

This is a repository copy of *Woven fabric model creation from a single image*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/153582/>

Version: Accepted Version

Article:

Guarnera, Giuseppe Claudio, Hall, Peter, Chesnais, Alain et al. (1 more author) (2017) Woven fabric model creation from a single image. ACM Transactions on Graphics. A165. ISSN 1557-7368

<https://doi.org/10.1145/3132187>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Woven Fabric Model Creation from a Single Image

GIUSEPPE CLAUDIO GUARNERA, NTNU - Norwegian University of Science and Technology

PETER HALL, University of Bath

ALAIN CHESNAIS, KiSP Inc.

MASHHUDA GLENCROSS, SwitchThat Technologies Ltd

We present a fast, novel image-based technique, for reverse engineering woven fabrics at a yarn level. These models can be used in a wide range of interior design and visual special effects applications. In order to recover our pseudo-BTF, we estimate the 3D structure and a set of yarn parameters (e.g. yarn width, yarn crossovers) from spatial and frequency domain cues. Drawing inspiration from previous work [Zhao et al. 2012], we solve for the woven fabric pattern, and from this build data set. In contrast however, we use a combination of image space analysis, frequency domain analysis and in challenging cases match image statistics with those from previously captured known patterns. Our method determines, from a single digital image, captured with a DSLR camera under controlled uniform lighting, the woven cloth structure, depth and albedo, thus removing the need for separately measured depth data. The focus of this work is on the rapid acquisition of woven cloth structure and therefore we use standard approaches to render the results.

Our pipeline first estimates the weave pattern, yarn characteristics and noise statistics using a novel combination of low level image processing and Fourier Analysis. Next, we estimate a 3D structure for the fabric sample using a first order Markov chain and our estimated noise model as input, also deriving a depth map and an albedo. Our volumetric textile model includes information about the 3D path of the center of the yarns, their variable width and hence the volume occupied by the yarns, and colors.

We demonstrate the efficacy of our approach through comparison images of test scenes rendered using: (a) the original photograph, (b) the segmented image, (c) the estimated weave pattern and (d) the rendered result.

CCS Concepts: • **Computing methodologies** → *Reflectance modeling; Shape modeling;*

Additional Key Words and Phrases: Appearance modeling, reverse engineering, textiles, weave pattern, depth map, Pseudo BTF

Authors' addresses: G.C. Guarnera, Dept of Computer Science, Norwegian University of Science and Technology (NTNU), Teknologivgen 22, Gjøvik, 2815, Norway; email: giuseppe.guarnera@ntnu.no; P. Hall, Dept of Computer Science, University of Bath, Bath, BA2 7AY, United Kingdom; email: p.m.hall@bath.ac.uk; A. Chesnais, KiSP Inc., 151 Placer Court, Toronto, ON, M2H 3H9 Canada; email: alain_chesnais@siggraph.org; M. Glencross, SwitchThat Technologies Ltd, Buxton Court 3, West Way, Botley, Oxford, OX2 0JB, United Kingdom; email: mashhuda@switch-that.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 0000 Association for Computing Machinery.

0730-0301/0000/0-ART00 \$15.00

https://doi.org/0000001.0000001_2




Fig. 1. In this image, we show (a) a photo of the fabric swatch, (b) a rendering of a sofa using our recovered model, (c) our weave pattern segmentation, and (d) a detailed view of the sofa with the recovered model.

ACM Reference format:

Giuseppe Claudio Guarnera, Peter Hall, Alain Chesnais, and Mashhuda Glencross. 0000. Woven Fabric Model Creation from a Single Image. *ACM Trans. Graph.* 0, 0, Article 00 (0000), 14 pages.

https://doi.org/0000001.0000001_2

1 INTRODUCTION

Acquiring Bidirectional Texture Function (BTF) [Dana et al. 1999; Filip and Haindl 2009] models is particularly useful in a wide range of applications including interior design, games and visual special effects (VFX). Measuring the full 6 dimensional BTF is a lengthy and storage intensive process and estimation from a single picture is an under-constrained problem. Currently, most real time applications use artist generated bump maps and normal maps as these are performance and memory efficient. There remains an ongoing need for fast acquisition of these models from images. In the case of woven cloth, reverse engineering the 3D structure and thread colors from a photograph of a cloth sample and using this to automatically generate bump and depth maps (for rendering) reduces the burden on artists to

create these, and also allows for creation of detail at a range of scales for level of detail.

Our work is motivated by the material acquisition needs of the contract furniture business in which woven cloth model acquisition must be rapid, the volume of cloth swatches that need to be routinely processed is huge, and results must be renderable in high quality on a mobile device. Since our method is specifically designed to be used in this industry, where high quality renderings are necessary to show clients realistic representations of their newly designed interior spaces, this imposes high quality demands together with competing robustness and performance constraints and necessitating a fast approach. Our work is equally applicable to reverse engineering woven cloth models for other applications.

The desire to automatically recover a textile model for woven fabrics presents a number of challenges due to cloth appearance being dependent on both the structure and reflectance properties of the yarns used in the fabric. Self shadowing effects, weave pattern, variation in yarn widths, texture and loose fibres etc. contribute significantly to visual perception of variation and resulting visual richness of fabrics. Although woven cloth seems fairly regular in terms of structure, the yarns used can significantly impact upon the appearance of regularity. Consequently, to realistically recover a woven fabric model from an image we need to estimate the appearance of the yarns as well as the weave pattern itself.




Fig. 2. Plain (a) and more complex weave patterns (b). Warps are the vertical yarns (colored gray), whereas wefts traverse horizontally (in magenta). The pattern is a binary matrix describing how warps and wefts are interlaced by a loom: 0 indicates the warp overlaps the weft at the corresponding yarn crossing, 1 means it lies below. The patterns are shown in the matrix.

In recent years, within the computer graphics and vision research communities, there has been much interest in acquisition methods. A number of these works specifically focus on weave pattern recognition and synthesis of volumetric fabric appearance [Schwartz et al. 2013; Zhao et al. 2012]. Typically, these methods either rely on costly scanning equipment, are computationally intensive, or are data intensive [Schwartz et al. 2013]. Synthesized databases of various weave patterns are also expanding with the inclusion of new samples and typically require more storage space. Generation of new weave patterns from models within an existing database can reduce storage space and manual effort [Zhao et al. 2012]. The merits and limitations of these approaches will be described further in the related work in Section 3. Ultimately, however, due to the limitations of current acquisition methods, creating woven fabric models is still largely a labor-intensive process involving effort from artists to construct,

from reference images, normal maps and bump maps for all fabric models contained within a scene.

Woven textiles possess spatially repeating structure defined by the warp and weft (Figure 2), which control the design and strength of the fabric. We hypothesise that, in common with crystallographic analysis [Thompson 1978], this regularity can be exploited to infer the weave pattern through Fourier domain analysis. Transforming an image into the frequency domain enables it to be expressed in terms of amplitude and phase of the frequencies contained within the original photograph and so the Fourier Transformed image signature will show clear peaks which can be used to understand the yarn crossover points in woven textiles.

Building upon this foundation, the four novel contributions of this paper are:

- An algorithm to address inhomogenous fabrics (i.e. not all yarns with the same color);
- An algorithm to address specific cases such as homogeneous and dark fabrics, for which the Signal to Noise Ratio (SNR) does not allow a proper analysis of the woven structure;
- A noise model characterizing the noise from real sample textiles (i.e. variations in yarn width, misalignments, etc.);
- An algorithm to insert the observed noise parameters from real samples into the generated depth maps;

Prior to discussing related work, we present an overview of our image processing pipeline to provide context for the discussion of the relative merits and limitations of previous approaches and to clarify the novel aspects of our work.

2 PIPELINE OVERVIEW

Our software takes as input a photograph of a fabric sample, captured with a standard DSLR camera. In common with other approaches we assume that the entire weave pattern repeat is contained in the photograph of the swatch. Although some exceptions exist, this condition is comparatively easy to obtain for a broad range of woven fabrics.

Recall that the visual characteristics of woven textiles are highly influenced by their resulting structural properties: fabric production machinery and the spinning technology of yarn production further complicates this causing an irregular distribution of warp and weft of yarns, which can result in skewed, twisted or partially overlapping yarns attributable to the slightly variable yarn width. For these reasons, we observe that naively analyzing the Fourier spectrum of images of woven cloth is insufficient for a robust algorithm.

This necessitates our novel multi-step approach (outlined in Figure 3) consisting of the following serial steps:

Data Capture Acquire input from textile swatches.

Initial Image Processing For each swatch, axis align yarns. Construct *Fourier Signatures* to use as unique identifiers




Fig. 3. Flow Chart of the System.

and store these in a database. Automatically classify the input swatch into homogeneous or inhomogeneous.

Inhomogeneous Fabric Analysis Apply our novel weave pattern detection and noise estimation algorithms described in Section 4.3. Save output to our own database for matching irregular (homogeneous) samples. Automatically classify the input swatch into homogeneous or inhomogeneous samples.

Homogenous Fabric Analysis Match to samples in database to deal with fabrics that are constructed from a single yarn color.

Depth Map Creation Create the 3D structure from the weave pattern incorporating the noise model to represent the natural yarn variations; derive a depth map from the 3D representation. Create an approximate albedo map.

In the following section we discuss related work. There is a large body of literature on acquisition, rendering and synthesis of cloth. However, since our focus is primarily on techniques to recover woven cloth models from images, we constrain our survey of related works to this area alone.

3 RELATED WORK

The methods discussed in this section can be classified in terms of Fourier Transform (FT) based, spatial domain based and database indexing approaches. As can be inferred from Figure 3 above, our method spans all three of these categories.

3.1 Fourier Transform (FT) Based Approaches

In this category, we discuss three approaches in turn. The first was presented by Kinoshita *et al.*, the second by Ravandi and Toriumi, and the third by Ralló *et al.*

Kinoshita *et al.* developed an early method which used image processing to determine the weave pattern of fabric samples. The authors used a television camera, with a concave lens placed in front, for magnification to capture their image data. Incident lighting was via a tungsten lamp and the camera gray levels were pre-set so that only features of interest were measured. This method focused on the spatial distribution of peaks in the Fourier Spectra and related this to the weave pattern of the cloth. A key limitation of this approach was that it could only recover basic weave types (plain, twill, satin) [Kinoshita *et al.* 1989].

Ravandi and Toriumi demonstrated the use of image processing techniques for yarn density estimation. Their approach was based on the Angular Fourier Power Spectrum (APS) and autocorrelation function of fabric surfaces [Ravandi and Toriumi 1995]. To obtain these measurements, the values of the two-dimensional power spectra taken over defined sectors, centered at the origin, were computed. A key limitation of this work is that it can only be used to recover the weave pattern of plain-woven cotton fabrics.

Ralló *et al.* proposed a method that provides information about the structure of textiles consisting of more than 30 yarns in a conventional square $n \times n$ weave-repeat [Ralló *et al.* 2003]. The authors used the convolution theorem (also widely used in crystallography [Thompson 1978] to analyze the structure of different crystalline materials) to decompose the Fourier Spectra of the fabric into the convolution of the FS of the corresponding conventional weave-repeat base and FS of the pattern of repeat. They identify the minimal weave repeat by finding slanted lines along which the peaks seem to be predominantly aligned and connecting a subset of them into a parallelogram, with two brightest peaks on the corners of one side and one peak in the middle of the opposite side. This main parallelogram is then subdivided into smaller ones according to the distribution of minor peaks contained within it. The number of parallelograms in each direction gives the size of the minimal weave repeat. A key limitation of this approach is that if the woven cloth pattern itself is not basic then this algorithm can only recover the size of the pattern and not the pattern itself.

3.2 Spatial Domain Based Approaches

In this category, we refer to a number of approaches based on integral projections of the gray levels or gradient magnitude of the textile image to construct a regular grid aligned with the axis before detailing a recent approach that is closest to our work.

Due to natural irregularities of yarn distribution, integral projections are unreliable by themselves. However, this approach has been used with input either based on single [Huang *et al.* 2000; Kuo *et al.* 2004] or double-sided views [Xin *et al.* 2009] of fabric swatches. Xin *et al.* proposed an active grid model that adapts to yarn contours following local gradient information [Xin *et al.* 2009], but they did not model

the statistical distribution of yarn irregularities necessary for realism in the rendered output. However, these works are important, as we use this idea in our work purely as an initial bootstrap step to construct a first guess about yarn crossings and their spacing.

A recent work by Schröder *et al.* is the closest current work to the algorithm we propose [Schröder *et al.* 2015]. The authors capture local irregularities such as thinning and thickening of yarns also using a single image as input. They also model the woven textile structure together with yarn parameters. Their method uses an iterative process starting with an estimate of the weave repeat obtained through auto-correlation analysis. Next they apply optical flow to register a tiled image of the weave pattern repeat to a shear corrected input image. This process is applied at both coarse and fine levels of granularity to locate yarn deformations. Horizontal and vertical structures are detected using cylindrical filters with an *a priori* requirement of size. The algorithm uses belief propagation and expectation maximization, on a graphical model, in which each node has a four-connected neighborhood. The textile image is then segmented using information from previous steps, assuming that yarns can be represented as tubes and minimizing a function which accounts for yarn color, background color and edge position.

Our approach, and the work by [Schröder *et al.* 2015], both aim to fully reverse engineer models of woven cloth at the yarn level. There are some elements in common purely because both algorithms address the same problem of recovering woven cloth models from a single image. However, Schröder *et al.* adopt a spatial domain approach. In contrast to all the prior works mentioned, our novel approach spans all three categories in our literature classification. This results in some important advances over previous methods. First, our method is fast compared with that of Schröder *et al.*'s which (according to the authors) takes 30-45 minutes to process each sample [Schröder *et al.* 2015], whereas our method takes seconds to a few minutes. Unlike other methods, the creation of our database indexed through the Fourier Signature allows us to rapidly obtain all the required information for rendering by simply matching the feature vector derived from the Fourier spectra, thus bypassing all the possible sources of errors and speeding up the reverse engineering process. Unique to our method, we are also able to construct woven pattern depth maps in a generative fashion and match observations in our image captures to these generated patterns. This allows us to build up a rich variety of datasets without measuring every possible woven fabric pattern we may encounter. A summary of the differences between our approach and Schröder *et al.*'s is reported in Table 1.

3.3 Database Indexing Approaches

Zheng *et al.* proposed a weave pattern database indexing method, based on a smallest weave repeat unit. The authors suggest a classification of patterns into three categories: plain, twill and satin. Their algorithm uses the Fast Fourier Transform (FFT) to describe the distribution of the weave pattern crossover points, followed by an entropy based method to index each pattern. In contrast to our approach, theirs requires that weave patterns are manually classified by experts [Zheng *et al.* 2009].

Reverse Engineering Steps	Our Approach	[Schröder <i>et al.</i> 2015]
Capture Setup	Diffuse lighting, single picture.	Diffuse lighting, single picture.
Filtering	Fourier domain.	Spatial domain.
Weave Pattern Size Detection	After partitioning the image into horizontal and vertical strips. Measuring the pairwise IMED distance for all of them. Size in pixels and in n warps \times n wefts.	Auto-correlation analysis, size in pixels.
Segmentation	Seeded region growing like, using color and edges as cues.	Active yarn model.
Weave Pattern Estimation	Using cues like color and texture. Robustness achieved by means of a voting approach.	SANCC algorithm. Cues: color, response to horizontal and vertical templates, spatial relationship in a graphical model. MAP approach.
Noise Estimation	During segmentation, modelled with a Markov model.	Analysis of fine optical flow.

Table 1. Main differences between the proposed algorithm for reverse engineering the cloth and Schröder *et al.*'s method.

Computed Tomography (CT) scans were used by Zhao *et al.* to produce high resolution synthesized material models. The authors used CT scanning to acquire a database of volumetric 3D models of complex woven textiles. Their algorithm used photographs of fabric to recover albedo alone and fit this to the CT scanned volumetric model. The method is very effective for synthesizing large regions of fabric free from tiling artifacts, but is limited to modelling the appearance of a single type of yarn and can only reproduce textiles that have been previously scanned [Zhao *et al.* 2011]. In subsequent work, Zhao *et al.* [Zhao *et al.* 2012] generate new samples from similar or simpler exemplars of a textile. Since their algorithm is able to predict appearance of specific weave patterns, this method is very useful for generating volumetric models of complex and spatially varying woven textiles. Realistic close up renderings are achieved from two types input, the description of the material weave pattern, and a few examples of simpler fabrics. Their method allows quick creation of new exemplars of fabrics due to a fast synthesis algorithm. Our acquisition method shares the same limitation to woven patterns (*i.e.* as opposed to knitted textiles) and removes the requirement on CT scan data [Zhao *et al.* 2012]. In [Zhao *et al.* 2016] Zhao *et al.* fitted CT data to a procedural model built upon [Schröder *et al.* 2015], further augmented with a measurement-based flyaway fiber model.

To summarize, our approach advances the state of the art of woven fabric appearance estimation by: removing dependencies on expensive scanning equipment, relaxing the need for *a priori* yarn dimensions, increasing the variety of woven fabric patterns that can be estimated through our ability to generate and match to synthetic weave patterns, and considerably speeds up processing compared with alternative approaches as required in our pipeline for high volume fabric analysis for use in industry.

4 A NOVEL FABRIC MODEL RECOVERY PIPELINE

In the following sections we describe in more detail the stages in our pipeline, as previously outlined in Section 2.

Conceptually a woven fabric image can be viewed as the juxtaposition of a number of weave patterns, created by interlacing yarns with a limited number of colors, hence it should contain a significant amount of self-similarities. However, the appearance of real fabric images is affected by incident lighting, choice of acquisition device, and other factors such as dust and lint over the fabric sample. Hence, the image I of a fabric will contain all these factors, making the automatic analysis of woven fabric challenging.

Our pipeline relies on the properties of the 2D Discrete Fourier Transform (DFT) \mathcal{F} of the input image I of size $M \times N$, which represents the image content as a sum of an infinite number of complex exponentials (sinusoids) with different frequencies, thus resulting in a suitable way to identify regular structures:

$$\mathcal{F}(u, v) = \sum_{x=0}^{M-1} \sum_{y=0}^{N-1} I(x, y) e^{-j2\pi(ux/M + vy/N)} \quad (1)$$

where $x = 0, \dots, M-1$, $y = 0, \dots, N-1$ and $u = 0, \dots, M-1$, $v = 0, \dots, N-1$ are the frequency variables which span the frequency domain. The power spectrum of the transform is directly related to the bayesian probability that a sinusoidal frequency is present in the data [Fougère 2012]. Roughly speaking, repeating structures in the input data give rise to high energy peaks in the power spectrum of the DFT: the brighter the peak, the more frequently a repeating pattern has been observed in the input data. The value of the transform at the origin of the frequency domain $\mathcal{F}(0, 0)$, called *DC* coefficient, is equal to MN times the average value of $I(x, y)$.

Given the fabric formation process on a weaving machine, we consider the image of a fabric sample to be a linear combination of three different signals:

$$\mathcal{F}(I) = \mathcal{F}(I_r + I_l + I_n) = \mathcal{F}(I_r) + \mathcal{F}(I_l) + \mathcal{F}(I_n) \quad (2)$$

where I_r , I_l and I_n respectively are the ideal image of the fabric as it would appear if produced by a perfect loom and captured by a perfect acquisition device, the image of the irregularities introduced by the real loom and the image of the noise due to the acquisition device, dust, lint and so on. The second equality holds thanks to the linearity of the DFT.

It has been observed in previous work [Schröder et al. 2015] that individual repetitions of the weave patterns may not have the same size, yarns may be partially occluded by neighboring yarns. However, when averaging over several warped repetitions, the yarn level structure becomes clearly visible and allows us to handle the non-rigid deformations in the cloth. In probabilistic terms, this implies that if a big enough fabric sample is observed, thus containing several repetitions of the weave pattern, the probability of observing data belonging to I_r is higher than the probability of observing loom deformations and noise due to fibers and other factors. Similarly we observe that noise due to dust, lint and so on is spatially invariant and does not display regular structures, whereas the noise due to the loom and yarn friction (*i.e.* deformations of the fabric lattice) tends

to display spatial regularities. Moreover, fine-scale details and noise in the spatial domain correspond to high frequency terms, far away from the DC coefficient and the weave pattern frequencies. From these observations we safely assume that the lowest non-zero energy carried by a frequency term in $\mathcal{F}(I_r)$ is greater than the highest energy associated with any frequency term in $\mathcal{F}(I_l)$. If we denote $\mathcal{F}_+(Im)$ as the set of frequencies of the DFT of an image Im with associated energy greater than zero, we have that $\mathcal{F}_+(I_r) \cap \mathcal{F}_+(I_l) = \emptyset$.

The steps in our algorithm, devoted to separate the terms in Equation 2 and recover the textile model, are presented in pipeline order as this is the same as the processing order. We begin with our data capture setup.

4.1 Data Capture

We capture woven textile swatches using a consumer Nikon DSLR camera, fitted with a Nikon macro lens $105mm f/2.8AI - s$. The sample swatch is placed horizontally inside a hemispherical acquisition rig, employed to obtain uniform incident illumination. Note that any uniform illuminant could be used instead, and our method is not dependent on the use of our specific capture setup. The camera is placed 50 cm above and orthogonal to the sample. A set of round fiducial markers (with known size) are positioned at the corners of swatches, in order to recover the scale of the yarns. Although our algorithm does not require this scale information to reverse engineer the yarn model, we use it to obtain renderings with the correct scale of the woven cloth pattern. This extra information allows us to use different capture setups with similar lighting, by allowing proper scaling of all the samples in the same way, even when the camera distance to the sample differs (*e.g.* in close-up shots and for swatches with very fine threads). Images are saved in the camera's raw format and white balancing is performed, taking into account the known and controllable color temperature of the illuminant and the gamut of the camera CCD [Reinhard et al. 2008]. To correct for vignetting and variations in incident lighting along the sample, we simply capture a flat white reference which is used as a correction map to compensate for lens falloff. While only one photograph is required for weave pattern estimation, if desired, our pipeline allows optional acquisition of more images of the same sample with exposure variation, in order to recover a higher fidelity albedo map.

4.2 Initial Image Processing

In this section we describe a set of preliminary image manipulations carried out in advance of estimating the woven textile pattern. These include, image alignment, creating an albedo map and extracting a Fourier Signature.

4.2.1 Image Alignment. Since the capture process involves placing a fabric swatch on a surface, placement errors or discrepancies in the yarn orientations within the swatch orientation itself will result in the warp and weft directions in the input image I not being perfectly aligned with the x and y axis of the image formation plane. Thus, we perform a rotation of α degrees to achieve the necessary

alignment for later stages in our algorithm. We exploit the properties of the Radon Transform (RT) [Deans 1983], which allows us to reconstruct a 2D function f , by calculating the line integral on the possible space of straight lines L in \mathbb{R}^2 (\mathbb{R} refers to the set of real numbers):

$$RT_{s,\phi} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) \delta(x \cos \phi + y \sin \phi - s) dx dy \quad (3)$$

Where the line of integration is specified in polar coordinates (s, ϕ) and δ is the Kronecker delta.

As explained in Section 4, on the whole fabric image we expect to more frequently observe the regular structure of the weave pattern rather than the deformations due to the weaving process. This gives rise to regular peaks in the u, v domain, hence applying the Radon Transform to the image of the log magnitude of the Discrete FT $\mathcal{F}(I_r + I_l)$ is more reliable than applying it directly on the input image I . We search for peaks of the function:

$$\Gamma_I = \int \left| RT_{s,\phi}(\log \mathcal{F}(I_r + I_l)) \right| ds \quad (4)$$

For search, we employ the following criteria [Jafari-Khouzani and Soltanian-Zadeh 2005]:

$$\alpha = \arg \max_{\phi} \frac{d\sigma_{\phi}}{d\phi} \quad (5)$$

where σ_{ϕ} is the variance of the projection of the Radon Transform at ϕ . To remove the term $\mathcal{F}(I_n)$ from Equation 2, before computing the DFT we convolve I with a Gaussian mask G , to suppress the noise due to the acquisition device characteristics and lint, thus suppressing unwanted high frequency terms in the Fourier domain. The radius of the filter preserves the frequencies associated with I_r and I_l , since the structural information of the weave pattern is encoded in the low frequency terms. The FT of G is still a Gaussian, hence we can express the idea of this filtering with the following, where \odot denotes a spatial convolution:

$$\mathcal{F}(I \odot G) = \mathcal{F}(I) \cdot \mathcal{F}(G) \approx \mathcal{F}(I_r + I_l) \quad (6)$$

The log intensity image of the magnitude of $\mathcal{F}(I_r + I_l)$ is then filtered to further enhance the peaks and the projections along several directions are analyzed using the Radon Transform allowing estimation of α . The input image I is then corrected for unwanted rotation and further processed within the pipeline.

4.2.2 Fourier Signature (FS) Extraction. We want to find a compact way to characterize the relevant properties of a fabric swatch, namely its regular structure and the deformations introduced by the loom and the properties of the yarns. In particular, we aim to capture the systematic deformations introduced by the loom, *i.e.* those most commonly observed within the input picture, which carry a relatively high energy in the Fourier Domain.

The set of frequencies useful to characterize the fabric appearance hence includes the entire $\mathcal{F}(I_r)$ (*i.e.* the ideal regular structure of the fabric) and a subset of the highest ranking frequencies (in terms of associated magnitude in the power spectra) in $\mathcal{F}(I_l)$ (*i.e.* the most frequent loom deformations).

To extract the FS, we automatically select multiple square $N \times N$ pixel image patches (spanning the weave pattern) from random parts of the image and average the power spectrum of their FT. This is indicative of the variation of the signal across the image and particularly useful to boost the signal to noise ratio for dark swatches, which do not reflect enough light to the camera sensor. At this stage in the process, we do not know the weave pattern. Therefore we bootstrap the process via an initial assumption that the weave pattern repeat has unknown width Wr_w and height Wr_h (in pixels) and both are smaller than N . The highest ranking frequencies in the mean power spectra, in terms of their power, are selected such that the sum of their energy E_h is a fraction K of the total energy E_T of the spectra. Using our dataset as a benchmark, we found that a suitable value for K is 0.3, which maximizes the F_{score} [Rijsbergen 1979] for samples retrieval. The F_{score} is defined as:

$$F_{score} = 2 \times \frac{precision \times recall}{precision + recall} \quad (7)$$

The terms *precision* and *recall* can be defined as follows [Sokolova and Lapalme 2009]:

- *precision*: the number of correctly classified samples of a given class, divided by the number of examples labeled by the system as belonging to that class.
- *recall*: the number of correctly classified samples of a given class, divided by the number of samples of that class in the dataset.




Fig. 4. Four textile samples (top row) and their respective FSs (bottom row). The samples depicted in (a) and (b) have the same weave pattern, notice the similarities between their FSs (respectively in (e) and (f)). The samples reported in (c) and (d) have the same weave pattern and similar FSs (see (g) and (h)).

Natural irregularities in fabrics have implications on the shape of the FS, as can be seen in Figure 4. Each show two samples of the same cloth with their corresponding FS. Figure 5, shows the impact on the shape and intensity of peaks of the FS of a textile sample with black yarns.

We identify a suitable subspace to project the shape signature, by reducing the search space for features and so obtain a smaller set of feature vectors to be used for later classification. This is achieved by




Fig. 5. Two samples with the same weave pattern and their FSs. The sample reported in (b) has the same weave pattern as (a) but with fully black yarns.

analyzing the Fourier Signatures using Kernel Principal Component Analysis (KPCA) [Schölkopf et al. 1998]. This extension of classical PCA is ideal for dealing with non-linear cases through a non-linear function Φ , which maps the input space into a high dimensional feature space and then applies a linear method in the augmented space. In our case the input space is the set $\mathbb{F} = \{f_1, f_2, \dots, f_k\}$ of the k Fourier Spectra shapes in our dataset of inhomogeneous fabrics and $\Phi : \mathbb{R}^{N \times N} \rightarrow \mathbb{R}^{N_\Phi}$. In order to find the main axis of the augmented space \mathbb{R}^{N_Φ} KPCA diagonalises the following matrix M_Φ to find its eigenvalues:

$$M_\Phi = \frac{1}{k} \sum_{l=1}^k \left[\left(\Phi(f_l) - \frac{1}{k} \sum_{i=1}^k \Phi(f_i) \right) \left(\Phi(f_l) - \frac{1}{k} \sum_{i=1}^k \Phi(f_i) \right)^T \right] \quad (8)$$

The basis shapes of the kernel space are the eigenvectors $\{\mathbf{e}_i^\Phi\}$. Each $f_j \in \mathbb{F}$ can be reconstructed in the kernel space with the following equation:

$$\Phi(f_j) = \frac{1}{k} \sum_{i=1}^k \Phi(f_i) + \sum_{l=1}^{N_\Phi} v_{j,l}^\Phi \mathbf{e}_l^\Phi. \quad (9)$$

By varying the values of the set of $v_{j,l}^\Phi$ we can generate unseen shapes, a key aspect for enlarging our database to be used for homogeneous sample analysis. We select a subset of q principal eigenvalues $\{\lambda_1^\Phi, \lambda_2^\Phi, \dots, \lambda_q^\Phi\}$, such that they are able to explain the main variations observed in \mathbb{F} , with $d < N_\Phi$ and:

$$\lambda_1^\Phi \geq \lambda_2^\Phi \geq \dots \geq \lambda_q^\Phi \geq \dots \geq \lambda_{N_\Phi}^\Phi. \quad (10)$$

The vector $\mathbf{v}_j = [v_{j,1}^\Phi, v_{j,2}^\Phi, \dots, v_{j,q}^\Phi]$ constitutes the FS of the fabric. In our experiments we selected a Gaussian kernel, with variance equal to 3 times the average distance between the samples in the feature space. The set \mathbb{F} is constructed by analyzing a number of inhomogeneous fabrics through the pipeline described in Section 4.3, which also gives us the yarn parameters and weave patterns. Once \mathbb{F} has been analyzed by KPCA and the FSs obtained, we can populate our database and associate each signature to the corresponding yarns and noise parameters derived through the pipeline. The number of elements in our database can be extended generatively, to create arbitrary samples that can be used to build up further exemplars in our database for enhanced matching.

Given two samples $f_i, f_j \in \mathbb{F}$, with FS being respectively equal to \mathbf{v}_i and \mathbf{v}_j , their distance D_{FS} is simply defined as:

$$D_{FS}(f_i, f_j) = \frac{\mathbf{v}_i \cdot \mathbf{v}_j}{\|\mathbf{v}_i\| \|\mathbf{v}_j\|} \quad (11)$$

A database query triggers the calculation of the distances between the input sample and all the others stored in the database; the sample with the smallest distance is then retrieved from the database and its parameters can be used for the depth map and albedo estimation.

To determine if a sample is homogeneous, our algorithm automatically analyzes the histogram of I_r in order to locate the number of peaks; each peak is the initial estimate of a color cluster center. The distance between each pair of neighbouring clusters is computed, and if it is less than an empirically determined threshold (determined using a data set of 20 samples) the two clusters are merged and a new cluster center is computed. This process is repeated until no more clusters can be merged; if only one cluster is left the sample is labelled as homogeneous.

4.3 Inhomogeneous Fabric Analysis

In this section we detail the algorithm for inhomogeneous fabrics, which takes the primary processing route as shown in Figure 3. These fabrics are fairly regular and have two or more colors of yarns. We begin by describing how we determine the warp, weft, crossover points and yarn dimensions. The proposed pipeline for this stage of the algorithm is described in detail in the following section.

4.3.1 Weave Pattern Detection. We aim to segment the fabric image using a set of features which includes the color, hence the aforementioned factors need to be handled using low-level image processing techniques both in the frequency and spatial domains. From the segmented image, it is possible to exploit self-similarity to identify repeat patterns in the fabric sample.

Recall that the regular structure of a fabric sample, contained in the image I_r , should be clearly noticeable in the frequency domain, where the brightest peaks give information about the fundamental frequencies related to the thread density in weft and warp directions.

Higher frequency terms without associated high energy can be explained by noise and fine details, contained in the image I_n , filtered in earlier stages of our pipeline (see Section 4.2.1) by means of a convolution with a Gaussian kernel G .

From Equation 6, taking into account the assumptions described in Section 4 we can obtain $\mathcal{F}(I_r)$ from $\mathcal{F}(I_r + I_l)$. An automatic selection of the peaks (*i.e.* strong local maxima) in $\mathcal{F}(I_r + I_l)$ is performed and each of them is added to the set of frequencies of $\mathcal{F}(I_r)$; thanks to the symmetry property of the DFT, the search space in the frequency domain is reduced to $1/2$, since the symmetric point of a peak in the frequency domain is also a peak (more precisely the real part is symmetric and the imaginary part is anti-symmetric). We do not search for a predefined number of peaks, as this would put a limitation on the typology of fabric we can handle. Instead, we aim to add as many strong peaks as possible, by thresholding the gradient values of the magnitude of the DFT to the 99.8 percentile of the gradient values. From the set of peaks we select the one(s) with the minimum energy, which gives a threshold to discriminate the frequencies in $\mathcal{F}(I_r)$ from the frequencies in $\mathcal{F}(I_l)$.

We can finally derive the estimate of the ideal fabric image I_r simply as $\mathcal{F}^{-1}(I_r) \odot G^{-1}$, where G^{-1} is the high-pass kernel obtained as the inverse of G . The described procedure is repeated for each color channel individually, combining the per channel estimated into a single RGB image. Examples of I_r are shown in Figure 6.




Fig. 6. Two fabric swatches (6(a) and 6(c)) and their corresponding estimated ideal images I_r (6(b) and 6(d)).

Under diffuse lighting we obtain a photograph of the fabric swatch from the normal direction, in which the centers of the yarns have the highest intensities and the intersections between yarns appear dark [Schröder et al. 2015]. This is particularly clear by analyzing I_r (see Figure 6), where the noise and the sharp specular highlights have been filtered out and enables us to derive a first estimate of the position of these intersections by examining the local minima of the projection profiles, along both the horizontal and vertical directions, of the ideal model of the fabric image. The projection profiles are calculated from the gradient values, since they allow us to obtain a more reliable estimate of the regular grid of yarn intersections.

At this stage in the algorithm, we now have a set of estimated intersection positions such that a yarn passes through two neighboring ones (*i.e.* their distance defines the yarn width). We model the relative set of yarn widths as gamma distributions (one per yarn color) filtering out the outliers. We roughly partition the image into a set of strips, each of which contains one weft. By measuring the distance between each pair of strips, we obtain a matrix of distances. The basis for this is that the minimum distance between the image of a

weft (with all the warps going over it) and another one is achieved periodically, when the weave pattern repeats itself. We employ the IMage Euclidean Distance (IMED) [Wang et al. 2005], a variation of the classic Euclidean distance which proves to be robust to local deformations by taking into account the relative position of pixels by considering the angles between them.

Let two images x and y , of size $M \times N$, be $x = (x^1, x^2, \dots, x^{MN})$ and $y = (y^1, y^2, \dots, y^{MN})$, where x^{kN+l} indicates the gray value in image x at location (k, l) . Their IMED distance $d(x, y)$ is given by
$$d^2(x, y) = \frac{1}{2\pi\sigma^2} \sum_{i,j=1}^{MN} \left(\exp \left\{ \frac{-|P_i - P_j|^2}{2\sigma^2} \right\} (x^i - y^i)(x^j - y^j) \right),$$

where $|P_i - P_j|$ is the distance between pixels P_i and P_j on the image lattice. Through examining the resulting matrix, we are able to locate the diagonals which show the lowest cumulative values. The distance between them gives us the number of wefts N_e in the weave pattern; an example is shown in Figure 7. This process is repeated for warps, obtaining their number N_a . From this we extract the weave repeat which enables us to roughly partition the image in terms of the elementary weave repeat pattern of size $N_e \times N_a$.

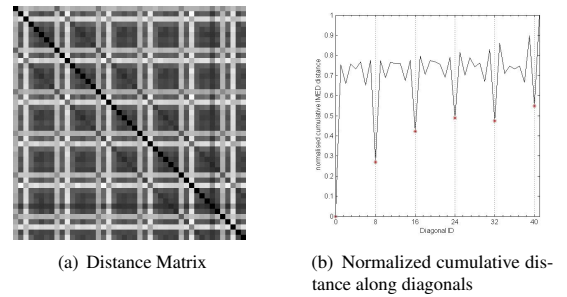


Fig. 7. Distance matrix among wefts in a sample area of Figure 4(a) and (b), left side. Due to symmetry, only the upper triangular part needs to be calculated. The plot of the normalized cumulative distance on the right, calculated along the diagonals, refers to the upper triangular matrix (*i.e.* diagonal 0 is the main diagonal).

As previously mentioned due to yarn and loom characteristics, the distribution of the yarns is irregular in the input image I . The estimated yarn arrangement thus requires further refinement to match the observed input image of the fabric. We aim to deform the regular grid obtained from previous steps by analyzing I_r to achieve a good match with the crossed areas of warp and weft yarns. The center positions of each cell in the regular grid are used to seed region growing [Adams and Bischof 1994] and the features we use for the algorithm are the CIELAB [Wyszecki and Stiles 1982] color values from the original input image. For the distance metric, we use the ΔE_{00}^* color difference formula, with a merging threshold set to $\Delta E_{00}^* = 2.3$ defined as the Just Noticeable Difference (JND) [Sharma 2002].

Each of the regions found in the previous step form an irregular grid $N_e \times N_a$, in which each cell encloses a weft/warp crossing. To determine whether it is the warp passing over the weft or vice versa, we analyze the texture features in a rectangular fixed window around

each detected crossing. We use Local Binary Pattern and Support Vector Machines [Ojala et al. 2002] to classify the type of crossing and bootstrap the classification exploiting the presence of elongated structure aligned along the warps or wefts directions, to solve the ambiguities between the two possible outcomes of the classification, namely the correct weave pattern $W = (w_{i,j}) \in \{0, 1\}^{N_e \times N_a}$ or the inverse (*i.e.* $1 - W$).

We solve each estimated weave repeat and output a sequence of pattern matrixes $W_l, l = 1, \dots, L$, where L is the number of repetitions of the weave pattern in the input image. The set of W_l are used as input to a voting scheme to ensure better reliability and minimize classification errors. The final estimated weave pattern W_E is then given by:

$$W_E(i, j) = \left(\sum_{l=1}^L W_l(i, j) \right) > L/2 \quad (12)$$

with $i = 1, \dots, N_e$ and $j = 1, \dots, N_a$.

4.3.2 Noise Estimation. At this stage in our pipeline, each individual yarn has been segmented and so we know if it is a warp or a weft yarn and its dimensions. We approximate the shape of these yarns with an ellipsoid with the yarn dimensions as it will puff out a little between crossing points. The traversal along one axis of the yarn, depending on its classification as a warp or a weft, gives us a measure of the degree of deviation of the yarn from the ideal horizontal or vertical. The other axis provides an estimate of yarn width at its thickest point, thus allowing us to refine the distribution of the yarn widths, previously characterized as Γ distributed, as described in Section 4.3.1. Hence, by tracking the path of the yarn (just following the estimated weave pattern) and its neighbours, we are able to gather information about how to model the noise in the distribution of the yarns. Each yarn Y_i is hence modeled as a tuple (a_i, b_i, c_i, d_i) , where the first two elements are respectively the shape parameter and the scale parameter of the Γ distribution controlling the width of Y_i , whereas c_i and d_i are mean and variance of the gaussian distribution controlling the angle of deviation of the yarn.

4.4 Homogeneous Fabric Analysis

We will inevitably encounter sample textiles for which the above technique cannot be used to determine weave patterns, due to the large variation of types of yarns and colors used in woven fabrics. In particular, we found homogenous fabrics to be challenging as well as those with dark colored yarns. To handle these cases, we take the second route through our pipeline overview (shown in Figure 3) instead.

We derive the FS of the sample and search our database for the closest match of woven pattern among the data saved from previously encountered samples and assign the associated depth map parameters to the recovered model for the input sample. The distance measure is given by the angular similarity between the input FS and the candidate's FS in the dataset. An ideal match would find the same weave pattern and allow to use the previously stored parameters to synthesise the depth map; if the same weave pattern has not been stored previously, the parameters associated with the FS displaying

the highest angular similarity in the dataset will be retrieved. For a densely populated dataset, this would mean to retrieve the most similar weave pattern.

Recall that we build a database including the estimated weave patterns and the inferred depth maps from all our samples. This contains the FS of weave patterns for every sample we encounter, extracted by examining the highest ranking frequencies in the Fourier Transform of the input fabric image. For matching purposes, and dimensionality reduction, we perform Kernel Principal Component Analysis (KPCA) on the space of signatures to obtain a small feature vector storing just the first N main modes. This allows for optimized lookup as the entire signature for each pattern is not stored.

In addition, we can also procedurally generate synthetic depth maps for woven fabrics using any arbitrary weave pattern, plausible yarn parameters, and noise with typical observed values for a class of materials. This enables us to obtain a FS matching those extracted from photographs of a real fabric woven with a similar yarn and employing a similar weave pattern to the one simulated. Once our database has been sufficiently populated (with at least 30 samples), we can simply use the photograph of the challenging swatch to extract its FS and query the database using its first N_K KPCA modes as a feature vector. In order to derive a dataset with a reasonable size we analysed the fabric samples provided by KiSP Inc. and found out that the most common weave samples have very similar weave patterns, hence we used them to bootstrap our database. Consequently, we are able to robustly assign a depth map to samples that cannot be reliably analyzed.

4.5 Depth Map Creation

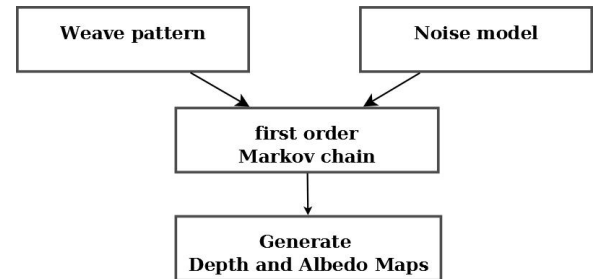


Fig. 8. Depth map generation.




Fig. 9. Depth map examples; from (a) to (c) the maps are obtained by re-injecting larger subsets of the estimated noise parameters.

From the parameters estimated in the previous step (weave pattern, yarn width, gaps among warp and weft) we can extract a regular




Fig. 10. Row 1: A fabric swatch with a 8×24 weave pattern; Row 2: A fabric swatch with a 6×10 weave pattern; Row 3: A fabric swatch with a 6×6 weave pattern. The estimated minimal weave patterns in 10(d), 10(k), 10(r), refer to the areas highlighted in blue respectively in 10(b), 10(i), 10(p), scaled to a 1:1 ratio between width and height.

depth map without noise. For visual richness, the noise needs to be accounted for. The measured yarn parameters are used to infer a depth map for each sample, approximating yarns with micro-cylinders [Sadeghi et al. 2013] [Irawan and Marschner 2012] and re-injecting our estimated noise in each of the parameters, including the cross-section, to avoid unnatural repetition of the texture when tiling is performed for final rendering. In common with other approaches, we assume the fabric to be single-layered [Zhao et al. 2012]. A 1^{st} order Markov chain is used to solve the positions of the yarns given the noise model and also determining the yarn tilt angle based on the width of the threads and the distance among the crossings. The output of this process allows us to infer a 3D model of the fabric sample, building a 3D convex envelop for each yarn. By assuming the sample is laid flat on a plane, the distance of each visible surface point from the plane provides us with the depth at that location of the sample. Figures 8 and 9 show the steps in the depth map estimation process and example depth maps respectively.

The 1^{st} order Markov model applies the following set of states $S = \{a_n, a_l, a_r, tn_n, tn_l, tn_r, tk_n, tk_l, tk_r\}$, such that: $\mathbf{P}\{S^{t+1}|S^1, \dots, S^{t-1}, S^t\} = \mathbf{P}\{S^{t+1}|S^t\}$. The set of states is derived from the information gathered as described in Section 4.3.2: $a \rightarrow$ average width yarn, $tn \rightarrow$ thinner yarn, $tk \rightarrow$ thicker yarn. The subscripts refer to the following yarn states: $n \rightarrow$ no deviation, $l \rightarrow$ tilted left, $r \rightarrow$ tilted right. The transition probability \mathbf{P} from a state S^t at time t to a state S^{t+1} at time $t + 1$ is derived from the input sample, as described in Section 4.3.2.

4.5.1 Normal map and Albedo Map. For many materials the visual complexity is mainly due to the geometric complexity, rather than the actual reflectance [Magda and Kriegman 2006] and the use of relatively simple albedo map estimation procedures has been demonstrated and perceptually validated for textured surfaces [Glencross et al. 2008] and planar samples [Riviere et al. 2016].

To estimate the diffuse albedo we first derive the surface normals at the yarn level, which have been previously approximated with the micro-cylinders (see Section 4.5). The knowledge of the surface normals allows us to account for the cosine foreshortening. We assume that a thread cannot change color along its path and it is made of homogeneous fibers with the same color. Hence, for each thread, the median of the estimated intensities provides us with the diffuse albedo, whereas the top 20% of the intensities is used for the specular level. Finally, we add 3D Perlin Noise [Perlin 1985] to the micro-cylinder derived normal map, to visually simulate the effect of the yarn fibers. If a higher quality albedo map is desired, we capture seven exposures of the swatch over a range of 3 f-stops and combine these using Greg Ward’s command line high dynamic range image builder [Ward 2001].

4.6 Renderings

If the input fabric shows a reasonable amount of specularities this does not pose a problem to our pipeline. If the cloth is very shiny and specular, then cross-polarized photography could be used to capture the input sample.




Fig. 11. A fabric swatch with a 26×100 weave pattern. Due to the size of the weave pattern the missing row in the estimated one does not affect visually the rendered result.

At this stage we have an albedo map, a specular level, normal map and a depth map for our fabric sample, which we render using Autodesk 3ds Max, although any rendering algorithm could be used. The light model we used is a Target Directional Light, which casts parallel light rays in a single direction, as this approximates well the light sources installed in our acquisition setup. We used an anisotropic BRDF model and set the Index of Refraction to 1.4 if information about the fibers used for the threads was not available.

5 RESULTS AND EVALUATION

We demonstrate the effectiveness of our pipeline using four textile samples processed through the inhomogeneous fabric pipeline, acquired from fabric swatches exhibiting very different weave patterns, yarn colors, and weave repeats. Specifically, 8×24 (wefts \times warps), 6×10 , 6×6 and 24×100 weave pattern swatches have been used. In the first three cases, we show that the weave pattern has been correctly recognized as can be seen in Figure 10. One row of the weave pattern, in the fourth sample (shown in Figure 11) has been misclassified. However, the generated depth map and corresponding rendering are not negatively affected due to the small size of the missing row compared with the weave pattern.

To evaluate the visual accuracy of our recovered textile model, in Figure 12, we show comparisons between our rendered images and photographs of the same fabric samples under identical lighting conditions. From the estimated depth map of the fabric we create Normal/Displacement/Bump Maps to use for the renderings. Our estimated noise, yarns and weave patterns have been used for all the renderings. In Figure 15 we show the renderings of recovered textile models, mapped on the surface of a cylinder to show the benefits of the depth map; close up images are shown in the insets. Recall that the effectiveness of the Fourier Signature was previously demonstrated in Figure 4 which shows that the same weave pattern can be easily recognized and that the Fourier Signature is very different for different classes of swatches. Moreover a completely black swatch, which cannot be correctly classified with our inhomogeneous fabric pipeline is correctly classified using our homogenous pipeline and shown in Figure 5; in Figure 13(e) we report the rendering for this swatch, in which the albedo estimation is negatively affected by the very dark color and low signal-to-noise ratio. We also included the results obtained by analyzing two more fabric swatches, other than the aforementioned black one, through the homogenous fabric pipeline (*i.e.* using the FS to retrieve the best matching parameters from our

database). The ground truth photographs and the rendering are reported in Figure 13. In particular, Sample 5 (Figure 13(a)) contains only one full repetition of its 48×200 weave pattern and our pipeline is not able to robustly estimate it; the use of the database allows use to find a close matching in Sample 4 (Figure 11(a)) and from this derive the model parameters. As for Sample 6 (Figure 13(c)), its weave pattern is ambiguous and even with a careful visual inspection it is not possible to determine the number of repetitions; however, our homogeneous fabric pipeline found the closest matching sample in the database and estimated a suitable model in about 29 seconds. In Table 2 the processing times for each of the samples is reported, estimated using an unoptimized single-core Matlab implementation on a laptop. Finally, we also compare with a recently published alternative method proposed by Schröder *et al.* [Schröder et al. 2015], using the same fabric photographs (kindly provided by the authors).

In Figure 14, we show the resulting correctly estimated weave patterns and corresponding renderings, it is important to note that these are for visual comparison only. These images are rendered under uniform lighting and by visually matching parameters (such as the camera lens and distance to the swatches) as the complete set of rendering and modeling parameters are not reported in Schröder *et al.* [Schröder et al. 2015]. I_r is the simplified image (ideal loom, perfect acquisition).

Sample ID	N° of weave repeats	Processing time
Sample 1 (Figure 12(e))	90	375 sec
Sample 2 (Figure 12(i))	443	85 sec
Sample 3 (Figure 12(a))	201	132 sec
Sample 4 (Figure 11(a))	9	318 sec
Sample 5 (Figure 13(a))	1	55 sec
Sample 6 (Figure 13(c))	Unknown	29 sec
Sample 7 (Figure 13(e))	190	47 sec

Table 2. Processing times for estimating a fabric model for the sample reported in Figures 11- 13, excluding renderings. Samples 1-4 have been processed through the inhomogeneous pipeline, whereas samples 5-7 through the homogenous pipeline.

6 LIMITATIONS

Our algorithm assumes that our capture is performed under diffuse lighting, the entire weave repeat is contained within the image, the camera position is fixed, and the sample swatch is planar without significant deformations. Consequently, our approach has some limitations. First the input image needs to contain the whole weave pattern repeat in order for our algorithm to obtain a robust Fourier Signature that is invariant to translations and rotations of the samples. The second limitation is that for our voting algorithm to properly classify the patterns observed in the swatch, we require multiple observations in the image, similar to [Schröder et al. 2015], meaning that the weave repeat needs to be small with respect to the size of the captured swatch. Third, in common with other methods [Schröder et al. 2015; Zhao et al. 2012], we currently cannot estimate woven




Fig. 12. Photographs of three fabric samples and corresponding renderings, obtained using our recovered textile model (inhomogeneous fabric analysis). The lighting condition used for the top view (first two columns) is different from the one used in the oblique view (last two columns).

patterns for textiles that do not observe a grid like woven structure, such as knitted textiles. The first two limitations can be addressed in future work by incorporating automatic techniques to stitch together photographs of sub-areas of the sample textile [Brown and Lowe 2007; Szeliski 2010] or by generating synthetic woven patterns that are larger than a swatch and matching an observed portion to the larger pattern structure. The third limitation can be addressed in future work by using a von Mises distribution instead of the current integral projection used in our implementation [Dhillon and Sra 2003].

Finally, unlike Schröder *et al.* [Schröder et al. 2015], although we capture a volumetric model of the woven cloth (inferring the yarn path and volume) we do not model the light scattering at the fibre level. Our target application, necessitates certain trade-offs when compared with the work of Schröder *et al.* [Schröder et al. 2015]. While modeling the light scattering at the fibre level leads to a boost in visual quality, it is not critical for our application in Virtual Reality visualization of interiors for the contract furniture industry. However, conceptually, there is nothing preventing our algorithm from being extended in this way.

7 CONCLUSIONS

To conclude, we have demonstrated a pipeline for reverse engineering 3D textile models of woven fabrics, we have presented four novel algorithms within our pipeline to extract volumetric fabric models using a single image as input. Our method computes a model in up to a few minutes. The creation of our database indexed through the Fourier Signature allows us to rapidly obtain all the required information for rendering by simply matching the feature vector derived from the Fourier spectra. We are also able to construct woven pattern depth maps in a generative fashion and match observations in our image captures to these generated patterns. This means that our algorithm can build up a rich variety of datasets without needing to measure every possible woven fabric pattern we may encounter. Consequently, we can also generate woven patterns that are larger than a swatch and match an observed portion to the larger pattern structure. This means that our method is more suited to its target application in the contract furniture industry, in which thousands of swatches needs to be rapidly classified and might share the same weave pattern but have different colors. In particular, this also allows us to deal with completely black samples, provided a similar weave pattern has been previously classified. The results of our estimation pipeline are illustrated using several woven textile samples, containing very different weave patterns and yarns. We also compare with ground truth data and also with previous similar work. We have highlighted




Fig. 13. Photographs of three fabric samples and corresponding renderings under the same lighting condition, obtained using our recovered textile model (homogeneous fabric analysis - database access).




Fig. 14. Weave pattern detection from samples in Schröder et al. [2015]. Row 1: 8×8 plain weave with color effects; Row 2: 2×2 plain weave; Row 3: 10×10 complex weave with color effects.

limitations of our methods and also propose solutions to these. If sufficient numbers of fabric swatches are characterized (captured and/or synthesized), then our solution for homogenous textiles could

be used exclusively for fast inference of the textile model from a photograph of any woven fabric swatches.

ACKNOWLEDGMENTS

The authors wish to thank Robert Kendal from KiSP Inc. and Ian Hall from Yulio Inc. for their financial support for this research. We thank Dar'ya Guarnera for help with rendering. Andrew Dann and Dar'ya Guarnera gave invaluable support and assistance with proof reading many versions of the manuscript. Kai Schröder kindly made available the input photographs used in Figure 14. Thanks also to the anonymous ACM TOG reviewers for their astute comments and suggestions, which served to help us improve our paper.

REFERENCES

- R. Adams and L. Bischof. 1994. Seeded Region Growing. *IEEE Trans. Pattern Anal. Mach. Intell.* 16, 6 (June 1994), 641–647. <https://doi.org/10.1109/34.295913>
- M. Brown and D. G. Lowe. 2007. Automatic Panoramic Image Stitching using Invariant Features. *International Journal of Computer Vision* 74, 1 (2007), 59–73.
- Kristin J. Dana, Bram van Ginneken, Shree K. Nayar, and Jan J. Koenderink. 1999. Reflectance and Texture of Real-world Surfaces. *ACM Trans. Graph.* 18, 1 (Jan. 1999), 1–34.
- S.R. Deans. 1983. *The Radon transform and some of its applications*. Wiley.
- Inderjit S Dhillon and Suvrit Sra. 2003. *Modeling data using directional distributions*. Technical Report.
- J. Filip and M. Haindl. 2009. Bidirectional Texture Function Modeling: A State of the Art Survey. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 31, 11 (Nov 2009), 1921–1940. <https://doi.org/10.1109/TPAMI.2008.246>
- Paul F Fougère. 2012. *Maximum entropy and Bayesian methods*. Vol. 39. Springer Science & Business Media.
- Mashhuda Glencross, Gregory J. Ward, Francho Melendez, Caroline Jay, Jun Liu, and Roger Hubbard. 2008. A Perceptually Validated Model for Surface Depth Hallucination. *ACM Trans. Graph.* 27, 3, Article 59 (Aug. 2008), 8 pages.
- Chang-Chiun Huang, Sun-Chong Liu, and Wen-Hong Yu. 2000. Woven fabric analysis by image processing Part I: identification of weave patterns. *Textile Research Journal* 70, 6 (2000), 481–485.
- Piti Irawan and Steve Marschner. 2012. Specular Reflection from Woven Cloth. *ACM Trans. Graph.* 31, 1, Article 11 (Feb. 2012), 20 pages. <https://doi.org/10.1145/2077341.2077352>
- K. Jafari-Khouzani and H. Soltanian-Zadeh. 2005. Radon transform orientation estimation for rotation invariant texture analysis. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 27, 6 (June 2005), 1004–1008. <https://doi.org/10.1109/TPAMI.2005.126>
- Mizuho Kinoshita, Yositada Hashimoto, Ryuichi Akiyama, and Sei Uchiyama. 1989. Determination of Weave Type in Woven Fabric by Digital Image Processing. *Journal of the Textile Machinery Society of Japan* 35, 2 (1989), 1–4. https://doi.org/10.4188/jte1955.35.2_1
- Kuo, Chung-Yang Shih, and Jiunn-Yih Lee. 2004. Automatic Recognition of Fabric Weave Patterns by a Fuzzy C-Means Clustering Method. *Textile Research Journal* 74, 2 (2004), 107–111. <https://doi.org/10.1177/004051750407400204> arXiv: <http://trj.sagepub.com/content/74/2/107.full.pdf+html>
- Sebastian Magda and David J Kriegman. 2006. Reconstruction of Volumetric Surface Textures for Real-Time Rendering.
- T. Ojala, M. Pietikainen, and T. Maenpää. 2002. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 24, 7 (Jul 2002), 971–987. <https://doi.org/10.1109/TPAMI.2002.1017623>
- Ken Perlin. 1985. An Image Synthesizer. *SIGGRAPH Comput. Graph.* 19, 3 (July 1985), 287–296. <https://doi.org/10.1145/325165.325247>
- Miquel Ralló, Jaume Escofet, and María S Millán. 2003. Weave-repeat identification by structural analysis of fabric images. *Applied optics* 42, 17 (2003), 3361–3372.
- SA Hosseini Ravandi and K Toriumi. 1995. Fourier transform analysis of plain weave fabric appearance. *Textile Research Journal* 65, 11 (1995), 676–683.
- Erik Reinhard, Erum Arif Khan, Ahmet Oguz Akyz, and Garrett M. Johnson. 2008. *Color Imaging: Fundamentals and Applications*. A. K. Peters, Ltd., Natick, MA, USA.
- C. J. Van Rijsbergen. 1979. *Information Retrieval* (2nd ed.). Butterworth-Heinemann, Newton, MA, USA.
- J. Riviere, P. Peers, and A. Ghosh. 2016. Mobile Surface Reflectometry. *Computer Graphics Forum* 35, 1 (2016), 191–202. <https://doi.org/10.1111/cgfg.12719>

- Iman Sadeghi, Oleg Bisker, Joachim De Deken, and Henrik Wann Jensen. 2013. A Practical Microcylinder Appearance Model for Cloth Rendering. *ACM Trans. Graph.* 32, 2, Article 14 (April 2013), 12 pages. <https://doi.org/10.1145/2451236.2451240>
- Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Müller. 1998. Nonlinear Component Analysis As a Kernel Eigenvalue Problem. *Neural Comput.* 10, 5 (July 1998), 1299–1319. <https://doi.org/10.1162/089976698300017467>
- K. Schröder, A. Zinke, and R. Klein. 2015. Image-Based Reverse Engineering and Visual Prototyping of Woven Cloth. *Visualization and Computer Graphics, IEEE Transactions on* 21, 2 (Feb 2015), 188–200. <https://doi.org/10.1109/TVCG.2014.2339831>
- Christopher Schwartz, Ralf Sarlette, Michael Weinmann, and Reinhard Klein. 2013. DOME II: A Parallelized BTF Acquisition System. In *Eurographics Workshop on Material Appearance Modeling: Issues and Acquisition*, Holly Rushmeier and Reinhard Klein (Eds.). Eurographics Association, 25–31. <https://doi.org/10.2312/MAM.MAM2013.025-031>
- Gaurav Sharma. 2002. *Digital Color Imaging Handbook*. CRC Press, Inc., Boca Raton, FL, USA.
- Marina Sokolova and Guy Lapalme. 2009. A Systematic Analysis of Performance Measures for Classification Tasks. *Inf. Process. Manage.* 45, 4 (July 2009), 427–437. <https://doi.org/10.1016/j.ipm.2009.03.002>
- Richard Szeliski. 2010. *Computer vision: algorithms and applications*. Springer.
- B.J. Thompson. 1978. Optical transforms and coherent processing systems with insights from crystallography. In *Optical Data Processing*, David Casasent (Ed.). Topics in Applied Physics, Vol. 23. Springer Berlin Heidelberg, 17–52. <https://doi.org/10.1007/BFb0057982>
- Liwei Wang, Yan Zhang, and Jufu Feng. 2005. On the Euclidean distance of images. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 27, 8 (Aug 2005), 1334–1339. <https://doi.org/10.1109/TPAMI.2005.165>
- Gregory J. Ward. 2001. HDR Image Builder. <http://www.anywhere.com>. (2001). Accessed: 2015-01-16.
- Gunter Wyszecki and Walter Stanley Stiles. 1982. *Color science: Concepts and Methods, Quantitative Data and Formulae*. Vol. 8.
- Binjie Xin, Jinlian Hu, George Baci, and Xiaobo Yu. 2009. Investigation on the Classification of Weave Pattern Based on an Active Grid Model. *Textile Research Journal* 79, 12 (2009), 1123–1134. <https://doi.org/10.1177/0040517508101459> arXiv:<http://trj.sagepub.com/content/79/12/1123.full.pdf+html>
- Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. 2011. Building volumetric appearance models of fabric using micro CT imaging. In *ACM Transactions on Graphics (TOG)*, Vol. 30. ACM, 44.
- Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. 2012. Structure-aware Synthesis for Predictive Woven Fabric Appearance. *ACM Trans. Graph.* 31, 4, Article 75 (July 2012), 10 pages. <https://doi.org/10.1145/2185520.2185571>
- Shuang Zhao, Fujun Luan, and Kavita Bala. 2016. Fitting Procedural Yarn Models for Realistic Cloth Rendering. *ACM Trans. Graph.* 35, 4, Article 51 (July 2016), 11 pages. <https://doi.org/10.1145/2897824.2925932>
- Dejun Zheng, George Baci, and Jinlian Hu. 2009. Accurate indexing and classification for fabric weave patterns using entropy-based approach. In *Cognitive Informatics, 2009. ICCI'09. 8th IEEE International Conference on*. IEEE, 357–364.

Received August 2015; revised July 2016; final version June 2017; accepted July 2017




Fig. 15. Woven fabric models shown on a cylinder for illustration. (a) Rendering of the fabric shown in Figure 10(a); (b) Rendering of the fabric shown in Figure 10(h); (c) Rendering of the fabric shown in Figure 10(o);