

Artistic Relighting of Paintings and Drawings

Bernardo Henz · Manuel M. Oliveira

Abstract We present a practical solution to the problem of subject relighting in paintings and drawings. Our interactive technique uses 3-D shading proxies and can be applied to objects with arbitrary geometries. Given a user-provided guess for the shading of an object in a painting/drawing and its corresponding target shading, we refine them using shading-color correlation and a multi-scale scheme. These refined shadings are then used to create a multi-channel shading-ratio image to perform relighting, while taking into account the colors used by the artists to convey shading information. We demonstrate the effectiveness of our solution on a variety of artistic styles, including paintings with strong brush strokes and unconventional shading encodings, drawings, and other types of artwork. Our method is the first to perform relighting of paintings and drawings and, in addition to relighting, can transfer smooth normal and depth maps from 3-D proxies to images.

Keywords Image Relighting; Painting Relighting; Drawing Relighting · Normal and Depth Map Transfer

1 Introduction

Image relighting tries to recreate the appearance of a pictured object or scene under new illumination. This is useful, for instance, when one would like to obtain

Bernardo Henz
Instituto de Informática, UFRGS
Porto Alegre - RS, Brazil
Email: bhenz@inf.ufrgs.br

Manuel M. Oliveira
Instituto de Informática, UFRGS
Porto Alegre - RS, Brazil
Email: oliveira@inf.ufrgs.br

different lighting effects, or direct attention to other areas of the image, but the picture cannot be (easily) re-taken or repainted under the desired conditions. This is a difficult problem, as a single image carries no explicit information about the scene's original lighting, or the objects' shapes and material properties. The problem becomes even harder in the case of paintings and drawings, as a combination of style, brush strokes, and use of colors often leads to *physically-inconsistent shadings*.

We present a practical solution to the problem of subject relighting in paintings and drawings. Image relighting requires information about the scene geometry, reflectance properties, and lighting. Instead of automatically recovering geometry using computer vision techniques [5, 10, 46] (which tend to not work well with such images), we use a 3-D model as a proxy to the subject's geometry. By illuminating the proxy with a guess to the original lighting conditions, the user provides a rough approximation to the *source shading*, which we call the *guess shading*. Since estimating source shading is a hard task even for experienced individuals, we present a technique to refine the user-provided source-shading. This is achieved by correlating shading values with colors in the input image. The proxy is also used to specify the *target shading*, which is refined using a multiscale scheme. An *RGB shading-ratio image* is then computed using the refined target and source shadings, and used for relighting. Such ratio image takes into account the artist's use of colors to encode shadings.

Figure 1 illustrates the flexibility of our approach to relight paintings and drawings. On top, it shows a van Gogh's painting with three relit versions: from the left, from the front, and from the top right. This is a challenging example, as the artist used colors to convey shading information (proportionally more red to indicate dark, and yellow for lit regions). The bottom

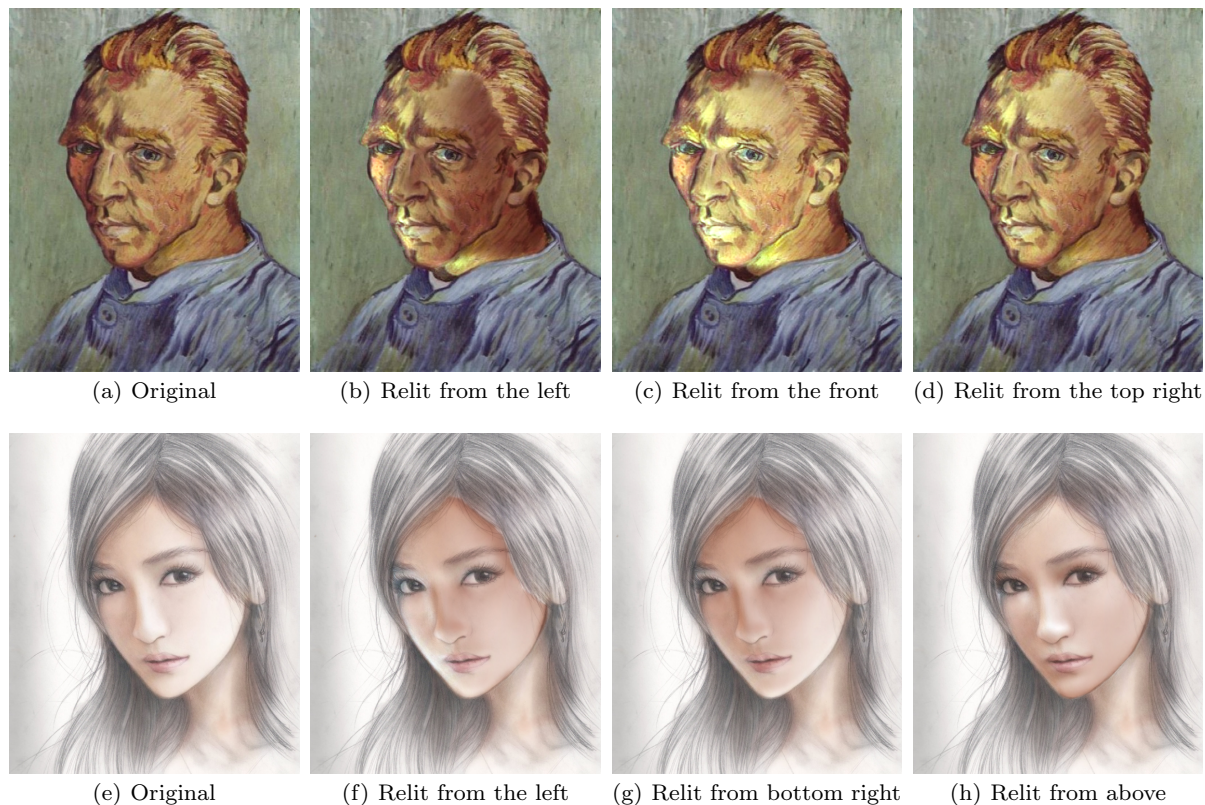


Fig. 1 Relighting of a painting and a pencil drawing from different directions using our technique. (top) van Gogh’s “Portrait de l’Artiste sans Barbe” and (bottom) a color pencil drawing by Sunnyrays.

row shows a color pencil drawing by Ray Sun (Sunnyrays). Note how the relit images preserve the artist’s shading style. These two examples cover quite distinct techniques and demonstrate the effectiveness of our solution. Section 6 presents several additional examples demonstrating the robustness of our approach to also handle photorealistic paintings, caricatures, and other types of artwork. We also compare relighting results, as well as albedos and normal maps recovered by our method, with the ones obtained with commercial and state-of-the-art techniques, demonstrating that paintings and drawings indeed represent a big challenge for previous methods.

The **contributions** of our paper include:

- A practical solution to the problem of subject relighting in paintings and drawings. Our technique is robust to variations in artistic styles, brush strokes, use of colors to convey shading information, and to physically-inconsistent shadings (Section 3);
- A technique to refine a user-provided rough specification of the source shading, which is based on the correlation of color and shading values (Section 4.1);
- A multiscale scheme to refine the target shading, allowing our method to compute appropriate shading-ratio images (Section 4.2);

- A technique that mimics the artists’ original use of colors to represent dark and bright shades (Section 5);
- A technique to transfer smooth normal and depth maps to single images, which can be applied to photographs, paintings, and drawings (Section 6).

2 Related Work

Image relighting has been addressed in different ways over the years. To the best of our knowledge, all previous techniques have focused on photographs and videos, and the relighting of paintings and drawings has remained unexplored. Essentially, existing techniques can be classified as *geometry-based* and *image-based*. Geometry-based methods use a 3-D representation for rendering an object under new illumination, while image-based ones rely on multiple images for relighting.

2.1 Geometry-Based Methods

Inverse-lighting techniques try to estimate and then modify the illumination of a scene based on geometric and reflectance descriptions. In order to acquire ge-




Fig. 2 Image-model registration. (a) Input image I . (b) Segmented reference object O (face). (c) 3-D face model M used to create a proxy for O . (d) Initial image-model registration by translating, rotating, and scaling M to fit O . (e) Deformed model using Green coordinates to improve the fitting of M to O . (f) Resulting transformed and deformed model (reference proxy P).

ometry and reflectance information, these techniques rely on 3-D scanners [13, 25], use multiple photographs taken under controlled settings [23], or create a virtual scene using 3-D models [9]. These techniques require access to the original scene and/or information about scene geometry and reflectance.

Multi-view stereo methods can be used to improve geometry recovery for scenes under general illumination [43, 44]. These techniques have been used for estimating illumination and texture maps [47], and for relighting of human performances [19]. However, they require multiple images as input and, therefore, cannot be used for the relighting of paintings and drawings.

For *face relighting*, many methods use simple models, such as ellipses [7] or a generic face model [42]. Others use morphable models to recover an approximated 3-D face representation [8, 40].

Some interactive techniques allow the user to specify or reconstruct *normal maps*. Okabe et al. use a pen-based interface to draw sparse surface normals [27], while other approaches are based on simple markups for the normal reconstruction [45, 46]. Lopez-Moreno et al. [22] estimate depth information from images and use it for achieving stylization (NPR) effects. These normal/depth-map recovering methods are heavily dependent on image gradients, making them inappropriate for paintings with heavy brush strokes and other high-frequency contents.

In the area of hand-drawn illustrations, several techniques focus on adding lighting effects to cartoons and outline drawings. These methods use normals assigned to edge lines [17] or user-specified relative depth [35, 36] to create geometric representations by inflating outlines. The resulting geometry is then shaded using computer graphics techniques.

2.2 Image-Based Methods

To perform relighting, some image-based techniques use a sparse set of images taken under controlled illumina-

tion [15, 24, 38], create a mosaic combining parts of various images [1, 2], or compute a linear combination of some basis images [3, 37]. In the case of paintings and drawings, however, we are restricted to a single image.

Several techniques have used *ratio images* to perform image relighting [28, 31, 39]. As opposed to previous uses, our *RGB shading-ratio images* take into account the artistic use of colors to encode shading.

For the specific case of *face relighting*, Li et al. proposed a logarithmic total-variation model to retrieve the illumination component and transfer it between faces [20]. Other methods use edge-preserving filters [11] or templates copied by artists [12] to transfer illumination from a reference to a target face. These techniques depend on the alignment, working for nearly-frontal faces.

3 Relighting Using Shading Proxies

Computer vision techniques used for extracting geometry or recovering shading from images make strong assumptions about the scene material properties (*e.g.*, full Lambertian surfaces, uniform albedo, etc.), and lighting conditions. Unfortunately, such assumptions are not satisfied by paintings and drawings, which often contain physically-inconsistent shadings.

We use an interactive approach to retrieve a coarse approximation for the imaged geometry. Unlike methods that focus on 3-D object manipulation on images [10, 18], we only require a geometric representation for the visible surfaces that should be re-illuminated. Our technique is simple and intuitive, using warped 3-D models as shading proxies to such surfaces. Similar to Kholgade et al. [18], we make use of available 3-D models to approximate the image objects. The ever increasing number of available 3-D models makes our technique capable of relighting virtually any object. We provide interactive tools to create shading proxies from approximate 3-D models, as well as to fit these proxies to the input images. The details are presented next.




Fig. 3 Our relighting framework. (a) Input image. (b) and (c) Image and proxy with color-coded matched features. (d) and (e) Guess (S_{sg}) and target (S_t) shadings over the warped proxy. (f) and (g) Refined source (S_{sr}) and target (S_{tr}) shadings. (h) RGB shading-ratio image. (i) Relit image.

3.1 Image-model Registration

Given an input image I containing a segmented reference object O to be relit, we choose an approximate 3-D model representation M to create a proxy for O (Figure 2). M can be obtained from private databases, from the Internet, or from any free or commercial 3-D model repositories. Our system allows the user to superimpose M and I , and align them through a series of translation, rotation, and scaling operations applied to M . This provides some initial registration (Figure 2 (d)). Since M is just an approximate representation for O , the user can deform M to better match the image. This is achieved using Green coordinates [21] associated with the vertices of the model’s bounding box (Figure 2 (e)). At the end of the registration process, the transformed and deformed model resembles O and is called the *reference proxy* P (Figure 2 (f)).

3.2 Feature-Correspondence Mapping

Since P only represents a rough approximation to the geometry of the segmented object O , we further refine their matching. For this, we establish correspondences among silhouettes and other relevant edges visible in both O and in P_o – the perspective projection image of P as seen from the same viewpoint used for image-model registration. Such correspondences are used to create a coherent feature-correspondence mapping (similar to the dense correspondence in Shih et al.’s work [33]). The silhouette edges of O and P_o are automatically retrieved using the technique by Suzuki and Abe [34]. Both silhouettes are divided into segments defined by pairs of matched corners (retrieved using Shi and Tomasi’s method [32]). The user can decide to delete some of the automatically-retrieved segments and corners. In addition, one can interactively define new pairs of corresponding segments using a technique similar to *intelligent scissors* [26]. Figures 3 (b) and (c) show a color-coded representation for the corresponding pairs of feature segments defined for the example shown in Figure 2. Since the sizes of corresponding

segments are unlikely to match, all segments are parameterized in the $[0, 1]$ interval, so that pairs of corresponding points are easily matched. We call this the **feature-correspondence mapping**, which is used to warp (using Delaunay triangulation) the selected features from P_o to the ones in O (Figure 3 (d)), thus defining a mapping from the reference proxy P to O . Figure 3 shows all the steps of our relighting pipeline.

4 Shading

Image relighting is obtained as a pixel-wise multiplication of the input image I by a shading-ratio image:

$$I_r(u, v) = I(u, v) \left(\frac{S_t(u, v)}{S_s(u, v)} \right), \quad (1)$$

where I_r is the relit image, S_s is the original shading of the painting/drawing, and S_t is the target shading. Existing intrinsic-image decomposition techniques [6, 16] cannot be used for estimating the shading of paintings and drawings, as in these representations shading is often encoded by colors, and tends to be physically inconsistent. Thus, besides the target shading, we also request the user to provide a *guess shading* (S_{sg}) for the input image. Both shadings are specified on the reference proxy using OpenGL renderings. Given S_{sg} , our technique estimates how colors change over the painting/drawing as a result of variations in S_{sg} . This information is then used to refine S_{sg} , obtaining a *refined source shading* S_{sr} (Eq. (3)), which in turn is also used to adjust the provided target shading (S_t).

4.1 Guess Shading Refinement

To refine S_{sg} , our technique starts by filtering the segmented object image O using a low-pass edge-aware filter to smooth brush strokes, obtaining a filtered image \hat{O} . For each color channel $\beta \in \{R, G, B\}$ in \hat{O} , we find the function f_β that minimizes

$$\sum |f_\beta(\hat{O}_\beta(u, v)) - S_{sg}(u, v)|^2. \quad (2)$$

We have tested many models of functions for f_β (logarithmic, higher-degree polynomials, and others), but the use of linear functions provided the most attractive alternative. Besides their simplicity, they also provide an immediate monotonic inverse (a nice property, which will be explored in Section 5). Our method then uses f_β to map the β channel value of each pixel in O to a corresponding shading value, producing a reconstructed shading image $S_\beta = f_\beta(O_\beta)$ for each channel β (Figures 4 (a)-(c)). Finally, the *refined source shading* (Figure 4(d)) is obtained as a linear combination of S_R , S_G , and S_B :

$$S_{sr} = \sum_{\beta} \frac{\lambda_{\beta}}{\omega} S_{\beta}, \quad (3)$$

where $\lambda_{\beta} = \text{corr}(S_{sg}, O_{\beta})^2$ is the squared cross-correlation coefficient obtained from all color-shading pairs, and $\omega = \sum \lambda_{\beta}$ is a normalization factor. Eq. (3) ensures that color channels that better explain the shading will have higher contributions for the refined shading.




Fig. 4 Refined source shading. (a)-(c) Reconstructed shading images using f_β for the R, G, and B channels, respectively, of van Gogh’s self portrait. (d) Refined source shading (S_{sr}) obtained as a linear combination of (a) to (c).

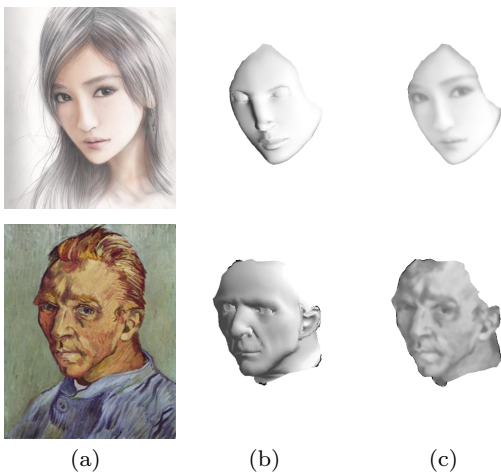


Fig. 5 Obtaining shading from paintings and drawings. (a) Input images. (b) User-provided guess shadings. (c) Refined source shadings obtained with our technique.



Fig. 6 Pairs of albedos recovered from the images shown in Figure 5(a) using the user-provided guess shading (Figure 5(b)) and the refined shadings (Figure 5(c)), respectively.

Figure 5 compares examples of guess shadings and their corresponding refined source shadings obtained with our technique. van Gogh’s example is quite challenging, as the artist used colors to convey lit and in-shade areas. Note that since our technique refines the guessed shading based on shading-color correlation, it may incorrectly consider part of the albedo as shading. This is visible, for example, in Figure 5, where the colors of the eyes are quite different from the colors of the rest of the faces. Figure 6 shows the corresponding recovered albedos using the guessed and refined shadings.

4.2 Target-Shading Refinement

The refined source shading may contain some high-frequency details due to brush strokes and subtle color changes (Figure 5(c)). Since a target/source shading ratio is used to obtain a relit image, we also need to adjust (*i.e.*, refine) the user-provided target shading accordingly to avoid introducing artifacts in the relit image. For this, we use a multiscale approach (similar to Shih et al.’s [33]) to introduce these high frequencies in the target shading, while preserving its low-frequency content. Our method decomposes each of the three shadings (S_{sg} , S_{sr} , and S_t) using a separate Laplacian pyramid. We then combine their information to obtain a fourth pyramid, from which a refined target shading (S_{tr}) is reconstructed. We only decompose these pyramids until the level $\lfloor \log_2(\min(w, h))/2 \rfloor$, where w is *width* and h the *height* of the input image. The coarsest level of the pyramid built for S_{tr} is copied from the one for S_t . Decomposing until this level preserves the target low-frequency shading content, while our multiscale method refines high-frequency details. We now explain how we introduce such details into the S_{tr} pyramid.

For all pyramid levels, we use a gain map similar to Shih et al.’s [33]. However, unlike Shih et al., we compute the gain map based on the user-informed shadings (S_t and S_{sg} – see Eq. (4)), and apply it to each level of the refined-source-shading pyramid to obtain the corresponding level of the refined-target-shading pyramid (Figure 7). Additionally, the gain sign is important for relighting, as this takes into account Laplacian edge

flips due illumination changes on the reference proxy. Therefore, for the l^{th} level of the Laplacian pyramids, we compute the gain map Γ as:

$$\Gamma(l) = \frac{\Pi_l(S_t) \otimes G_\sigma}{\Pi_l(S_{sg}) \otimes G_\sigma}, \quad (4)$$

where $\Pi_l(X)$ is the l^{th} level of the Laplacian pyramid $\Pi(X)$ obtained from image X . G_σ is a Gaussian kernel with standard deviation σ . Unlike Shih et al., we only apply gain to pixels from $\Pi_l(S_{sr})$ that satisfy the condition

$$\frac{\Pi_l(S_{sr})(u, v) \otimes G_\sigma}{\Pi_l(S_{sg})(u, v) \otimes G_\sigma} < \alpha. \quad (5)$$

This ensures that the gain map will only modulate pixels where Laplacian values in S_{sg} and S_{sr} are similar. This prevents possible smoothing of details from S_{sr} not found in S_{sg} , which would happen for pixels whose gain-map value is in the $[0, 1)$ interval. Since we down-sample each Laplacian level, we use a single Gaussian kernel for all pyramid levels. For the results shown in the paper, we used $\alpha = 2$ in Eq. (5). This value was determined experimentally. The finer levels are obtained as illustrated in Figure 7: for each level l , if pixel at position (u, v) satisfies Eq. (5), $\Gamma(l)(u, v)$ is obtained according to Eq. (4); otherwise, $\Gamma(l)(u, v) = 1$.

The refined target shading is then reconstructed from a Laplacian pyramid $\Pi(S_{tr})$. This guarantees that the low-frequency content of the target shading will be used in the computation of its refined version. The refined target level is then computed as the pixel-wise product:

$$\Pi_l(S_{tr}) = \Pi_l(S_{sr}) \times \Gamma(l). \quad (6)$$

Figure 8 compares the refined target shading S_{tr} with S_t and S_{sr} for the example of the van Gogh’s painting shown in Figure 5 (a) (bottom).

4.2.1 Matching Intensity Levels

While the previous scheme is good for refining high-frequency details in shading, the low-frequency content also needs adjustment. As S_{sr} was estimated using our refinement procedure based on shading-color correlation, the range and intensities of its shading values may differ when compared to the ones of S_{tr} . This must be addressed to prevent unwanted changes in illumination intensity and contrast. While one might consider using techniques for histogram matching [29, 30, 41], they are not suited for this problem. For instance, the user could want to change the intensity of the illumination (or the proportion between direct and ambient illumination), something that such histogram-transfer methods would not allow. We present two simple approaches




Fig. 7 Target shading refinement at pyramid level l . The refined target shading is computed as a pixel-wise multiplication between the refined source shading and a gain map.



Fig. 8 Target-shading refinement. (left) Target shading specified by the user on the warped proxy. (center) The refined target shading computed by our Laplacian-pyramid scheme. (right) The refined source shading.

for performing this low-frequency adjustment between S_{sr} and S_{tr} . While one modifies the already computed S_{sr} , the other changes S_{tr} . Both approaches work over the coarsest level of the Laplacian pyramids, where the low-frequency content can be properly manipulated.

Adjusting the S_{sr} low-frequency content: such difference in intensities between S_{sr} and S_{tr} could be solved if both coarsest levels of pyramids derive from the same place (e.g., the rendering under user-specified illuminations). This can be obtained replacing the coarsest level l_c of the S_{sr} pyramid by its corresponding level in S_{sg} :

$$\Pi_{l_c}(S_{sr}) = \Pi_{l_c}(S_{sg}). \quad (7)$$

Let \hat{S}_{sr} be the result of the reconstructed pyramid. The low-frequency content of both \hat{S}_{sr} and S_{tr} come from user-specified illumination over the proxy. Although such adjustment could introduce shading errors, such low-frequency errors are barely noticed in the final result.

Adjusting the S_{tr} low-frequency content: An alternative is to compute a new \hat{S}_{tr} that takes into account changes in intensity values. This is achieved by

modulating the coarsest level l_c of the S_{tr} pyramid using its corresponding gain map $\Gamma(l_c)$:

$$\Pi_{l_c}(S_{tr}) = \Pi_{l_c}(S_t) \times \Gamma(l_c) \quad (8)$$

Note that the reason for modulating the coarsest level of the pyramid is to change the user-specified target shading in the same way as the user-specified S_{sg} was changed. For the other levels of the pyramid, the motivation of Γ was to compute finer details to S_{tr} .

5 Artistic Relighting

One could consider obtaining the shading-ratio image simply by dividing the refined-target by the refined-source shadings. However, in the case of paintings and drawings, artists often use color to convey shading. Recall from Section 4.1 that f_β maps the β channel value of each pixel in O to a corresponding shading value. Thus, the shading ratio has to be computed on a per-channel basis, according to the functions f_β^{-1} . Each f_β^{-1} maps (through an affine transform) the range of refined shading values to a new shading range represented by the β -channel values in input image I , with $\beta \in \{R, G, B\}$. The shading-ratio image for channel β is then computed as the pixel-wise ratio

$$S_{\beta_ratio} = \frac{f_\beta^{-1}(S_{tr})}{f_\beta^{-1}(S_{sr})}, \quad (9)$$

and the **relit image** is obtained as:

$$I_{r\beta}(u, v) = S_{\beta_ratio}(u, v) \times I_\beta(u, v), \quad (10)$$

where $I_{r\beta}$ is the β channel of the relit image I_r , I_β is the β channel of the input image I , and (u, v) are pixel coordinates.

As an alternative to the automatically computed f_β functions (Section 4.1), we allow the user to specify an *in-shade* and a *lit* region by directly pointing on the image (at pixels $I(u_d, v_d)$ and $I(u_l, v_l)$, respectively). Let R_d, G_d, B_d be the average R, G, and B values computed at a 10×10 pixel neighborhood around $I(u_d, v_d)$. Likewise, let R_l, G_l, B_l be the R, G, and B average values computed at a similar neighborhood around $I(u_l, v_l)$. Thus, the new function g_β is defined as the line passing through (R_d, G_d, B_d) and (R_l, G_l, B_l) . This procedure defines a set of three linear functions (Figure 9). Figure 10 shows the colored versions of the shadings ($g_\beta(S_{tr})$ and $g_\beta(S_{sr})$), as well as the corresponding colored shading-ratio image. This ratio ensures that the relighting process preserves the shading encoding.

Although the automatic fitting of functions f_β might seem more convenient for novice users, the specification of in-shade and lit regions gives the user more control




Fig. 9 Artistic shading-color interpolation for the van Gogh painting. These functions encode how the values of each color channel vary based on the guessed shading values. Shadow regions have proportionally more red than lit regions, where the red and green values ensure a yellowish color.



Fig. 10 Colored shadings (by g_β) for *van Gogh's* relighting. (left) Colored refined source shading. (center) Colored refined target shading. (right) Colored shading-ratio image.

over the final results. In our system, we allow both options. Figure 11 compares the results obtained when relighting the painting in Figure 1(a) using three different approaches. On the left, one sees the result obtained with the use of a monochromatic shading-ratio image, as done in conventional image relighting (Eq. (1)). Note that the use of the same scaling factor for all channels of a given pixel leads to unpleasant results. The image in the center shows the relight obtained with the use of an RGB shading-ratio image using Eq. (10) with automatically computed f_β functions (Section 4.1). The image on the right shows the recoloring result obtained also using Eq. (10) with g_β functions computed based on the specification of in-shade and lit regions. The




Fig. 11 Relighting using different shading-ratio strategies: (left) Monochromatic. (center) RGB using automatically computed f_β functions. (right) RGB using g_β functions computed using user-specified in-shade and lit regions. The color maps represent the mappings between shading values in the $[0,1]$ range and RGB triplets.

color maps next to the two rightmost images represent the mappings between shading values in the $[0, 1]$ range to RGB triples, for van Gogh’s face. Due to the least-squares fitting, the automatically-computed color map tends to produce smoother changes in the relighting.

6 Results

We have implemented the techniques described in the paper using C++ and OpenGL, with the help of OpenCV and MATLAB functions. For all examples shown in the paper, we used OpenGL to illuminate the shading proxies using a single point light source. We used our approach to relight a large number of images including paintings, drawings, and other kinds of artwork. A typical interactive model-registration session, including feature-correspondence mapping, takes from 2 to 7 minutes (in the case of more complex shapes). Relighting is performed in real time using such mappings, which can be saved and reused. The times reported for performing relighting tasks were measured on a 3.4 GHz i7 CPU with 16 GB of RAM.

This section presents relighting examples for distinct styles of paintings, drawings, digital paintings, and even unusual kinds of artworks. The examples shown in Figures 1, 13, and 17 were produced with user-specified in-shade and lit regions. The other ones use f_β functions computed as described in Section 4.1. The examples in Figures 13, 16, and 28 were produced adjusting the low frequency content of the refined-target shading (Section 4.2.1). For the other examples, the adjustment was applied to the refined-source shading. For all examples, the Gaussian filters used in Eqs (4) and (5), as well as for creating the Laplacian pyramids, use $\sigma = 0.5$. Since our shading proxies can be any 3-D models, our image-model registration step (Section 3.1) allows us to relight images of objects with arbitrary geometries. We, however, show mostly face examples due to the great diversity of artistic styles that can be found in portraits. Our method is robust to strong brush strokes, painting styles, and use of colors to encode shading. For all examples in this section, we also show, together with the original image, an inset exhibiting the reference proxy P (see Section 3.1) with the user-provided guess shading. Note that a single 3-D model was used for the relighting of all male faces shown in the paper. In each case, the model has been transformed and warped to fit the individual characteristics of the subjects. By transforming and warping 3-D models, our approach provides great flexibility, allowing us to handle a large number of images using a small number of models. Our results illustrate not only changes in illumination di-



Fig. 12 Relighting of the painting *William I, Prince of Oranje*, by Adriaen Key. (left) Original. (center) Relighting from above. (right) Relighting from right.



Fig. 13 Relighting of a Tim Benson’s painting: *Dad Looking Down*. (a) Original. (b) Relighting from left with no low-frequency adjustment. (c) Relighting with the same lighting condition as (b), but adjusting the low-frequency over the refined source shading. (d) Relighting from the left adjusting the refined-target-shading low-frequency content.

rection, but also in its intensity, achieving new moods while preserving the original color-shading encodings.

Figure 1 illustrates the relighting of two very distinct styles: a van Gogh’s painting (top) and a color pencil drawing (bottom). Relighting van Gogh’s work is quite challenging. Note how our technique preserved the artist’s use of colors to represent shading information, even when using distinct light intensities. The pencil drawing example has a very smooth shading.

Figure 12 illustrates the relighting of a photorealistic painting by the Flemish renaissance painter Adriaen Key. Figure 13 shows the relighting of Tim Benson’s portrait *Dad Looking Down*. The image in (b) shows the relighting with no low-frequency content adjustment. In (c), the low-frequency content was adjusted by manipulating the refined source shading. In (d), the image has been relit using the same lighting condition, but the low-frequency content was adjusted by manipulating the refined target shading. This is another challenging example due to color discontinuities and sharp brush strokes. Also, note how both adjustments generate pleasing results.

Figure 14 shows the relighting of a digital drawing by Mark Hammermeister. This example illustrates the use of our technique with caricatures, which employ very distinctive geometry and illumination styles.



Fig. 14 Relighting of a caricature by Mark Hammermeister. (left) Original. (center) Relighting from left using a stronger light source. (right) Relighting from the right.



Fig. 15 Relighting of a digital painting by J. P. Targete. (left) Original. (center) Frontal relit. (right) Relit from left.




Fig. 16 Relighting of the digital painting *Big Face Dragon Head*, by Vincent Hie (*wallace*). (left) Original. (center) Relit from left. (right) Relit from below. Note how the proxy (inset on the left image) is very different from the painting geometry.

Figure 15 shows another example of relighting applied to a digital painting. This is another challenging example, as the character contains lots of high-frequency surface details in his face, ears, and horns. This should be contrasted with the smooth reference proxy used for specifying the guess and target shadings (see inset in Figure 15 (left)). Nevertheless, the shading-ratio image obtained from our refined source and target shadings leads to proper relighting of even highly-detailed surfaces, such as this one. Figure 15 (center) and (right) show the character relit from the front and from the left. Note the surface details. Similarly, note how our method generates appropriate shadings for highly-complex geometries (Figure 16), even using a shading proxy that is fairly different from the imaged character (see inset).

Figure 17 shows a portrait by Andrew Myers created using screws. This image contains high-frequency



Fig. 17 Relighting of a portrait made with screws, by Andrew Myers. (left) Original. (center) Relit from left. (right) Relit from above using a stronger light source.



Fig. 18 Relighting of a painting of a rubber duck, by Elizabeth Edwards. (left) Original. (center) Relighting from below. (right) Relighting from right.

details that are likely to hamper shading estimation by traditional intrinsic-decomposition techniques. Despite that, our method is able to relight it, preserving the essence of this unique piece of artwork.

Figure 18 shows the relighting of a painting of a rubber duck, by Elizabeth Edwards. Note how the color transitions between yellow and orange are smooth.

Figure 19 shows the relighting of a painting of a red BMW, by *Don4x*. Figure 20 shows the relighting of a Camaro drawing also by Elisabeth. The original images contain several details that cannot be directly mapped from the reference proxies (insets). These details are handled by our shading refinement procedures. We can also use different mapping functions for each shading (target and source), achieving results that mimics change of the illuminant color. Figure 21 compares the results achieved by using the same encode-mapping functions (on the center), and by using different mapping functions for each shading (right).

Many results were created by significantly changing the lighting conditions, exhibiting the flexibility provided by our method. Nonetheless, subtle modifications on the illumination can significantly change the interpretation of images. We encourage the readers to check out the supplemental videos illustrating this observation.

6.1 Normal and Depth-Map Transfer

The feature-correspondence mapping generated by our method can be useful for additional applications. For instance, it makes it straightforward to transfer information from the 3-D proxies to pictures. We exploit this possibility to transfer normal and depth maps from the 3-D proxies to paintings and drawings. Figure 22 shows



Fig. 19 Relighting of a Red BMW drawing, by *Don4x*. (left) Original. (center) Relit from right. (right) Relit from above.



Fig. 20 Camaro drawing, by Elisabeth K. (Lizkay). (left) Original. (center) Relit from right. (right) Relit from above.



Fig. 21 Mustang drawing, by Elisabeth K. (Lizkay). (left) Original. (center) Relit from above using the same mapping functions to encode color in both (target and source) shadings. (right) Using the same shadings as (center), but encoding colors using different mapping functions to convey an illuminant with reddish color.




Fig. 22 Examples of recovered normal maps from paintings and drawing. All proxies derive from the same 3-D model.

examples of normal maps transferred to Tim Benson's and van Gogh's paintings, to a digital caricature and to an outline drawing. Conventional techniques for estimating normal maps from images are based on image gradients and, therefore, do not work for these examples. Our approach is capable of transferring smooth

normal maps that capture the essence of each of these images, even though the used proxies all derive from the same 3-D model. This observation also applies to transferred depth maps (Figure 23).

6.2 Comparisons

This section compares our relighting results with the ones obtained with the state-of-the-art technique [6] and with a commercial software [4]. We also compare our intrinsic decompositions of paintings and drawings against the ones produced by [6]. This is intended to show that the problem of estimating shading and albedo for such images is significantly harder than for pho-




Fig. 23 Examples of recovered depth maps for the paintings and drawing shown in Figure 22.




Fig. 24 Comparison with Chen et al.’s technique [12]. (a) Reference image along with the light and shadow templates. (b) Original drawings, by David Hajari (top) and Judy Wise (bottom). (c) Results of applying the templates in (a) using Chen et al.’s technique. (d) Result of our technique simulating similar light conditions.

torrealistic pictures. In addition, we compare normal maps obtained with our method with the ones recovered by [6] and by a commercial technique [14].

As stated before, to the best of our knowledge, ours is the first method designed for the relighting of paintings and drawings. Unfortunately, we could not find any publicly available code or implementation for any of the state-of-the-art image relighting methods. Thus, to do an as fair as possible comparison, we have taken results from these methods and tried to relight them under similar light conditions using our technique. Figure 24 compares our results to the ones of Chen et al.’s technique [12] (the only technique for which we have found results with sufficient resolution for comparison). Note that their method requires light and shadow templates drawn by professional artists specifying the areas to be lit and shaded. Ours, on the other hand, allows users to freely relight the images. In addition, our results exhibit smoother shadings and more coherent colors, specially in shadow regions.

We have only found a single commercial software for performing relighting: *Portrait Professional* [4]. As its name suggests, this software is specific for portraits, not working for arbitrary objects. In addition, it does not support a wide range of lighting conditions. Figure 25 (c) shows relighting examples produced by the software. Note that, besides incorporating only small changes in illumination, it does not discount the original one. For instance, in the relighting of the *orc* image (third row), the proxy in the inset shows an illumination from the left while the resulting image appears to be lit frontally. In addition, this method uses tone-skin colors, generating unrealistic results for paintings with unusual

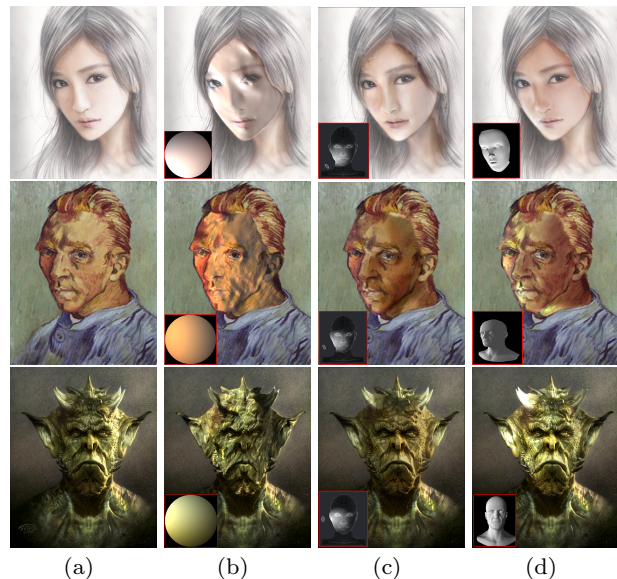


Fig. 25 Comparing relighting results. (a) Original paintings and drawings. (b) Relightings achieved by changing spherical-harmonics parameters of Barron and Malik’s method [6], illumination on insets. (c) Relighting achieved using a commercial software [4], target illumination on inset. (d) Our method results with *proxies* on insets.

shading encodings (van Gogh’s portrait), or faces with peculiar colors (*orc*).

We also present comparisons with Barron and Malik’s SIRFS method [6], one of the most effective techniques for estimating shape, reflectance, and shading from single images. Figure 25 (b) shows the use of SIRFS for relighting. This was done by carefully changing the spherical-harmonics coefficients that encode the object illumination (see the insets). The results look unnatural and displeasing due to failure in properly estimating shape and illumination.

We have also compared other features of our technique against Barron and Malik’s. While their method works well with natural (photorealistic) images, it is not nearly as effective for paintings and drawings. Figure 26 (top) shows the recovered normal maps for three of the images shown in Figure 22, and recovered by using a commercial software [14] (bottom). As traditional techniques for normal-map estimation heavily depend on image gradients, they are likely to fail when used with paintings containing brush strokes and other artistic features.

Figure 27 compare the albedos retrieved by Barron and Malik’s technique [6] (top) and by our method (bottom). While their approach indeed yields good results for photorealistic images (Figure 27 (left)), it does not account for unusual shading encodings, which in turn leads to poor relighting results.




Fig. 26 Normal maps obtained for some of the images shown in Figure 22 using Baron and Malik’s technique [6] (top), and a commercial software [14] (bottom).



Fig. 27 Examples of albedos retrieved using Barron and Malik’s technique [6] (top) and our method (bottom).

6.3 Discussion and Limitations

We used OpenGL’s point light sources to illuminate the reference proxy because this lends to real-time rendering, but other lighting models can be used. Our target-shading-refinement procedure ensures appropriate ratios that allow one to experiment with different illumination intensities. We experimented our technique with other color spaces (CIE Lab, HSV, HSL), but according to our experience the RGB color space yields the most pleasing results.

For the specific case of portrait relighting, one could replace the image-model registration step by an automatic procedure for registering 3-D models to face images [8, 40]. However, since our procedure for shading refinement does not require that the 3-D model or its shading exactly matches the input image’s shape or shading (see, for example, Figure 15 and Figure 16), we argue that it is still advantageous for the user to perform approximate registration using just a few mouse clicks. Even simple shapes can be used as shading prox-

ies. In Figure 28, we relight a painting of an *angry bird* using a sphere as shading proxy. Nonetheless, the refined source shading (inset in the left image) still approximates the painted object well. Our approach not only provides more creative freedom, but also allows for the relighting of virtually any object.

Since we only relight the segmented objects, the users should be careful to not introduce too noticeable inconsistencies in the scene illumination. Although one can use an arbitrary number of reference proxies to re-illuminate an image, our technique does not handle shadows or other global illumination effects. Finally, as mentioned in Section 4.1, our technique may incorrectly treat part of the albedo as shading.

7 Conclusion

We have presented a practical solution to the problem of subject relighting in paintings and drawings. It works by correlating color and shading values and using this information to refine a user-provided source-shading guess. Given a target shading, we use a multiscale approach to add to it possible high-frequency information found in the refined source shading. These new source and target shadings are used to compute a per-channel shading-ratio image, which takes into account the colors used by the artist to encode shading information. The relit image is then obtained through pixel-wise multiplication between the input color image and the obtained shading-ratio one. We have demonstrated the effectiveness of our technique on a variety of artistic styles, including the use of strong brush strokes, unconventional shading encodings, drawings, and other types of artwork. Our method is the first to perform relighting of paintings and drawings, and is capable of preserving the artist’s intended use of colors to encode shading information. We have compared our results to the ones obtained with several other techniques, showing that indeed none of them properly works for paintings and drawings.



Fig. 28 Relighting the painting *Big Bro Angry Bird*, by Sylvia W. (left) Original painting. (center) Relighting from above. (right) Relighting from the front. We have used a simple sphere as the shading proxy, the refined source shading is shown in inset.

Our technique can also be used to transfer smooth normal and depth maps from 3-D proxies to both photo- and non-photorealistic images. The resulting maps capture the essence of the represented underlying surfaces. Given its flexibility, robustness, and ease of use, our technique can help artists and casual users to experiment with various lighting effects on existing artworks, enabling new and creative applications.

Acknowledgments This work was sponsored by CAPES and CNPq-Brazil (grants 134134/2013-3, 306196/2014-0, and 482271/2012-4). We would like to thank the artists for allowing us to use their artworks.

References

- Agarwala A, Dontcheva M, Agrawala M, Drucker S, Colburn A, Curless B, Salesin D, Cohen M (2004) Interactive digital photomontage. *ACM TOG* 23(3):294–302
- Akers D, Losasso F, Klingner J, Agrawala M, Rick J, Hanrahan P (2003) Conveying shape and features with image-based relighting. In: *VIS'03*, pp 46–51
- Anrys F, Dutré P (2004) Image-based lighting design. *CW Reports CW382*, K. U. Leuven
- Anthropics (2015) PortraitPro v12.4 - Trial Version. <http://www.portraitprofessional.com/>
- Barron JT, Malik J (2012) Color constancy, intrinsic images, and shape estimation. In: *ECCV'12*, pp 57–70
- Barron JT, Malik J (2015) Shape, illumination, and reflectance from shading. *TPAMI*
- Basso A, Graf H, Gibbon D, Cosatto E, Liu S (2001) Virtual light: digitally-generated lighting for video conferencing applications. In: *ICIP'01*, vol 2, pp 1085–1088
- Blanz V, Vetter T (1999) A morphable model for the synthesis of 3d faces. In: *SIGGRAPH '99*, pp 187–194
- Boivin S, Gagalowicz A (2002) Inverse rendering from a single image. In: *CGIV'02*, pp 268–277
- Chen T, Zhu Z, Shamir A, Hu SM, Cohen-Or D (2013) 3Sweep: Extracting editable objects from a single photo. *ACM TOG* 32(6):195:1–195:10
- Chen X, Chen M, Jin X, Zhao Q (2011) Face illumination transfer through edge-preserving filters. In: *CVPR '11*, pp 281–287
- Chen X, Jin X, Zhao Q, Wu H (2012) Artistic illumination transfer for portraits. *CGF* 31(4):1425–1434
- Costa A, Sousa A, Nunes Ferreira F (1999) Lighting design: A goal based approach using optimisation. In: *EGWR'99*, pp 317–328
- CrazyBump (2014) Demo Version 1.2. www.crazybump.com/
- Debevec P, Hawkins T, Tchou C, Duiker HP, Sarokin W, Sagar M (2000) Acquiring the reflectance field of a human face. In: *SIGGRAPH '00*, pp 145–156
- Hu K, Liu Y, Dong Q, Liu H, Xing G (2014) Color face image decomposition under complex lighting conditions. *The Visual Computer* 30(6-8):685–695
- Johnston SF (2002) Lumo: Illumination for cel animation. In: *NPAR '02*, pp 45–ff
- Kholgade N, Simon T, Efros A, Sheikh Y (2014) 3d object manipulation in a single photograph using stock 3d models. *ACM TOG* 33(4)
- Li G, Wu C, Stoll C, Liu Y, Varanasi K, Dai Q, Theobalt C (2013) Capturing relightable human performances under general uncontrolled illumination. *CGF* 32(2):275–284
- Li Q, Yin W, Deng Z (2009) Image-based face illumination transferring using logarithmic total variation models. *The Visual Computer* 26(1):41–49
- Lipman Y, Levin D, Cohen-Or D (2008) Green coordinates. *ACM TOG* 27(3):78:1–78:10
- Lopez-Moreno J, Jimenez J, Hadap S, Reinhard E, Anjyo K, Gutierrez D (2010) Stylized depiction of images based on depth perception. In: *NPAR '10*
- Loscos C, Drettakis G, Robert L (2000) Interactive Virtual Relighting of Real Scenes. *IEEE TVCG* 6(3):289–305
- Malzbender T, Gelb D, Wolters HJ (2001) Polynomial texture maps. In: *SIGGRAPH'01*, pp 519–528
- Marschner SR (1998) Inverse rendering for computer graphics. PhD thesis, Cornell University
- Mortensen EN, Barrett WA (1995) Intelligent scissors for image composition. In: *SIGGRAPH '95*, pp 191–198
- Okabe M, Zeng G, Matsushita Y, Igarashi T, Quan L, yeung Shum H (2006) Single-view relighting with normal map painting. In: *Pacific Graphics*, pp 27–34
- Peers P, Tamura N, Matusik W, Debevec P (2007) Post-production facial performance relighting using reflectance transfer. *ACM TOG* 26(3):52:1–52:10
- Reinhard E, Adhikhmin M, Gooch B, Shirley P (2001) Color transfer between images. *IEEE Computer Graphics and Applications* 21(5):34–41
- Shapira D, Avidan S, Hel-Or Y (2013) Multiple histogram matching. In: *Proc. IEEE ICIP*, pp 2269–2273

31. Shashua A, Riklin-Raviv T (2001) The quotient image: Class-based re-rendering and recognition with varying illuminations. *IEEE TPAMI* 23(2):129–139
32. Shi J, Tomasi C (1994) Good features to track. In: *CVPR'94*, pp 593–600
33. Shih Y, Paris S, Barnes C, Freeman WT, Durand F (2014) Style transfer for headshot portraits. *ACM TOG* 33(4):148:1–148:14
34. Suzuki S, Abe K (1985) Topological structural analysis of digitized binary images by border following. *CVGIP* 30(1):32 – 46
35. Sýkora D, Sedlacek D, Jinchao S, Dingliana J, Collins S (2010) Adding depth to cartoons using sparse depth (in)equalities. *CGF* 29(2):615–623
36. Sýkora D, Kavan L, Čadík M, Jamriška O, Jacobson A, Whited B, Simmons M, Sorkine-Hornung O (2014) Ink-and-ray: Bas-relief meshes for adding global illumination effects to hand-drawn characters. *ACM TOG* 33(2):16
37. Tunwattanapong B, Debevec P (2009) Interactive image-based relighting with spatially-varying lights. *ACM SIGGRAPH'09 Poster*
38. Tunwattanapong B, Ghosh A, Debevec P (2011) Practical image-based relighting and editing with spherical-harmonics and local lights. In: *CVMP'11*, pp 138–147
39. Wang O, Davis J, Chuang E, Rickard I, de Mesa K, Dave C (2008) Video relighting using infrared illumination. *CGF* 27(2):271–279
40. Wang Y, Liu Z, Hua G, Wen Z, Zhang Z, Samaras D (2007) Face re-lighting from a single image under harsh lighting conditions. In: *CVPR'07*, pp 1–8
41. Welsh T, Ashikhmin M, Mueller K (2002) Transferring color to greyscale images. *ACM TOG* 21(3):277–280
42. Wen Z, Liu Z, Huang TS (2003) Face relighting with radiance environment maps. In: *CVPR'03*, pp 158–165
43. Wu C, Varanasi K, Liu Y, Seidel HP, Theobalt C (2011) Shading-based dynamic shape refinement from multi-view video under general illumination. In: *IEEE ICCV*, pp 1108–1115
44. Wu C, Wilburn B, Matsushita Y, Theobalt C (2011) High-quality shape from multi-view stereo and shading under general illumination. In: *IEEE CVPR*, pp 969–976
45. Wu TP, Tang CK, Brown MS, Shum HY (2007) Shapepalettes: interactive normal transfer via sketching. *ACM TOG* 26(3):44:1–44:5
46. Wu TP, Sun J, Tang CK, Shum HY (2008) Interactive normal reconstruction from a single image. *ACM TOG* 27(5):119:1–119:9
47. Yu T, Wang H, Ahuja N, Chen WC (2006) Sparse lumigraph relighting by illumination and reflectance estimation from multi-view images. In: *EGSR*, pp 41–50