table 2. Percentage error calculated for stiffness and density estimates when features were computed from the motion's magnitude versus grayscale intensity values. Percentage error is calculated by taking the average percentage difference between a predicted measurement for each video segment and all ground truth log measurements for a specific fabric.

Stiffness	Density
-5.8% $\pm$ 4.0%	-4.6% $\pm$ 5.2%

Table 3. The sensitivity of our model to the wind strength in estimating the material properties of fabric. The average percentage change (and standard deviation) of a pairwise score for every wind strength increase applied to a given fabric. The sensitivity of our model to the wind force is comparable to the sensitivity of human observers.

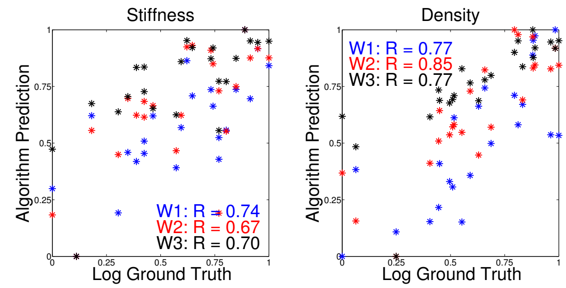

(a) (b)

Figure 8. Comparison of ground truth stiffness (a) and density (b) versus our model's predictions computed from differences in the value predicted for fabrics exposed to the same wind strength. Results are colored according to the wind strength applied (indicated by W1, W2, and W3).

ing to the same fabric were removed from the training set.

6. Results and Discussion

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 this section, we demonstrate the power of the features introduced in Section 5.2 for predicting the stiffness and density of fabrics from video.

We compare predicted measurements of stiffness and density from our algorithm to the ground truth measurements (Section 3) and perceptual estimates (Section 4) in Figure 7. This figure suggests that our estimates of the material properties of the fabric in a video are well correlated with the log of ground truth material property values. Thus, our model is able to find a general trend of increasing stiffness and density in the fabric videos.

Percentage error for stiffness and mass of our results can be seen in Table 2. To evaluate the usefulness of extracting the motion magnitude from the videos, as a baseline we have also calculated the percentage error when features were computed from the grayscale intensity values of the video rather than the the motion's magnitude. The error is significantly larger when features are computed from the grayscale intensity values. This supports our claim that it is necessary to decompose the video into printed texture and motion in order to estimate material properties using our proposed features.

To evaluate our model's sensitivity to wind strength in predicting a fabric's material properties, we computed the average change in pairwise score for every increase in wind strength difference as described in Section 4.2. Results (Table 3) show that our model's sensitivity to the wind force is comparable to that of human sensitivity (Table 1) in estimating the stiffness and density of fabric. For completeness, Figure 8 shows how relative predictions made by our model correlated with ground truth material properties when the videos contained fabrics moving under the *same* wind strength.

Sensitivity Analysis To evaluate the importance of each of our feature-types (eg. marginal statistics, autocorrelation, etc.) in the estimation of material properties, we have computed the total sensitivity due to each feature-type. The total sensitivity of the prediction due to the set of features F in a single feature-type is computed as

$$S(F) \propto \sum_{f \in F} \sqrt{\sum_{m=1}^{M} (\beta_m E_m^f)^2} \tag{1}$$

where E_m^f is the fth feature of the mth eigenvector and β_m are the regression weights from our model. A bar graph of the normalized sensitivities can be found in Figure 9. These sensitivities indicate that the autocorrelation is the most important feature for prediction of both the stiffness and density of fabric from video.

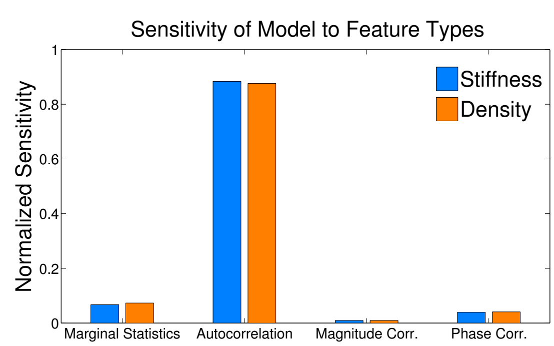

Figure 9. The average normalized sensitivity of each feature-type in our proposed model for the prediction of stiffness and density. Features related to the autocorrelation have the largest effect on the estimation of stiffness and density for videos from our database.

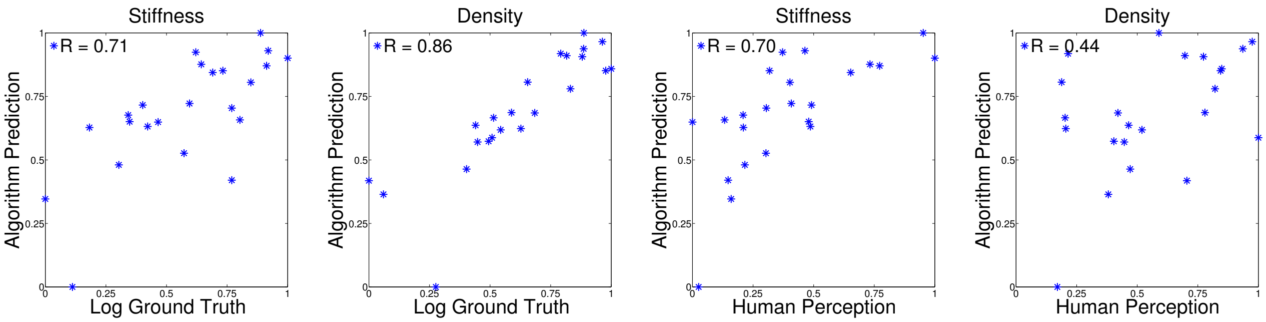

(a) (b) (c) (d)

Figure 7. Comparisons of model predictions for material properties against (a) ground truth stiffness, (b) ground truth density, (c) perceptual stiffness scores, and (d) perceptual density scores. Each star in the plots represents a single fabric. The Pearson product-moment correlation coefficient (R-value) is shown for each comparison. Results are shown scaled to values in the range of 0 to 1.

7. Conclusion

We have developed an approach for estimating the material properties of fabric from video 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, publicly available dataset on dynamic fabric movement and ground truth material parameters that we constructed. Our method recovers estimates of the stiffness and density of fabrics that are well correlated with the log of ground truth measurements. Both our method and humans were able to partially discount the intensity of applied forces when forming judgments about material properties. We believe our dataset and algorithmic framework is the first attempt to passively estimate the material properties of deformable objects moving due to unknown forces from video. More generally, our work suggests that many physical systems with complex mechanics may generate image data that encodes their underlying intrinsic material properties in a way that is extractable by efficient discriminative methods.

Acknowledgments We would like to thank Lowell ACMTL, particularly Patrick Drane, for their help in collecting data for this project. We would also like to thank Adrian Dalca for all of his helpful discussions and feedback.This work was partially supported by NSF CGV-1111415 and NSF CGV-1212928. Katherine Bouman was partially supported by an NSF Graduate Fellowship. Bei Xiao was supported by an MIT Intelligent Initiative Postdoctoral Fellowship. Peter Battaglia was supported by (IARPA) - D10PC20023.

References

[1] K. Bhat, C. Twigg, J. Hodgins, P. Khosla, Z. Popovi, and S. Seitz. Estimating cloth simulation parameters from video. Eurographics Association, 2003. 2

[2] R. W. Fleming, R. O. Dror, and E. H. Adelson. Real-world illumination and the perception of surface reflectance properties. *Journal of Vision*, 2003. 2

[3] D. J. Heeger and J. R. Bergen. Pyramid-based texture analysis/synthesis. SIGGRAPH '95, 1995. 6

[4] Y.-x. Ho, M. S. Landy, and L. T. Maloney. How direction of illumination affects visually perceived surface roughness. *Journal of Vision*, 2006. 2

[5] N. Jojic and T. S. Huang. Estimating cloth draping parameters from range data. In *In International Workshop on Synthetic-Natural Hybrid Coding and 3-D Imaging*, 1997. 2

[6] B. Julesz. Visual pattern discrimination. *IRE Transactions on Information Theory*, 1962. 6

[7] S. Kawabata and M. Niwa. Fabric performance in clothing and clothing manufacture. *Journal of the Textile Institute*, 1989. 2

[8] F. B. Leloup, M. R. Pointer, P. Dutr, and P. Hanselaer. Geometry of illumination, luminance contrast, and gloss perception. *Journal of the Optical Society of America*, 2010. 3

[9] C. Liu. *Beyond pixels : exploring new representations and applications for motion analysis*. Thesis, Massachusetts Institute of Technology, 2009. 5

[10] C. Liu, L. Sharan, E. Adelson, and R. Rosenholtz. Exploring features in a bayesian framework for material recognition. 2010. 2

[11] 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 (Proc. of Eurographics)*, 2012. 2

[12] J. Portilla and E. P. Simoncelli. A parametric texture model based on joint statistics of complex wavelet coefficients. *IJCV*, 2000. 5, 6

[13] C. Schuldt, I. Laptev, and B. Caputo. Recognizing human actions: a local SVM approach. In *ICPR 2004*, volume 3, Aug. 2004. 6

[14] L. Sharan, Y. Li, I. Motoyoshi, S. Nishida, and E. H. Adelson. Image statistics for surface reflectance perception. *Journal of the Optical Society of America. A, Optics, image science, and vision*, Apr. 2008. 2

[15] E. P. Simoncelli and W. T. Freeman. The steerable pyramid: A flexible architecture for multi-scale derivative computation. In *IEEE ICIP*, 1995. 5

[16] D. Soteropoulos, K. Fetfatsidis, J. A. Sherwood, and J. Langworthy. Digital method of analyzing the bending stiffness of NonCrimp fabrics. *AIP Conference Proceedings*, 2011. 2

[17] H. Wang, J. F. O'Brien, and R. Ramamoorthi. Data-driven elastic models for cloth: modeling and measurement. SIGGRAPH, 2011. 2

[18] S. C. Zhu, Y. Wu, and D. Mumford. Filters, random fields and maximum entropy (FRAME) towards a unified theory for texture modeling. *IJCV*, 1998. 6