

Capturing and Stylizing Hair for 3D Fabrication

Jose I. Echevarria^{1,2} Derek Bradley¹ Diego Gutierrez² Thabo Beeler¹
1) Disney Research Zurich 2) Universidad de Zaragoza



Figure 1: We present the first method to capture an individual’s hair style (left) in a manner suitable for miniaturization and physical reproduction while still faithfully preserving the essential visual look of the style (right).

Abstract

Recently, we have seen a growing trend in the design and fabrication of personalized figurines, created by scanning real people and then physically reproducing miniature statues with 3D printers. This is currently a hot topic both in academia and industry, and the printed figurines are gaining more and more realism, especially with state-of-the-art facial scanning technology improving. However, current systems all contain the same limitation - no previous method is able to suitably capture personalized *hair-styles* for physical reproduction. Typically, the subject’s hair is approximated very coarsely or replaced completely with a template model. In this paper we present the first method for *stylized* hair capture, a technique to reconstruct an individual’s actual hair-style in a manner suitable for physical reproduction. Inspired by centuries-old artistic sculptures, our method generates hair as a closed-manifold surface, yet contains the structural and color elements stylized in a way that captures the defining characteristics of the hair-style. The key to our approach is a novel multi-view stylization algorithm, which extends feature-preserving color filtering from 2D images to irregular manifolds in 3D, and introduces abstract geometric details that are coherent with the color stylization. The proposed technique fits naturally in traditional pipelines for figurine reproduction, and we demonstrate the robustness and versatility of our approach by capturing several subjects with widely varying hair-styles.

CR Categories: I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture—Scanning

Keywords: Hair Capture, Stylization, Printable Figurines.

Links: [DL](#) [PDF](#)

1 Introduction

A mainstream goal in computer graphics is to create data-driven methods for building geometric models of humans. In recent years we have seen advances in 3D face and body scanning, motion capture and real-time performance capture. While human scanning has many applications in fields like video games, films and medical analysis, a fast-growing field is the physical reproduction of miniature statues or figurines. Physical reproduction in general, and particularly of humans, has become a hot topic both in academia [Li et al. 2013; Sturm et al. 2013; Tena et al. 2013] and industry [3DSystems 2013; 3D-u 2010; D-Tech Me 2012; Omote 3D 2012; PocketSize Me 2013]. Recently, *3D-Systems* even announced the release of a 3D-photobooth, which will facilitate 3D-portraits for the masses. The underlying pipeline of all these systems is essentially the same: A person is scanned, the resulting mesh is processed often with artist interaction, and the figurine is printed using a 3D printer. Consequently, all systems have similar drawbacks, and in particular, no previous approach can capture personalized *hair-styles* with adequate details, while being suitable for physical reproduction.

Almost as much as the face, a person’s hair-style is a defining characteristic of an individual. The reproduction of figurines without properly capturing the hair is a severe limitation of current systems, since the hair-style contributes so substantially to the person’s identity. Existing research in hair capture methods either focus on reconstructing highly-detailed individual wisps or hair strands [Luo et al. 2013; Beeler et al. 2012; Paris et al. 2008], which do not meet the physical manufacturing constraint of a closed manifold surface, or they produce coarser reconstructions [Luo et al. 2012] that lack the level of stylization, detail or colors required to produce appealing 3D-printed models. In this work we present the first method for *stylized* hair capture, which addresses current limitations of physical reproduction systems, enabling the faithful miniaturization and physical reproduction of figurines with drastically varying hair-styles. Our method automatically reduces the complexity of hair to an abstract, printable 3D surface, while still capturing the essential structural and color elements that define its style. The proposed method fits naturally into existing physical reproduction pipelines, and so our work has the potential to significantly impact the growing industry of figurine reproduction.

Our work is inspired by existing artistic sculptures of hair. For centuries, artists have shown that the essence of a hair-style can be represented on a continuous manifold surface, such as marble or clay, through stylized sculpting of geometric details (Figure 2, left). Today in the digital world, CG modelers follow the same principles by virtually sculpting hair structure on a 3D mesh (Figure 2, right). Our goal is to computationally achieve a similar level of abstraction in the captured hair-style of an individual.



Figure 2: *Our work is inspired by artistic sculptures of hair in the real world and those created by digital artists in professional modeling software.*

To this end, we start by obtaining a smooth surface representation of the hair-style from multi-view stereo reconstruction. As our primary contribution, we introduce a novel color stylization operator that works directly over the geometric mesh domain, and can be applied over non-uniform manifold surfaces. This way, color information can be sampled, stored and processed in a consistent way with respect to the input views. As it is important to retain the appearance of directional wisps and the overall flow of hair, the color is stylized over the mesh using a combination of directional smoothing and orthogonal shock filters, inspired by analogous 2D image stylization [Kang et al. 2009; Kyprianidis and Kang 2011]. The per-vertex stylized color is then used to generate coherent geometric displacements over the surface, effectively stylizing the shape as well. The final result is a printable surface that can be miniaturized so that both geometry and color convey the hair-style of the captured person (see Figure 1). Our method allows the user to adjust the level of abstraction to match the scale of the final printout and to achieve different visual styles, which behave in a consistent way no matter the complexity of the original hair-style. Ultimately, the stylized hair is combined with state-of-the-art face scanning and traditional 3D printing methods for fabricating full-head figurines. We show the flexibility of our approach by reconstructing a large number of varied, complex hair-styles, and even non-human furry objects.

2 Related Work

Our work is related to methods for creating personalized figurines, reconstructing hair, geometry abstraction and image stylization.

Personalized Figurines. A large variety of methods are being employed to capture a person in 3D, ranging from depth sensors such as the Kinect to photogrammetric systems. Li et al. [2013] present a system to capture 3D self portraits from a 3D sensor, based on non-rigid registration of several partial scans and Poisson texture blending for smooth colors from the input views. Sturm et al. [2013] propose a signed distance function that is updated at interactive rates directly from the sensor feed, obtaining dense 3D models. Tena et al. [2013] develop a semi-automated, commercial system to seamlessly integrate customer faces into figurines. However, all methods have problems when it comes to reconstructing

hair, either approximating the hair-style in low-resolution or replacing it completely with a pre-modeled template.

In addition to research efforts, the consumer market is exploding with products and services that offer physical reproduction of personal figurines [3DSystems 2013; 3D-u 2010; D-Tech Me 2012; Omote 3D 2012; PocketSize Me 2013]. However, these systems suffer from the same drawback that the hair-style is not sufficiently captured. Our work is complementary to these efforts, proposing the first approach to address this limitation and provide personalized hair-styles suitable for 3D reproduction.

Hair Reconstruction. Early work on reconstructing hair from images targeted simple hair-styles [Kong et al. 1997] or reconstructed only partial hair [Grabli et al. 2002]. More recently we have seen several advances in generating photo-realistic hair reconstructions [Paris et al. 2004; Wei et al. 2005; Paris et al. 2008; Jakob et al. 2009; Beeler et al. 2012; Herrera et al. 2012; Chai et al. 2013; Luo et al. 2013; Hu et al. 2014] and synthesis [Wang et al. 2009]. These methods aim to reconstruct individual strands of hair, which do not meet the physical reproduction constraint of a manifold surface. In contrast, we aim to create 3D-printable hair with enough geometric detail to convey the same *style* as the original. Luo et al. [2012] obtain a hair surface from multiple images, based on the observation that orientation fields are reasonably coherent across views. However, hair-style results tend to lack the level of stylization, detail and color required to produce appealing 3D-printed models. Finally, coarse hair reconstructions have been obtained from a single image for the application of advanced image editing [Chai et al. 2012], which differs significantly from our goal of stylized 3D hair capture.

Geometry Abstraction. Several works have recently appeared on the topic of abstraction of geometrical features for simplification of complex models [Yumer and Kara 2012; Nan et al. 2011; Mehra et al. 2009]. However, those focus on man-made shapes, which are very different from the geometry of hair-styles.

Image and Video Stylization. The field of non-photorealistic rendering is traditionally very active in studying the problem of stylizing and abstracting 2D images and videos (refer to Kyprianidis et al. [2013] for a detailed survey). Specifically related to hair stylization are methods that stylize images while preserving the directionality of the most prominent features [Kang et al. 2009; Kyprianidis and Kang 2011]. These techniques are designed for single images represented as 2D regular grids, and are thus unsuited for 3D hair stylization. Even in the context of multi-view styling, consistency of existing 2D stylization algorithms across views is not guaranteed. We draw inspiration from these 2D operators, and propose a novel extension to irregular 3D geometric domains and multi-view settings. This allows us to stylize hair while handling its view-dependent appearance and geometric complexity in a coherent way.

In summary, our work represents the first approach for reconstructing hair that is suitable for manufacturing personal figurines, yet stylized with subject-specific details that capture the identity of the individual.

3 Method Overview

Figure 3 shows a diagram of our pipeline. Our algorithm takes as input several color images of a person with a given hair-style from a multi-view capture system (Section 4). Using these images, an initial smooth and coarse proxy geometry is obtained using multi-view stereo [Beeler et al. 2010]. This proxy is then initialized with

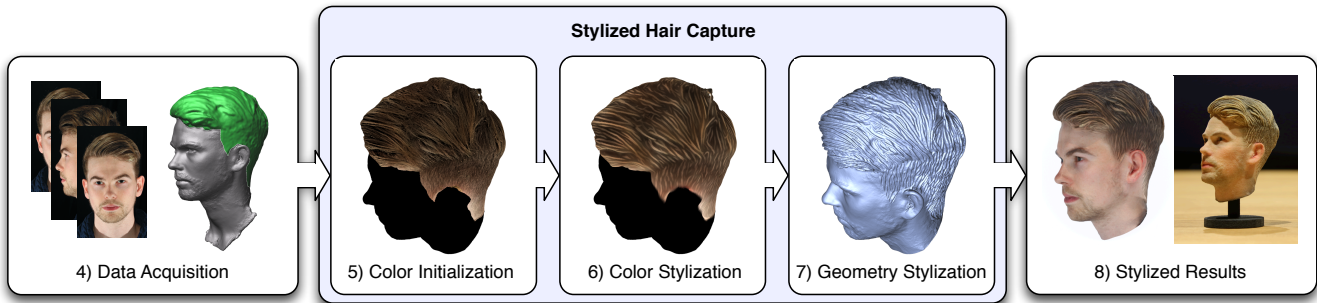


Figure 3: Our approach starts with a data acquisition process to construct a coarse proxy surface of the hair from multiple images. Stylized hair capture then begins by initializing, and then stylizing color information over the proxy surface using a novel stylization filter. We then extract detailed structure from the stylized colors in order to consistently stylize also the geometry. The result is a 3D-printable surface that captures the defining features of the hair-style.

per-vertex color information from the input images, split into multiple frequencies (Section 5). Next, stylization operators are applied over the color information to achieve the desired level of abstraction (Section 6). From the stylized color information, new geometric details for the proxy are synthesized giving the artist control over the geometric appearance (Section 7). The final output of our method is a stylized mesh which abstracts the complexity of a real hair-style while still preserving its defining features (Section 8). This mesh is suited for miniaturization and 3D printing (Figure 1, right).

4 Data Acquisition

As a domain for computing stylized effects, our algorithm requires a coarse, low-resolution geometric proxy surface of the hair. This proxy may be generated with any 3D capture system that provides geometry and images, such as Li et al. [2013]. In this work, we use a multi-view reconstruction setup as described in the following.

Ideally, the geometry proxy would be captured in a single shot from sufficient viewpoints to cover the full head. Since this approach requires a significant amount of camera and lighting equipment, we describe a technique that makes use of only limited hardware. We place ten digital SLR cameras in a quarter-spherical setup and photograph the subject under four consecutive orientations defined by 90-degree rotations (Figure 4).



Figure 4: Our capture setup consists of 10 cameras placed in a quarter-sphere and we capture 4 different orientations of the subject.

Although we take care not to change the hair-style or alter the facial expression drastically between the four takes, minor differ-




Figure 5: Partial reconstructions from four sequential orientations (left) are combined to form a proxy geometry (right) shown with the hair masked in green.

ences are tolerated since we require only a low-resolution proxy shape. We use a multi-view stereo reconstruction algorithm [Beeler et al. 2010] to compute partial reconstructions from each of the four orientations (Figure 5, left). These reconstructions are aligned rigidly through the Iterative Closest Points (ICP) algorithm [Besl and McKay 1992], and a single surface is obtained through Poisson reconstruction [Kazhdan et al. 2006] of the combined point cloud. This surface represents our geometric proxy, including both hair and face (Figure 5, right). We manually identify the hair region through simple masking. The proxy will serve as a base for synthesizing stylized details in both shape and color. The four rigid transformations computed from ICP are also applied to the calibrated camera views to produce virtual cameras surrounding the proxy. In total we could obtain 40 virtual cameras; however, while such a dense view sampling is advantageous for multi-view stereo, we found that a subset of eight views is sufficient for hair stylization once the geometry is reconstructed. To this end, we select the eight views that cover the hair volume from front, back, and both sides at two different elevations. In the rest of this paper we refer only to this subset of inputs.

5 Color Initialization

Once the proxy geometry has been obtained, the next step is to assign colors to its vertices. Contrary to face colors, which come from [Beeler et al. 2010], coloring the hair requires special treatment because the proxy surface only poorly approximates the volumetric nature of hair and is not geometrically accurate, so there will inevitably be inconsistencies between different views. Furthermore,




Figure 6: a) All-frequency colorization from the single-best view introduces strong seams but locally preserves sharpness. b) All-frequency colorization by averaging multiple views leads to blurry results, but reduces seams. c) Seams are most apparent in low-frequency bands, while they are masked in the high-frequency bands (shown in false color given they are offsets from the DoG). d) Loss of sharpness on the other hand is most apparent for high frequencies. e,f) We colorize the proxy by averaging low-frequency information from multiple views but sample high-frequency information from the single-best view, leveraging to the complementary nature of the two approaches. g) This approach might cause a change in contrast, which we equalize as described in Subsection 5.3. h) Contrast may be further boosted by scaling the intensity of the high-frequency band to help stylization. (Dataset courtesy of Luo et al. [2012]).

hair has a very complex appearance, with strong view dependent effects, such as specular reflection, translucency and occlusions. Obtaining a sharp and seamless colorization from multiple views is therefore extremely challenging. On the one hand, assigning colors per vertex based on a single-best viewpoint will lead to strong color seams (Figure 6 (a)). On the other hand, computing color by averaging multiple viewpoints will lead to blurry results (Figure 6 (b)). As shown in Figure 6 (c), color seams due to single-best viewpoint selection are most apparent in the lower frequencies, while they are masked in the high frequency bands. In contrast, blurriness due to color averaging has the biggest impact on the high-frequency components (Figure 6 (d)). Following from this observation, we separate color information into two frequency bands, low and high, using a Difference of Gaussians (DoG) filter. These color bands are processed separately according to their complementary nature and combined back together on the mesh (Figure 6 (e-f)).

Subsequent steps in the pipeline will make use of the intensity values to enhance contrast (Section 5.3) and guide geometric stylization (Section 7). To be able to directly operate on intensity values we convert the images from RGB to HSV color space as a preprocessing step.

5.1 Low-Frequency Color

For low-frequency color information, we assign a color \mathbf{c} to each vertex of the proxy mesh by averaging color samples \mathbf{c}^j from the set of views \mathcal{V} , weighted by their foreshortening angle ω^j

$$\mathbf{c} = \frac{1}{\sum_{j \in \mathcal{V}} \omega^j} \sum_{j \in \mathcal{V}} \omega^j \mathbf{c}^j, \quad (1)$$

with

$$\omega^j = \max(\langle \mathbf{n}, \mathbf{v}^j \rangle, 0),$$

$$\mathbf{c}^j = I^j(P^j \mathbf{p}),$$

where \mathbf{n} is the normal and \mathbf{p} is the position in world space of the vertex. I^j , \mathbf{v}^j and P^j are the low-frequency HSV image, view vector and projection matrix of view j respectively. Equation 1 is

applied to each HSV channel separately. This approach effectively removes visible color seams and attenuates view-dependent color changes (Figure 6 (d)).

5.2 High-Frequency Color

As mentioned previously, the proxy geometry is only a coarse approximation of the hair volume, and averaging high-frequency color from multiple views as described for the low-frequency components would lead to blurry results and exhibit ghosting (Figure 6 (d)). We therefore sample the high-frequency details only from the single-best view j^* , which we consider to be the one with the highest foreshortening angle ω^{j^*} (Figure 6 (e)).

5.3 Contrast Equalization

Local contrast or dynamic range is directly related to incident light intensity. When averaging the low-frequency color from multiple views, the resulting intensity l might be very different from the low frequency intensity l^{j^*} in the single-best view used to extract the high-frequency component h^{j^*} due to view dependent effects; this results in a perceived loss of contrast when frequencies are recombined (Figure 6 (f)). To alleviate this problem, we perceptually adjust the intensity value of the high frequency components h as

$$h = \xi h^{j^*} = \frac{l}{l^{j^*}} h^{j^*}. \quad (2)$$

In the case that the combined l is darker than l^{j^*} we use $1/\xi$ instead, since the goal here is to increase the contrast.

Figure 6 (g) shows the effect of equalizing the contrast when compared to Figure 6 (f). The impact is most apparent in areas around the seams, where the averaged low-frequency intensity l differs substantially from the intensity found in the single-best view l^{j^*} .

Finally, contrast may be boosted globally by uniformly scaling h to help color stylization (Figure 6 (h)).

6 Color Stylization

As motivated in the introduction, our intent is to create a representation of the hair that can be miniaturized and printed. However, the level of detail of real hair is overwhelmingly high, and we thus need to find a means to reduce the complexity while preserving its defining features. We approach this problem by employing a specialized stylization filter.

Recently, Kyriandis and Kang [2011] have proposed an anisotropic, feature-preserving stylization filter for images, which exhibits interesting color stylization effects in hair regions. Their 2D filter can successfully reduce hair complexity, while still maintaining the overall appearance of a hair-style. Unfortunately, their method is not immediately suited for stylizing hair in 3D, since (i) it requires a regular 2D-domain, while our hair proxy constitutes an irregular manifold embedded in 3-space; and (ii) it is designed for single-image processing, while we operate within a multi-view scenario, where consistency between different views is essential. Nevertheless, the 2D filter of Kyriandis and Kang contains several of the properties we desire in our application of stylized hair capture, and thus we propose to employ a similar mathematical foundation in a novel multi-view stylization algorithm that operates on irregular manifolds in 3D. Our approach will be to couple 2D and 3D stylization. Therefore, we will first provide an overview of the foundations for feature-preserving stylization in 2D (as presented by Kyriandis and Kang), and then proceed to describe the key extensions that enable color stylization of hair in 3D.

6.1 Feature-Preserving 2D Filter

Feature-preserving or enhancing directional filters are based on three main components [Kang et al. 2009; Kang and Lee 2008]: The estimation of the local structure, an integration/smoothing operation to reduce complexity, and a sharpening operator to enhance the desired features. In the following, we describe how these components are achieved, focusing first on a single image.

Local Structure Estimation. Estimating the directionality of features in a 2D image has been studied before, with some specific applications to hair [Paris et al. 2004; Kang and Lee 2008; Luo et al. 2012]. Given its fast computation and proven effectiveness in stylization contexts, we first compute the structure tensor $\mathcal{S}(\mathbf{x})$ [Brox et al. 2006] at each pixel \mathbf{x} . To attenuate noise, we filter the structure tensor using an isotropic 2D Gaussian [Kyprianidis and Kang 2011] with standard deviation $\sigma_d = 4$.

Based on the structure tensor we introduce the *orientation tensor* $\mathcal{O}(\mathbf{x}) \in \mathbf{R}^{2 \times 2}$, which consists of the two eigenvectors of \mathcal{S} scaled by their respective eigenvalues. For our purposes $\mathcal{O}(\mathbf{x})$ is equivalent to $\mathcal{S}(\mathbf{x})$ but has some advantages when extending to multiple views, as we will see in Section 6.3. For ease of notation, we will omit the pixel (\mathbf{x}) in the orientation tensor and refer simply to \mathcal{O} in the rest of this discussion. The first row of the orientation tensor \mathcal{O}_g , corresponds to the image gradient, while the second row \mathcal{O}_t coincides with the tangent. The gradients and tangents will be used for directional smoothing and directional sharpening as described in the following. We will use $\vec{\mathcal{O}}_{g,t}$ to indicate the normalized gradient and tangent directions, respectively.

Directional Smoothing/Integration. Once we have a smooth orientation field, stylization starts by performing directional smoothing following the tangents $\vec{\mathcal{O}}_t$. This operation consists of a line integral convolution [Cabral and Leedom 1993]. To obtain the resulting color $I'(\mathbf{x}_0)$ of a given pixel \mathbf{x}_0 , additional color samples $I(\mathbf{x}_k^\pm)$ are interpolated from both directions of the flow, and then averaged using a 1D Gaussian function $G(d_k)$ of standard deviation $\tilde{\sigma}_t$, where d_k represents the geodesic distance from the sampling point \mathbf{x}_k to \mathbf{x}_0 . We aim to filter with a standard deviation $\sigma_t = 10$, and account for local anisotropy by computing $\tilde{\sigma}_t$ as

$$\tilde{\sigma}_t = \frac{1}{4} \sigma_t \left(1 + \frac{\|\mathcal{O}_g\| - \|\mathcal{O}_t\|}{\|\mathcal{O}_g\| + \|\mathcal{O}_t\|} \right)^2. \quad (3)$$

We then have:

$$I'(x_0) = \frac{1}{w} \left[G(0)I(\mathbf{x}_0) + \sum_{k=1}^l G(d_k) [I(\mathbf{x}_k^+) + I(\mathbf{x}_k^-)] \right], \quad (4)$$

where $I(\mathbf{x})$ is the original color of a pixel \mathbf{x} , $w = \sum_{k=-l}^l G(d_k)$ is a normalization factor, and $l = \lceil 2\tilde{\sigma}_t \rceil$ represents the cut-off of $G(d_k)$. Equation 4 is applied to each HSV channel independently. Several ways to compute $I(\mathbf{x}_k^\pm)$ have been proposed [Kang et al. 2009; Kyprianidis and Kang 2011]; we have found that a second-order Runge-Kutta integration scheme provides the best results in our context. Figure 7 (center) shows the result after this step for an image patch of hair (left).

Directional Sharpening. To further stylize certain features, a directional *shock filter* can be applied in the direction of the gradient $\vec{\mathcal{O}}_g$ [Osher and Rudin 1990; Kang and Lee 2008]. This filter



Figure 7: A schematic visualization of the smoothing and shocking operations. Left: Input image. Center: Line integral convolution following the tangents (in blue). Orange dots show integrated samples $I(\mathbf{x}_k^\pm)$ for pixel \mathbf{x}_0 (green). Right: Application of the shock operator in the direction of the gradients (shown in red).

consists of morphological operations of dilation and erosion, which create ruptures between local maxima and minima, while enhancing flow-like patterns in the image [Weickert 2003]. For efficiency, it can be approximated by a min/max filter applied over a neighborhood of radius r , depending on the sign of a luminance-based Laplacian of Gaussian (LoG) with standard deviation σ_g [Kyprianidis and Kang 2011]. In our case, we apply the LoG over the value channel in HSV space. We use $r = \sigma_g = 3$ for the results shown in this paper. Figure 7 (right) shows the final result, after both directional smoothing and sharpening.

6.2 Extension to Irregular Manifolds in 3-Space

This section describes how the stylization filter defined on a regular 2D image domain can be extended to operate on an irregular manifold embedded in 3-space. Color processing over manifolds has been studied before for applications like de-noising [Sochen et al. 2003]. Here we extend a different set of operators that additionally deal with per-vertex information obtained from several input images, which need to be processed in a coherent way.

A possible alternative would be to compute a UV-parameterization that maps the manifold into 2-space, and then use the 2D stylization filter. However, we discarded such an approach for two reasons. First, the mesh proxy is far from being a developable surface, which would cause distortions in the mapping. These distortions would vary spatially and so the stylization operators would have to take them into account. Second, the proxy may have arbitrary genus, which can occur due to curls or ponytails (for example, see Figure 12), which would cause a discontinuous parameterization and require cutting the mesh - a challenging problem on its own. Furthermore, the stylization operators would have to be adapted to handle the discontinuity caused by such cuts. While these approaches can be found in previous work on texture synthesis [Lefebvre and Hoppe 2006; Wei and Levoy 2001], we avoid these challenges by operating directly in 3-space.

As described in Section 5, we compute and store color information directly on the mesh, and will do the same for orientation tensors. Operations such as directional smoothing and shock filtering are then performed on a per-vertex basis, using the geodesic distance on the mesh surface. To do so, analogous to 2D directional smoothing, we require to look up the color value along a direction \mathbf{t} at a geodesic distance δ from the current vertex \mathbf{x} on the mesh. Since the mesh is not planar, the point $\mathbf{x}'_k = \mathbf{x} + \delta\mathbf{t}$ might be off surface and needs to be projected back on to get \mathbf{x}_k (Figure 8 (a)). If δ is larger than the local tessellation of the mesh, then \mathbf{x}_k will not reside in the one-ring neighborhood of \mathbf{x} and the process is repeated recursively using the closest point \mathbf{x}_e on the edge of the one-ring as new starting point and subtracting $\|\mathbf{x}_e - \mathbf{x}\|$ from $\delta\mathbf{t}$, effectively rolling

down the original vector $\delta\mathbf{t}$ onto the mesh. Figure 8, b shows this roll-down schematically. Once \mathbf{x}_k is computed, the value is interpolated using barycentric coordinates.

The important parameter in this process is the step size δ , which should be similar to the local vertex density to avoid sampling issues. Fortunately, our meshes have very uniform vertex density allowing us to use a global δ that corresponds to the average edge length of the mesh.




Figure 8: a) A point \mathbf{x}'_k in the direction of the tangent \mathbf{t} at distance δ from \mathbf{x} may be off surface, and thus needs to be projected back on (\mathbf{x}_k) in order to sample interpolated vertex information. b) If the sampling distance δ (blue) is larger than the local tessellation, the sampling point \mathbf{x}_k will reside outside of the direct neighbourhood of \mathbf{x} . In this case the process is repeated recursively using the closest point \mathbf{x}_e on the edge of the one-ring as new starting point \mathbf{x} (green) and adjusting δ accordingly, effectively rolling down the original vector $\delta\mathbf{t}$ onto the mesh.

6.3 Extension to Multiple Views

As motivated in Section 5, information from different views will be naturally misaligned since the proxy geometry is an approximation of the hair volume and the appearance of hair may differ substantially in different views. Consequently we merge low-frequency color information by averaging the contributions of the individual views to avoid color seams (Section 5.1), but we must combine high-frequency components using the single-best view to avoid blurring (Section 5.2). As a result, the high-frequency color information on the mesh will contain seams, and this can adversely impact orientation tensor computation if performed directly on the mesh. This is demonstrated in Figure 9 on a synthetic mesh patch. The high-frequency components are sampled from three different views (color-coded for better visualization) using the single-best view per vertex (Figure 9 (a)). The seams are clearly visible due to the change in orientation of the features. If we compute the orientation tensor on the mesh, it will inevitably follow such seams and so will do the directional smoothing (Figure 9 (b) and (f)). As a consequence, we will compute the orientation tensors from the images and transfer them to the mesh vertices.

Since the orientation tensors capture the continuous style of the hair, their spatial variation is essentially low-frequency, even though they are computed from the high-frequency components of the images. Analogously to vertex colorization discussed in Section 5, computing per-vertex orientation tensors from the single-best view would lead to strong seams (Figure 9 (c) and (g)), and thus we need to combine the tensors from multiple views using an averaging scheme instead (Figure 9 (d)). In the following section we will explain how to combine them to produce a smooth tensor field on the mesh.

Orientation Tensor Backprojection. Similar to Paris et al. [2008], transferring an orientation tensor \mathcal{O}^j from a view j to a point \mathbf{x} in space happens via backprojection, although our proxy geometry makes the process more straightforward. We project the normalized gradient direction $\vec{\mathcal{O}}_g$ onto the tangent plane at \mathbf{x} . The tangent direction $\vec{\mathcal{O}}_t$ is computed as the orthogonal vector

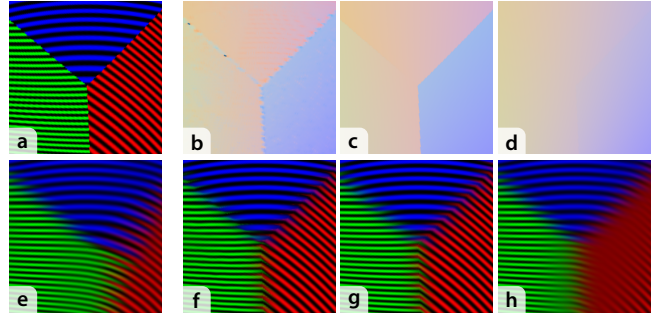


Figure 9: Synthetic example with high-frequency values from different views (color-coded for better visualization). a) Using only the single-best view will cause sharp transitions (blue-red, green-red). b) These transitions would adversely impact orientation tensor computation (colors encode the dominant orientation). c) Computing and transferring the tensor from the single-best view avoids these artifacts, but exhibits the same discontinuities. d) Using the proposed tensor combination yields a smooth tensor field. e, f, g, h) Tensor fields (b, c, d) applied to (a). The orientation is only continuous in (h) but the spatial frequencies are attenuated around seams as the colors don't match the tensor anymore. e) The final gathering step re-establishes this consistency and smoothly combines the structural elements from the different views without degradation.

in the tangent plane. Both vectors are re-scaled to their original magnitudes $\|\mathcal{O}_{g,t}^j\|$. Note that the sign of these vectors is arbitrary. We thus ensure consistency of the direction vectors of all views by reversing vectors that do not agree with the orientation of the single-best view. In the following, \mathcal{O} will denote this backprojected orientation tensor.

Orientation Tensor Combination. Unlike structure tensors, which are oriented absolutely with respect to each view, orientation tensors from different views can be directly combined since they are oriented with the gradients. The combined gradient direction $\vec{\mathcal{O}}_g$ is thus computed as

$$\vec{\mathcal{O}}_g = \frac{1}{w} \sum_{j \in \mathcal{V}} \omega^j \theta^j \vec{\mathcal{O}}_g^j, \quad (5)$$

where \mathcal{V} is the set of views and $w = \sum_{j \in \mathcal{V}} \omega^j \theta^j$ is the normalization factor. The contributions are weighted using foreshortening ω as defined in 5.1 as well as the misalignment θ , which is computed as the discrepancy between the orientation tensors of the view j and the single-best view j^*

$$\theta^j = \langle \vec{\mathcal{O}}_g^j, \vec{\mathcal{O}}_g^{j^*} \rangle. \quad (6)$$

Given the gradient direction $\vec{\mathcal{O}}_g$ on the mesh, the tangent direction $\vec{\mathcal{O}}_t$ can again be derived since they are orthogonal.

Finally, the magnitudes of the vectors are computed as

$$\|\mathcal{O}_{g,t}\| = \frac{1}{w} \sum_{j \in \mathcal{V}} \omega^j \theta^j \|\mathcal{O}_{g,t}^j\|, \quad (7)$$

using the same weights and normalization as in Equation 5.




Figure 10: The proposed color stylization pipeline couples 2D image (bottom row) with 3D mesh stylization (top row). The final output is the mesh-proxy with stylized per vertex colors.

Directional Gathering And Smoothing. In the previous section we described how to combine the orientation tensors from the different views to produce a continuous orientation tensor field on the mesh (Figure 9 (d)). Unfortunately, since the high-frequency color information has been computed from the single-best view and the orientation tensors are produced by averaging multiple views, the two will be inconsistent around seams. This discrepancy will cause the directional smoothing to filter across color boundaries in these areas and produces blurry results (Figure 9 (h)). To overcome this problem, we propose an approach to update the color information, using a measure of the discrepancy between the gradient computed from the mesh colors $\vec{\mathcal{O}}'_g$ and the gradient from the combined multi-view orientation tensor $\vec{\mathcal{O}}_g$ as

$$\beta = 1 - \langle \vec{\mathcal{O}}_g, \vec{\mathcal{O}}'_g \rangle. \quad (8)$$

Given this measure, we introduce a *gathering* step, which ensures that color is consistent with the combined orientation tensor field. For a vertex with discrepancy β larger than a given threshold τ , we search along the tangent field in both directions for color samples with low discrepancies, using the same second-order Runge-Kutta scheme as for directional smoothing. The vertex color \mathbf{c} is then updated by linearly interpolating between those color samples based on their geodesic distance from the current vertex. The discrepancy β is then updated in the same way. Overall, discrepancy is usually lower than 0.1 in our examples, and we found a value of $\tau = 0.02$ produces sufficient consistency. The updated discrepancy β is further used as weighting factor to attenuate these samples during directional smoothing.

The final result is a smooth orientation tensor field with consistent coloring, the required basis for our stylizations (Figure 9 (e)) as it preserves the important structural elements from the different viewpoints while combining them together in a continuous way.

6.4 Coupled Mesh-View Stylization

Now that we have explained all the components required for multi-view hair stylization on a manifold in 3D, we present the complete stylization pipeline, as shown conceptually in Figure 10. It is important to see that the final result is based on coupled stylization of both the mesh and the input views. This is required to keep the orientation tensor, which is computed from the views, consistent with the colors on the mesh, which are stylized using this tensor.

The degree of stylization is emphasized by iteratively re-applying the method, which allows for direct artistic control. We use two to three iterations for the results presented in this paper. After stylization, an antialiasing step is employed to refine the color transitions. The final output is the mesh-proxy with stylized per vertex colors.

7 Geometry Stylization

Up to this point, we have stylized high-frequency information in HSV color space over the mesh. We wish to also stylize the geometric details of the hair, such that they are consistent with the color style. To this end, we will compute spatially-varying surface offsets $d(\mathbf{x})$, and displace the vertices of the proxy geometry along the normal direction $\mathbf{n}(\mathbf{x})$ by $d(\mathbf{x})\mathbf{n}(\mathbf{x})$.

Most of the perceived high-frequency contrast is encoded in the value channel while hue and saturation vary less substantially in general. This contrast can be largely attributed to shading changes caused by hair geometry, which is why we look to the stylized color intensity to determine the structural offsets. This is the underlying motivation for operating in HSV color space. Converting the high-frequency stylized intensity offsets $v(\mathbf{x})$ defined in the range of $[-1, 1]$ to displacements $d(\mathbf{x})$ is essentially a tone mapping application

$$d(\mathbf{x}) = \varphi \Phi(v(\mathbf{x})), \quad (9)$$

where φ is an artistic parameter that controls the strength of the shape stylization. For the results in this paper, φ varies between 3 and 6, and as tone mapping operator Φ , we apply a simple gamma correction with γ ranging between 0.5 and 0.7.

Wisp Profiles. The shock operation in color stylization has the property of creating uniform intensities (and thus uniform displacement values) for all the vertices within a shock, in the direction of the gradient (Figure 11, a), producing wisps with a locally flat appearance (Figure 11, c). While this is often sufficient to create successful hair stylizations, additional geometric styles can be obtained by modulating the displacements artistically.

Our idea is to find the relative position of each vertex within the cross-section of a wisp, and use this position to modulate displacement according to a user-defined *wisp profile*. To this end, for each vertex we search for discontinuities in intensity following the direction of the gradient $\vec{\mathcal{O}}_g$ (as done previously for the shock filter) by checking the relative change in intensity with respect to the starting vertex. Discontinuities in each direction mark the cross-sectional boundaries $b_{1,2}$ of the wisp (Figure 11, a). We then apply a 1D-window function to the geometry offsets to control its profile. In this paper, we employ three different wisp profile functions - box, sinusoidal and custom - which produce different artistic looks (Figure 11, b-d). Defining additional profile functions is straightforward, providing a simple way to artistically control the style of the hair. To reduce aliasing, the modulated displacements can be smoothed in the direction of the tangents using the same directional smoothing operator explained in Section 6.




Figure 11: a) As a consequence of the shock filter, vertices along the gradient will have the same color within the boundaries $b_{1,2}$ of a wisp. This enables to compute the relative position x of a vertex \mathbf{x} within the cross-section and to apply different wisp profile functions (b,c,d) to artistically alter the geometric appearance.

All-Frequency Color and Geometric Information. After applying the stylized high-frequency displacements, low- and high-frequency color information is combined and converted from HSV to RGB, obtaining the final colored and sculpted result, which is suitable for miniaturization and physical reproduction using traditional 3D printing processes.

8 Results

We demonstrate the robustness and versatility of our method by capturing several subjects with widely varying hair-styles, and show that the proposed technique fits naturally in traditional pipelines for personalized figurine creation by physically reproducing several figurines with stylized hair.

A primary goal of our algorithm is to reconstruct an individual’s hair in a way that captures the identifying characteristics of the hair-style. Not only should the result be recognizable as a replica of a particular person, but it should be clearly recognized as a reconstruction of that person with their hair in a particular style. Our technique successfully achieves this goal, which we demonstrate by capturing the same two subjects each with four different hair-styles, as shown in Figure 12. Each hair-style is clearly recognizable in both shape and color, even though it is captured in a stylized and 3D-printable way.

We show the robustness of our approach in Figure 13, by capturing several different subjects with a wide range of hair-styles and hair colors. Our method can even stylize the reconstruction of facial hair, as seen for the subjects in the bottom two rows.

In Figure 14 we further demonstrate the versatility of the proposed technique, by capturing stylized reconstructions of non-human hair. In these examples, we capture a fur collar and a stuffed toy dog. Despite their complexity, both reconstructions are stylized manifold surfaces, suitable for 3D printing.

In this paper we describe a data acquisition step (Section 4) to build a low-resolution proxy surface. However, our hair stylization method can be applied to any proxy geometry with available camera views. This is demonstrated in Figure 15, where we apply stylization to a dataset provided by Luo et al. [2012]. As can be seen in the zoom regions, the benefit of our approach over Luo et al. [2012] is that we capture the important structural and color elements of the hair-style, even starting from a proxy that only coarsely approximates the hair.

Different levels of stylization can be achieved by changing the shock filter width, as shown in Figure 17 (a-c). These styles can even be combined in any artistic way, which we show in Figure 17 (d). As mentioned, our technique can be applied naturally in physical reproduction pipelines for manufacturing personalized figurines. We miniaturized and 3D-printed several of our results, as shown in Figure 1 (right), Figure 3 (rightmost), and Figure 18. Physical printouts were created using a *ZCorp* printer. Our method can produce smaller figurines or handle lower-resolution printers by creating coarser hair wisps, as seen in Figure 17 (c).

For a better assessment of our contribution to capturing and stylizing hair, we illustrate the intermediate steps of our pipeline in Figure 16, for a subset of our datasets. This figure shows the proxy mesh generated by Beeler et al. [2010], the proxy with initialized colors, both our color and geometry stylization results on their own, and the final combined result of our algorithm. Our method increases the color and structural elements of the hair-style substantially compared to the initial proxy reconstruction.

We computed our results on a desktop PC with a six-core Intel 3930K CPU. It took our prototype implementation 4-7 hours to pro-



Figure 13: Our stylized hair capture technique can reconstruct a wide range of hair-styles, including the facial hair on the subjects in the bottom two rows. Here we show one of the input images (left), the stylized geometry without color (center), and the final result with stylized shape and color (right).

cess meshes between 300K and 600K hair vertices; however, we believe this could be substantially improved using optimized data structures and by exploiting the extremely parallelizable nature of the problem on the GPU (following Kyprianidis et al. [2011]). The parameters given in Subsection 6.1 were determined for input images of 2592 x 1728 pixels and transferred to the mesh domain for the 3D filter via re-projection onto the proxy geometry.

Discussion. In our implementation, the total processing time and the overall quality of the results depend on the mesh discretization. It is preferable to operate on uniform triangulations with reasonable resolution. In particular, we found that between 300K and 600K



Figure 12: Here we show four different captured hair-styles for two people (left and right). For each style we show the final result with both color and geometry stylization, as well as the result without color in order to best visualize the geometry stylization. Our method is able to faithfully capture the essence and identifying characteristics of each hair-style, even when the subject is the same.

hair vertices are sufficient to obtain sharp color and geometry details. Most multi-view capture techniques require many viewpoints when reconstructing complex objects with strong self-occlusions, which is why we used 40 viewpoints for the initial proxy reconstruction. For stylization, however, we found that a subset of about 8 views is sufficient. We expect the global shape of the hair to be provided by the proxy geometry, which might fail to reproduce some of the finer structures such as individual curls. Identifying and addressing such structures would be an interesting area for future research.

In conclusion, we present the first method for *stylized* hair capture, a technique to reconstruct an individual's actual hair-style in a manner suitable for physical reproduction. Our method generates hair

as a closed-manifold surface, yet contains the structural and color elements stylized in a way that captures the defining characteristics of the hair-style.

Acknowledgements

We wish to thank Linjie Luo for sharing his dataset. Thanks to all our hair models, and to Maurizio Nitti and Alessia Marra for artistic help throughout the project. This publication has been partially funded by the European Commission, Seventh Framework Programme, through projects GOLEM (Marie Curie IAPP, grant: 251415) and VERVE (ICT, grant: 288914).



Figure 14: Our approach is versatile, and can be used for stylized capture of more than just human hair. Here we show additional examples of a stuffed toy dog (top) and fur (bottom).




Figure 15: Although we do not share the same final goal, we show our method applied over a dataset from Luo et al. [2012]. From left to right: one of the input views. Initial coarse reconstruction. Reconstructed results by Luo et al. Our stylized results starting from the same low-resolution proxy. In addition to sculpted geometry, our method generates also stylized color information. Our detailed stylization improvements over Luo et al. are clearly visible in the zoomed regions.

References

- 3D-U, 2010. ThreeDee-You. <http://www.3d-u.es/>. Accessed: January 18, 2014.
- 3DSYSTEMS, 2013. 3DMe Photobooth. <http://www.3dsystems.com/>. Accessed: January 10, 2014.
- BEELER, T., BICKEL, B., BEARDSLEY, P., SUMNER, B., AND GROSS, M. 2010. High-quality single-shot capture of facial geometry. *ACM Trans. Graph.* 29, 3, 40:1–40:9.
- BEELER, T., BICKEL, B., NORIS, G., BEARDSLEY, P., MARSCHNER, S., SUMNER, R. W., AND GROSS, M. 2012. Coupled 3D reconstruction of sparse facial hair and skin. *ACM Trans. Graph.* 31, 4, 117:1–117:10.
- BESL, P. J., AND MCKAY, N. D. 1992. A method for registration of 3-D shapes. *IEEE Trans. Pattern Anal. Mach. Intell.* 14, 2, 239–256.




Figure 16: We illustrate the intermediate steps of our pipeline on two different datasets. From left to right: the input proxy geometry, color initialization on the proxy, color stylization alone, geometry stylization alone, and the final result with color and geometry stylization.




Figure 17: Here we show different levels of stylization (a-c) and a combination of styles (d) computed as $0.3(a) + 0.4(b) + 0.4(c)$.



Figure 18: Our method fits naturally in the physical reproduction pipeline for personalized figurines. A stylized capture result was printed with and without color.

- BROX, T., BOOMGAARD, R., LAUZE, F., WEIJER, J., WEICKERT, J., MRZEK, P., AND KORNPROBST, P. 2006. Adaptive structure tensors and their applications. In *Visualization and Processing of Tensor Fields*, Mathematics and Visualization. 17–47.
- CABRAL, B., AND LEEDOM, L. C. 1993. Imaging vector fields using line integral convolution. In *Proceedings of the 20th Annual*

- Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '93*, 263–270.
- CHAI, M., WANG, L., WENG, Y., YU, Y., GUO, B., AND ZHOU, K. 2012. Single-view hair modeling for portrait manipulation. *ACM Trans. Graph.* 31, 4, 116:1–116:8.
- CHAI, M., WANG, L., WENG, Y., JIN, X., AND ZHOU, K. 2013. Dynamic hair manipulation in images and videos. *ACM Trans. Graph.* 32, 4, 75:1–75:8.
- D-TECH ME, 2012. D-tech me. <http://www.insidethemagic.net/tag/d-tech-me/>. Accessed: January 18, 2014.
- GRABLI, S., SILLION, F., MARSCHNER, S. R., AND LENGYEL, J. E. 2002. Image-based hair capture by inverse lighting. In *Proc. Graphics Interface*, 51–58.
- HERRERA, T. L., ZINKE, A., AND WEBER, A. 2012. Lighting hair from the inside: A thermal approach to hair reconstruction. *ACM Trans. Graph.* 31, 6, 146:1–146:9.
- HU, L., MA, C., LUO, L., AND LI, H. 2014. Robust hair capture using simulated examples. *ACM Trans. Graph.* 33, 4.
- JAKOB, W., MOON, J. T., AND MARSCHNER, S. 2009. Capturing hair assemblies fiber by fiber. *ACM Trans. Graph.* 28, 5, 164:1–164:9.
- KANG, H., AND LEE, S. 2008. Shape-simplifying image abstraction. *Computer Graphics Forum* 27, 7, 1773–1780.
- KANG, H., LEE, S., AND CHUI, C. K. 2009. Flow-based image abstraction. *IEEE Transactions on Visualization and Computer Graphics* 15, 1, 62–76.
- KAZHDAN, M., BOLITHO, M., AND HOPPE, H. 2006. Poisson surface reconstruction. In *Symposium on Geometry Processing*, 61–70.
- KONG, W., TAKAHASHI, H., AND NAKAJIMA, M. 1997. Generation of 3D hair model from multiple pictures. In *Proc. Multimedia Modeling*, 183–196.
- KYPRIANIDIS, J. E., AND KANG, H. 2011. Image and video abstraction by coherence-enhancing filtering. *Computer Graphics Forum* 30, 2, 593–602.
- KYPRIANIDIS, J. E., COLLOMOSSE, J., WANG, T., AND ISENBURG, T. 2013. State of the Art: A taxonomy of artistic stylization techniques for images and video. *IEEE Transactions on Visualization and Computer Graphics* 19, 5, 866–885.
- LEFEBVRE, S., AND HOPPE, H. 2006. Appearance-space texture synthesis. *ACM Trans. Graph.* 25, 3, 541–548.
- LI, H., VOUGA, E., GUDYM, A., LUO, L., BARRON, J. T., AND GUSEV, G. 2013. 3D Self-portraits. *ACM Trans. Graph.* 32, 6, 187:1–187:9.
- LUO, L., LI, H., SYLVAIN, WEISE, T., PAULY, M., AND RUSINKIEWICZ, S. 2012. Multi-view hair capture using orientation fields. In *Proceedings of the 2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1490–1497.
- LUO, L., LI, H., AND RUSINKIEWICZ, S. 2013. Structure-aware hair capture. *ACM Trans. Graph.* 32, 4, 76:1–76:12.
- MEHRA, R., ZHOU, Q., LONG, J., SHEFFER, A., GOOCH, A., AND MITRA, N. J. 2009. Abstraction of man-made shapes. *ACM Trans. Graph.* 28, 5, 137:1–137:10.
- NAN, L., SHARF, A., XIE, K., WONG, T.-T., DEUSSEN, O., COHEN-OR, D., AND CHEN, B. 2011. Conjoining gestalt rules for abstraction of architectural drawings. *ACM Trans. Graph.* 30, 6, 185:1–185:10.
- OMOTE 3D, 2012. Omote 3D. <http://www.omote3d.com/>. Accessed: January 18, 2014.
- OSHER, S., AND RUDIN, L. I. 1990. Feature-oriented image enhancement using shock filters. *SIAM Journal on Numerical Analysis* 27, 4, 919–940.
- PARIS, S., BRICEÑO, H. M., AND SILLION, F. X. 2004. Capture of hair geometry from multiple images. *ACM Trans. Graph.* 23, 3, 712–719.
- PARIS, S., CHANG, W., KOZHUSHNYAN, O. I., JAROSZ, W., MATUSIK, W., ZWICKER, M., AND DURAND, F. 2008. Hair photobooth: Geometric and photometric acquisition of real hairstyles. *ACM Trans. Graph.* 27, 3, 30:1–30:9.
- POCKETSIZE ME, 2013. PocketSize Me. <http://www.pocketsizeme.ch/>. Accessed: January 18, 2014.
- SOCHEN, N., DERICHE, R., AND PEREZ, L. 2003. The beltrami flow over implicit manifolds. In *Ninth IEEE International Conference on Computer Vision (ICCV)*, vol. 3, 832–839.
- STURM, J., BYLOW, E., KAHL, F., AND CREMERS, D. 2013. CopyMe3D: Scanning and printing persons in 3D. In *Pattern Recognition*, vol. 8142, 405–414.
- TENA, J. R., MAHLER, M., BEELER, T., GROSSE, M., YEH, H., AND MATTHEWS, I. 2013. Fabricating 3D figurines with personalized faces. *IEEE Computer Graphics and Applications* 33, 6, 36–46.
- WANG, L., YU, Y., ZHOU, K., AND GUO, B. 2009. Example-based hair geometry synthesis. *ACM Trans. Graph.* 28, 3, 56:1–56:9.
- WEI, L.-Y., AND LEVOY, M. 2001. Texture synthesis over arbitrary manifold surfaces. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '01*, 355–360.
- WEI, Y., OFEK, E., QUAN, L., AND SHUM, H.-Y. 2005. Modeling hair from multiple views. *ACM Trans. Graph.* 24, 3, 816–820.
- WEICKERT, J. 2003. Coherence-enhancing shock filters. In *Pattern Recognition*, vol. 2781 of *Lecture Notes in Computer Science*. 1–8.
- YUMER, M. E., AND KARA, L. B. 2012. Co-abstraction of shape collections. *ACM Trans. Graph.* 31, 6, 166:1–166:11.