

UNIVERSIDADE FEDERAL DO RIO GRANDE DO SUL
INSTITUTO DE INFORMÁTICA
PROGRAMA DE PÓS-GRADUAÇÃO EM COMPUTAÇÃO

BERNARDO HENZ

Image Relighting using Shading Proxies

Thesis presented in partial fulfillment
of the requirements for the degree of
Master of Computer Science

Prof. Dr. Manuel Menezes de Oliveira Neto
Advisor

Porto Alegre, April 2014

CIP – CATALOGING-IN-PUBLICATION

Henz, Bernardo

Image Relighting using Shading Proxies / Bernardo Henz. –
Porto Alegre: PPGC da UFRGS, 2014.

69 f.: il.

Thesis (Master) – Universidade Federal do Rio Grande do Sul.
Programa de Pós-Graduação em Computação, Porto Alegre, BR–
RS, 2014. Advisor: Manuel Menezes de Oliveira Neto.

1. Image Relighting. 2. Image Manipulation. 3. Computa-
tional Photography. I. Oliveira Neto, Manuel Menezes de. II. Ti-
tle.

UNIVERSIDADE FEDERAL DO RIO GRANDE DO SUL

Reitor: Prof. Carlos Alexandre Netto

Vice-Reitor: Prof. Rui Vicente Oppermann

Pró-Reitor de Pós-Graduação: Prof. Vladimir Pinheiro do Nascimento

Diretor do Instituto de Informática: Prof. Luís da Cunha Lamb

Coordenador do PPGC: Prof. Luigi Carro

Bibliotecária-Chefe do Instituto de Informática: Beatriz Regina Bastos Haro

“Research is what I’m doing when I don’t know what I’m doing.”

— WERNHER VON BRAUN

ABSTRACT

We present a practical solution to the problem of single-image relighting of objects with arbitrary shapes. It is based on a shading-ratio image obtained from the original and target lighting applied to shading proxies (warped versions of 3-D models that approximate the objects to be relit). Our approach is flexible and robust, being applicable to objects with non-uniform albedos. We demonstrate its effectiveness by relighting a large number of photographs, paintings, and drawings containing a variety of objects of different materials. In addition to relighting, our technique can estimate smooth normal and depth maps from pictures, as well as perform intrinsic-image decomposition. Preliminary evaluation has shown that our technique produces convincing results, and novice users can relight images in just a couple of minutes.

Keywords: Image Relighting, Image Manipulation, Computational Photography.

Reiluminação de Imagens utilizando Shading Proxies

RESUMO

Esta dissertação apresenta uma solução prática para o problema de reiluminação de imagens para objetos com geometria arbitrária. Nossa técnica baseia-se no que chamamos de *shading proxies* (versões deformadas de modelos 3D que aproximam o objeto a ser reiluminado) para guiar o processo de reiluminação. Nosso método é flexível e robusto, podendo reiluminar fotografias, pinturas, e desenhos de diferentes objetos de maneira eficaz. Além de reiluminação, nossa técnica pode ser usada para estimar mapas de normais e profundidade, bem como realizar decomposição intrínseca de imagens, e transferir iluminação para desenhos delineados. Uma avaliação preliminar mostra que nossa técnica produz resultados convincentes, e usuários novatos podem reiluminar imagens facilmente em poucos minutos.

Palavras-chave: Reiluminação de Imagens, Manipulação de Imagens, Fotografia Computacional.

LIST OF FIGURES

Figure 1.1:	Image relighting of a photograph (Sharbat Gula, the "Afghan Girl", by Steve McCurry)	15
Figure 1.2:	Image relighting of a van Gogh's painting: <i>Portrait de l'Artiste sans Barbe</i>	16
Figure 1.3:	Transferring shading and normal map to an outline drawing.	17
Figure 3.1:	Image-model registration.	26
Figure 3.2:	Feature-correspondence mapping and actual relighting.	27
Figure 3.3:	Steps of the image-relighting process.	28
Figure 3.4:	Shading-ratio inpainting.	29
Figure 3.5:	Relighting of a girl's face pencil drawing.	29
Figure 3.6:	Interface for selecting regions of interest for relighting.	30
Figure 3.7:	Regular vs artistic relighting.	30
Figure 3.8:	Graph of artistic shading-color interpolation.	31
Figure 3.9:	Relighting of a painting of Isaac Newton by Godfrey Kneller using artistic relighting.	32
Figure 3.10:	Suggesting models for shading proxies.	33
Figure 3.11:	Examples of recovered normal maps from paintings and drawing and comparison with commercial software.	34
Figure 3.12:	Examples of recovered normal maps from drawings and photographs and comparison with commercial software.	35
Figure 3.13:	Examples of recovered depth maps for paintings and drawings.	35
Figure 3.14:	Comparison between our method and state-of-the-art technique for intrinsic-image decomposition.	36
Figure 3.15:	Transferring shading from proxies to 2-D drawings and coloring shading using artistic information.	36
Figure 4.1:	Relighting of Steve McCurry's photograph <i>Afghan Girl</i> using artistic and regular relighting.	39
Figure 4.2:	Relighting of a pencil drawing.	40
Figure 4.3:	Relighting of Tim Benson's portrait "Dad Looking Down".	41
Figure 4.4:	Relighting a car photograph from different directions.	42
Figure 4.5:	Relighting of two colored bottles.	42
Figure 4.6:	Relighting of a Botticelli's painting.	43
Figure 4.7:	Relighting of an Eric Daigh's painting " <i>Allyson</i> ".	43
Figure 4.8:	Relighting of a screw work of Andrew Myers.	44
Figure 4.9:	Relighting of Modigliani's painting <i>Portrait de Jeanne Hébuterne</i>	44
Figure 4.10:	Relighting of a caricature made by Mark Hammermeister.	45

Figure 4.11: Relighting of a black-and-white photograph of a juggler taken by Mike Cipriano.	45
Figure 4.12: Relighting of a black-and-white photograph of Albert Einstein. . . .	46
Figure 4.13: Relighting of a Tracey Costescu’s painting of a Nautilus shell. . . .	46
Figure 4.14: Relighting of a skull by light sources with different colors.	47
Figure 4.15: Relighting of a black vase by manipulating independently each channel of the shading-ratio image.	47
Figure 4.16: Relighting effect by performing color quantization in <i>new shading image</i>	48
Figure 4.17: Relighting effect by applying edge-preserving filters on the <i>new shading image</i>	48
Figure 4.18: Relighting effect by applying edge-preserving filters on the <i>new shading image</i> in our artistic relighting.	49
Figure 4.19: Image and 3-D model used for novice training.	49
Figure 4.20: Relighting hat examples produced by novice users.	50
Figure 4.21: Relighting face examples produced by novice users.	51
Figure 4.22: Relighting Oprah’s portrait that has smooth highlights.	52
Figure 4.23: Defining feature correspondence provides additional freedom to specify relighting effects.	53

LIST OF TABLES

Table 4.1:	Average times taken by the novices in the training.	50
------------	---	----

TABLE OF CONTENTS

1	INTRODUCTION	15
1.1	Thesis Structure	17
2	RELATED WORK	19
2.1	Image Relighting Techniques	19
2.1.1	Geometry-Based Methods	19
2.1.2	Image-Based Methods	20
2.2	Estimating Illumination and Geometry	21
2.2.1	Estimating Illumination	22
2.2.2	Estimating Geometry	22
2.3	Image Composition	23
3	IMAGE RELIGHTING USING SHADING PROXIES	25
3.1	Regular Relighting Method	25
3.1.1	Image-Model Registration	25
3.1.2	Feature-Correspondence Mapping	26
3.1.3	Shading-Ratio Image and Actual Relighting	26
3.1.4	Shading-Ratio Inpainting	28
3.1.5	Selecting Regions of Interest for Relighting	29
3.2	Artistic Relighting	30
3.3	Selecting 3-D Models for Shading Proxies	32
3.4	Normal and Depth-Map Recovery	34
3.5	Intrinsic-Image Decomposition	35
3.6	Shading Transferring	37
4	RESULTS	39
4.1	User Evaluation	49
4.2	Discussion and Limitations	51
5	CONCLUSIONS AND FUTURE WORK	55
5.1	Future Work	55
	REFERENCES	57
	APÊNDICE A REILUMINAÇÃO DE IMAGENS UTILIZANDO <i>SHADING PROXIES</i>	63
A.1	Trabalhos Relacionados	64
A.2	Reiluminação de Imagens	65
A.2.1	Reiluminação Artística	67

A.2.2	Decomposição Intrínseca de Imagens	67
A.2.3	Tranferência de Iluminação	67
A.2.4	Reconstrução de Mapas de Normais e Profundidade	67
A.3	Resultados	67
A.4	Conclusão	69

1 INTRODUCTION

Image relighting is a fundamental problem in image processing. It tries to recreate the appearance of a pictured object or scene under new illumination. This is useful, for instance, when a picture cannot be (easily) retaken in the desired lighting condition, when one would like to re-target a painting to draw attention to a certain area, or to achieve different perceptions of a photograph (Figure 1.1).

However, this is a difficult problem, as a single image carries no explicit information about the scene’s original lighting, or about the objects’ shapes and material properties (*e.g.*, BRDFs). Although in the past years some techniques were proposed to estimate scene illumination from single images under certain conditions (LOPEZ-MORENO et al., 2010; LALONDE; EFROS; NARASIMHAN, 2012), no general solution to this problem is available. Also, recovering object’s shape from a single image is an under-constrained task, for which satisfactory results are not attainable for arbitrary shapes and non-uniform albedo (WU et al., 2007, 2008; BARRON; MALIK, 2012; CHEN et al., 2013). In Section 2.2 we discuss in greater detail the most recent methods for estimating illumination and geometry.

Figure 1.1: Image relighting of a photograph (Sharbat Gula, the "Afghan Girl", by Steve McCurry) using our technique. Note how different lighting conditions produce distinctive effects, changing the perception of the photograph.




We present a practical solution to the problem of single-image relighting of objects with arbitrary shapes. Instead of recovering the object’s geometry from the input image, we use a *shading proxy*, an approximate 3-D model representation for the object, which is interactively transformed (*i.e.*, translated, rotated, and scaled) and warped to mimic the view of the pictured object. This makes our technique robust, since it does not need to make strong assumptions such as uniform albedo and Lambertian reflectance.

We then define correspondences among salient features in the image of the object and in the image/view of the proxy (which we call the *reference image*), from which we derive

a 2-D mapping between them. Conceptually, by illuminating the proxy with (an approximation to) the original lighting conditions as well as with the desired illumination, we can compute a *shading-ratio image* ($\frac{\text{desired}}{\text{original}}$), which is then used for pixel-wise relighting of the object. In practice, good results can be obtained even with poor approximations of the original illumination. This is possible because the relighting process provides real-time visual feedback. Thus, one can obtain the desired relighting effect interactively by changing the shading ratio in an interactive way. The accompanying video provides several examples of real-time re-illuminations that illustrate this concept.

When dealing with paintings, several artists use different colors to convey shading. Relighting them with our regular approach used for photographs would result in unpleasant relightings. Therefore, we also propose a technique to perform artistic relighting. We show that our technique is able to preserve artist’s color-usage intention while relighting, generating coherent relit images (Figure 1.2). To the best of our knowledge, it is the first technique to address this scenario in image relighting.

Figure 1.2: Image relighting of a van Gogh’s painting: *Portrait de l’Artiste sans Barbe*. Note how the colors used by the artist to represent dark and bright shades are preserved in our relit images.




Our technique benefits from a wide range of freely and commercially available 3-D models, which can be used as proxies. Moreover, while most single-image relighting techniques are specific to face relighting (BLANZ; VETTER, 1999; WEN; LIU; HUANG, 2003; WANG et al., 2007, 2008), our method is very flexible, being applicable to objects with arbitrary shapes, as well as to non-photorealistic depictions, such as paintings and drawings.

Our technique can also be used for intrinsic-image decomposition, and for normal and depth estimation, even for paintings and drawings. Furthermore, our technique is capable of transferring shading information, as well as normal and depth maps from the *proxy* to flat drawings (e.g., outline drawings like Figure 1.3), something unreachable by state-of-the-art and commercial techniques. In those cases, our technique can recover approximate 3-D information from 2-D matched features.

The **contributions** of this thesis include:

- *A practical solution to the problem of single-image relighting of objects with arbitrary shapes (Section 3.1).* Our technique uses shading proxies and user interaction to guide the relighting process. It handles objects with non-uniform albedo and produces relit images whose visual quality is indistinguishable from the original ones;

Figure 1.3: Transferring shading and normal map to an outline drawing. Note that the original image does not hold any information about its geometry.



- *A technique to perform artistic relighting of nonphotorealistic images, such as paintings and drawings (Section 3.2).* It can preserve the artists' original use of colors to represent dark and bright shades;
- *A technique to estimate smooth normal and depth maps from single images.* Our technique can be applied to photographs, paintings, and drawings (Section 3.4). In the case of flat drawings (outlines), our method transfers normals from the *proxy* of the image;
- *A technique for intrinsic-image decomposition (Section 3.5).* Our approach is robust to albedo gradients, being able to disambiguate them from shading gradients;
- *A technique to transfer shading to flat drawings (Section 3.6).* Our approach can give 3-D flavor to outline drawings even when images do not hold any information about their geometry.

1.1 Thesis Structure

The remaining of this thesis is organized as follows: Chapter 2 discusses previous image relighting techniques, as well as methods for estimating illumination and geometry from images. It also discusses image composition techniques that take illumination into account. Chapter 3 presents our approach for image and artistic relighting. Additionally, it explains how the technique can be used for normal and depth-map recovery, intrinsic-image decomposition, and shading transferring. Chapter 4 shows our results, demonstrating the robustness and flexibility of our technique. Finally, Chapter 5 summarizes this thesis, and suggests some ideas for future work.

2 RELATED WORK

This chapter reviews the state-of-the-art techniques for image relighting (Section 2.1), and techniques to estimate illumination and geometry (Section 2.2). It also discusses image composition techniques that take into account lighting conditions (Section 2.3), showing that they are still not sufficient to generate realistic results.

2.1 Image Relighting Techniques

Image relighting has been addressed in different ways over the years. Essentially, existing techniques can be classified based on how they approach the problem. We discuss different techniques while classifying them as *geometry-based* or *image-based*. Geometry-based techniques use a 3-D representation/approximation for rendering an object under new illumination, while image-based methods rely on multiple images in order to create the relit image.

2.1.1 Geometry-Based Methods

The fundamental of inverse lighting problems (a specific example of inverse rendering (PATOW; PUEYO, 2003)) is to find what light/illumination could be provided to a renderer that would cause an output similar to the input image (MARSCHNER, 1998). Considering it is given a description of the scene (geometry and reflectance) and camera viewing, these techniques recover a rendering scene and, by setting the light conditions of the scene, render a new image. In order to acquire geometry and reflectance information, they rely on 3-D scanners (MARSCHNER, 1998; COSTA; SOUSA; NUNES FERREIRA, 1999), use multiple photographs taken under controlled settings (LOSCOS; DRETTAKIS; ROBERT, 2000), or create a virtual scene by manually placing 3-D models (BOIVIN; GAGALOWICZ, 2002). These techniques require access to the original scene and/or information about scene geometry and reflectance, which is hard to recover from a single image.

For **face relighting**, many methods use simple models, such as ellipses (BASSO et al., 2001) or a generic face model (WEN; LIU; HUANG, 2003). Others use morphable models to recover an approximated 3-D face representation (BLANZ; VETTER, 1999). These approximations are modeled based on an example set of 3-D face models, which are used in an optimization procedure in order to match the image. Later, Wang et al. combined morphable models and spherical harmonics to relight faces under harsh lighting conditions (WANG et al., 2007). Their method decouples the texture, geometry and illumination by modeling them separately on an iterative minimization procedure. Unfortunately, all these techniques are specific for faces and will not work in relighting images of differ-

ent classes.

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, which are then propagated, using a smooth flood-fill tool, to generate a normal map (OKABE et al., 2006). As this method propagates normals using gradients, it is heavily dependent on the *single-albedo assumption*, and cannot be used with images with non-uniform albedos.

In the area of hand drawn illustrations, many methods have been proposed for approximating lighting. Johnston proposed a technique that interpolates normals assigned to edge lines and, using a confidence matte, the method blends multiple normal approximations (JOHNSTON, 2002). Sýkora et al. proposed an interactive approach to add relative depth to cartoons (SÝKORA et al., 2010). Sýkora et al. proposed a pipeline in order to create a relief-type mesh from hand-drawn images (SÝKORA et al., 2014). The last two methods treat the regions in a layered fashion, testing occlusions and inflating the layers to achieve a 3-D model. These methods only work for cartoons and some specific kinds of sketch drawings where parts can be segmented and relations among them can be established. Our shading transferring (Section 3.6) is more flexible, as it modifies a 3-D model to match the sketch.

Teng and Cham introduced an example-based image relighting approach (TENG; CHAM, 2011). The method makes use of a database of reference images and, matching image intensities and gradients, it corresponds patches which have the same surface slant. Although this method generates good results for synthetic images, it is not suitable for natural images, as it relies on strong assumptions such as: uniform albedo, and patches to be matched must be spherical and convex surfaces. Karsch et al. proposed an interactive approach to create a physical model of the scene in order to include synthetic objects (KARSCH et al., 2011). However, their method focuses on rendering a synthetic object, whose geometry and material properties are known, which is not the case in image relighting.

The main problem of geometry-based methods is that recovering the shape of an arbitrary object from a single image is an ill-posed problem, and is especially difficult when dealing with non-uniform albedo and arbitrary shapes. Furthermore, when the approximation is not accurate, poor relighting results are generated.

2.1.2 Image-Based Methods

Debevec et al. proposed a method that uses a light-stage device to acquire a sparse set of viewpoints from a face under a dense set of lighting directions (DEBEVEC et al., 2000). This data is used to re-render the face under any combination of the captured light directions. Later, Tunwattapong et al. presented a method to reduce the number of required photographs (TUNWATTANAPONG; GHOSH; DEBEVEC, 2011). Malzbender et al.'s polynomial texture maps (MALZBENDER; GELB; WOLTERS, 2001) generate images of captured objects under new illumination by interpolating data from a set of photographs. Their method stores coefficients of a biquadratic polynomial for each texel, which are used to reconstruct surface colors given a light direction. All these techniques require a large number of images taken under controlled illumination.

Some relighting techniques create a mosaic combining separate parts of various images (AKERS et al., 2003; AGARWALA et al., 2004). Others, compute a linear combination of some basis images to achieve desired lighting effects (ANRYS; DUTRÉ, 2004; TUNWATTANAPONG; DEBEVEC, 2009). Another technique addresses the problem by capturing several photos with a moving light source (BOYADZHIEV; PARIS; BALA,

2013), and uses them in a photo editing software to generate results comparable with ones generated in a professional studio. However, these techniques require a set of images taken from the same viewpoint, and the range of possible results is limited by the captured images.

Ratio images have been used for image relighting in many techniques. *Quotient images* (SHASHUA; RIKLIN-RAVIV, 2001) are defined as the ratio of the surface albedos of two objects. In this work, Sashua et al. address the problem of "class-based" image relighting by capturing a set of images taken under linearly independent lighting conditions in order to create their *Quotient images*, from which a relit image can be created by combining the basis images. Peers et al. map one's face illumination to another using ratio images obtained from a set of pictures captured on a light-stage device (PEERS et al., 2007). Wang et al. proposed a method that uses infrared (IR) basis in order to relighting video conferences (WANG et al., 2008). Their method computes an optimally weighted combination of basis to create a ratio image that is used to relit the frames with the ideal illumination. Liu et al.'s method maps facial expressions using *expression ratio images* (ERI) (LIU; SHAN; ZHANG, 2001). Their technique also takes into account subtle changes in illumination and appearance, generating more expressive face results. The main difficulty of all techniques that use ratio images, besides the necessity of capturing images under specific lighting conditions, is the alignment [between the original and reference images](#).

Li et al. proposed a logarithmic total variation model to transfer illumination between faces (LI; YIN; DENG, 2009). Their method uses radial basis functions, and the relit image is obtained by replacing the illumination component of the target image by the reference one. Chen et al. (2011) proposed a method to transfer face illumination from a reference frontal face to a near-uniformly-illuminated target image. It uses edge-preserving filters to decompose the image into large-scale and detail layers, transferring the large-scale (that represents luminance variance) from reference to target image. In a subsequent work, Chen et al. (2012) used *shadow* and *light templates* copied by artists from professional portraits to perform face-illumination transfer. The templates were warped to fit the face in the input image, which is expected to contain no shadows or other significant lighting effect. Face-illumination transfer is then performed by multiplying the templates by the input image. This method requires a database of shadows and light templates created by artists, and the results are limited by the available templates. Additionally, the two last techniques only work for frontal faces.

Lopez-Moreno et al. presented a method to develop non-photorealistic rendering (NPR) techniques for input images (LOPEZ-MORENO et al., 2011). They show how relighting the original image using recovered 2.5-D depth information can convey new moods to the scene, controlling its emotional expressiveness. Although their recovered depth is not accurate for some applications, non-photorealistic stylization of images provides more forgiving ground, masking possible inconsistencies and making abstraction process unhampered.

2.2 Estimating Illumination and Geometry

This section discusses the most recent techniques for estimating illumination and geometry from images, discussing their main advantages and drawbacks. This section only describes techniques that work for single images under unknown conditions ([lighting conditions and material properties](#)), as others are too restrictive for our [intended application](#).

2.2.1 Estimating Illumination

The position of the light source is a valuable information about the scene illumination. Lalonde et al. proposed a method to estimate the position of the sun given a single outdoor image, by combining weak clues extracted from the image (LALONDE; EFROS; NARASIMHAN, 2012). Their method computes the probability distribution of the sun’s position by observing four clues: sky, cast shadows on the ground, vertical surfaces shadings, and pedestrian appearance. Although each clue is weak by itself, when combined they proved to be robust. Unfortunately, these clues are not available in most images.

Moreno et al. proposed a method capable of detecting multiple light sources and use it to recover a rough approximation of the object’s surface. They use it to perform image composition in a lighting-consistent way (LOPEZ-MORENO et al., 2010). The method uses an arbitrary object as a virtual light probe and uses a simple approach to recover a plausible depth map (KHAN et al., 2006). This technique, however, makes strong assumptions on the properties of the light-probe, such as having uniform albedo and consisting of a continuous convex surface.

Although the techniques just described are good for estimating illumination in certain circumstances, they cannot be generalized for different lighting conditions. Our goal is not to estimate precisely the original illumination, but to generate coherent relit results. Thus, we provide a mechanism of allowing the user to create plausible relightings interactively.

2.2.2 Estimating Geometry

In order to estimate the geometry of an object, some methods focus on recovering a normal map by using image gradients. Wu et al. proposed an interactive approach based on a simple markup procedure for normal reconstruction from single images (WU et al., 2007). They extend this method to automatically produces faithful normal reconstruction for high-frequency components, and offering an interactive markup procedure to correct low-frequency errors (WU et al., 2008). These methods suffer from the problem of *single-albedo assumption*, i.e., they fail when trying to estimate normal maps from images with non-uniform albedo.

Barron and Malik presented a unified model for recovering shape, chromatic illumination and reflection called SIRFS (BARRON, 2012). In a subsequent work (BARRON; MALIK, 2012), they extended this work to address achromatic illumination. Their method uses an optimization procedure to find which shape, illumination, and reflectance combined would generate the input image. The method uses priors for reflectance and illumination; and a multi-scale scheme is also proposed to speed up the optimization. This technique works really well for Lambertian materials with illumination encoded using spherical harmonics, but its effectiveness goes down for natural images.

Chen et al. proposed an interactive technique to manipulate simple 3-D shapes to extract 3-D objects from photographs (CHEN et al., 2013). It provides a tool (which they call *3-sweep*) to allow the user to explicitly define three dimensions of the primitive using three sweeps: two to define the 2-D profile and the last one to define the sweep path. Although many objects can be decomposed into simpler parts, there are lots of shapes that cannot be represented by this technique.

For our method, we decided to prioritize robustness over precision. We show how poor approximations of the geometry can still generate pleasing and consistent lighting effects, even for objects with complex geometry and non-Lambertian materials.

2.3 Image Composition

Combining multiples images into a single composition is a fundamental problem in image processing. Although many techniques focus on combining the colors efficiently (PÉREZ; GANGNET; BLAKE, 2003; JIA et al., 2006; YANG et al., 2009; SUNKAVALLI et al., 2010), when the illumination of the images differ considerably, unrealistic results are obtained. Recently, illumination-aware image composition has called the attention of several researchers.

Lalonde et al. proposed a data-driven approach to insert images of objects with given illumination in outdoors photographs with similar lighting conditions (LALONDE et al., 2007). Their method evaluates the illumination of the scene and, given a vast database of segmented objects (from images) with known illumination, it retrieves the ones with similar illumination. The user can then decide which objects to use for generating a composition. Although the method generates good results, it is not able to change the illumination of a particular object. Furthermore, when the illumination of the scene cannot be estimate properly, the method fails to find good matches.

Zhang et al. proposed an outdoors environment-sensitive image cloning technique that accounts for illumination (ZHANG; TONG, 2011). This method creates a reference image considering the sun's position, and combines this reference image with the Poisson method (PÉREZ; GANGNET; BLAKE, 2003). The idea of this method is that solving the Poisson equation will ensure a smooth transition in the patch boundaries while the reference image diffuses the sun's color. This method has several limitations: it combines the Poisson and reference image in an additive way, which will generate unrealistic results when inserting non-neutral colors objects; the sun must be visible in the image, and the color is diffused equally over the object, despite its shape.

Xue et al. proposed a method to adjust the appearance of an inserted foreground object so it looks compatible with the appearance of the background (XUE et al., 2012). The method matches zones of the histogram of different characteristics independently, using machine learning to automatically choose the zones. As this method only works directly over the histogram, it does not account for the geometry of the object, limiting the realism of the final compositing..

The main problem with these two last image-composition techniques is that they only perform 2-D operations. Thus, all the processing does not take into account the geometry/shape of the object: Zhang et al.'s technique diffuses the reference image equally despite the shape of the object, and Xue et al.'s manipulate colors independent of geometry.

3 IMAGE RELIGHTING USING SHADING PROXIES

Estimating illumination and geometry from single images is a difficult problem, especially when dealing with arbitrary objects with non-uniform albedo. As we discussed, several techniques narrowed the problem by relighting specific classes of objects (*e.g.*, faces), while others simplified the problem by making strong assumptions (full Lambertian surfaces, uniform albedo, restricted lightings conditions). We present an interactive technique that can be used by artists, photographers and even casual users in a transparent way to achieve different lighting effects on existing images. The technique does not require any previous information about the image and pleasing results can be achieved even by novice users.

In this chapter we present our image relighting method (Section 3.1) as well as our approach for relighting of paintings (Section 3.2). We also present a simple technique to [assist in the selection of 3-D models](#) (Section 3.3) and explain how our method can be used for normal and depth-map estimation (Section 3.4), intrinsic-image decomposition (Section 3.5), and shading transferring (Section 3.6).

3.1 Regular Relighting Method


Our technique uses *shading proxies* (*i.e.*, approximate 3-D models) and user interaction to achieve plausible relighting of arbitrary shapes. It consists of three steps: (i) image-model registration; (ii) creation of a feature-correspondence mapping; and (iii) actual relighting. These steps are detailed in the following sub-sections.

3.1.1 Image-Model Registration

Given an input image I containing a reference object O to be relit (*e.g.*, Figure 3.1(a)), we choose an approximate 3-D model representation M (Figure 3.1(b)) to create a shading proxy for O . M can be obtained from private databases, from the Internet, or from any free or commercial 3-D model repositories. Section 3.3 discusses how to identify suitable models for a given shading proxy.

Figure 3.1 illustrates the image-model registration process. 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 3.1(c)). Since M is just an approximate representation for O , the user can warp M to better match the image. This is achieved using Green coordinates (LIPMAN; LEVIN; COHEN-OR, 2008) associated with the vertices of the model's bounding box (illustrated as green vertices in Figures 3.1(c) and (d)). By [interactively](#) moving these vertices, M is automatically deformed. At the end of the registration process, the transformed and warped model re-

Figure 3.1: Image-model registration. (a) Input image. (b) Hat model that approximates the image in (a). (c) Initial registration by model transformation (translation, rotation, and scaling). (d) Warped model using Green coordinates. (e) Transformed and warped model.



sembles O (from the input image’s viewpoint) (Figure 3.1(e)) and is then called a **shading proxy**.

3.1.2 Feature-Correspondence Mapping


After performing image-model registration, one needs to establish a pixel-wise correspondence between the object to be relit in I and the image of the proxy (which we call the *reference image* R). For this, our technique matches a few key points and interpolates their positions to create a coherent feature-correspondence mapping. We start by matching silhouette and other relevant edges visible both in I and in R (Figure 3.2 (top)). The silhouettes edges from I and R are retrieved using the technique of Suzuki et al. (SUZUKI; BE, 1985). Both silhouettes are divided into segments defined by pairs of matched corners (retrieved using Shi and Tomasi’s method (SHI; TOMASI, 1994)). The user can decide to delete some of the automatically-retrieved segments and corners. In addition, one can interactively define pairs of corresponding segments between I and R . These new segments are obtained using a technique similar to *intelligent scissors* (MORTENSEN; BARRETT, 1995), enforcing segments to follow salient features in both images, one can also set different parameters to be able to define different segments. Figure 3.2 (top row) shows a color-coded representation for the corresponding pairs of feature segments defined for the hat example. Since the sizes of the corresponding segments are unlikely to match, all segments are parametrized in the $[0, 1]$ interval, so that we can easily match pairs of corresponding points.

Let S_R and S_I be the sets of feature segments from R and I , respectively (Figure 3.2 (top)). Also let P_R be a set of selected points from S_R , and let P_I be their corresponding points in S_I , obtained using the discussed parametrization. Currently, we select P_R as every third pixel on the feature segments in S_R . P_R and P_I are used to obtain Delaunay triangulations for the corresponding regions of interest of R and I (Figure 3.2 (middle)). These triangulations define a one-to-one mapping between both images, which we call **feature-correspondence mapping**. Triangles mapped outside the object’s mask are ignored.

3.1.3 Shading-Ratio Image and Actual Relighting

An image can be expressed as the product of illumination (*i.e.*, *shading*) and reflectance (*i.e.*, *albedo*) (BARROW; TENENBAUM, 1978). Thus, in our method, we perform image relighting by retrieving the object’s albedo and multiplying it by the new shading. This is equivalent to multiplying the original image by a pixel-wise shading ratio:

Figure 3.2: Feature-correspondence mapping and actual relighting. (top) Reference and input images with color-coded corresponding feature lines. (middle) Delaunay triangulations obtained from the corresponding feature lines. (bottom) Relituminated proxy (left) and corresponding relit image (right).




$$I_r(u, v) = I(u, v) \left(\frac{S_{new}(u, v)}{S_I(u, v)} \right), \quad (3.1)$$

where I and I_r are the original and relit images. S_I and S_{new} are instances of the reference image R rendered using the original and target lighting, respectively. We call S_I and S_{new} the *original shading image* and the *new shading image*, respectively. $S_{ratio}(u, v) = (S_{new}(u, v)/S_I(u, v))$ is the *shading-ratio image*, which is related to Shashua and Riklin-Raviv's *quotient image* (SHASHUA; RIKLIN-RAVIV, 2001). One should note, however, that a quotient image is defined as the ratio of the surface albedos of two objects and, therefore, is independent of illumination. A shading-ratio image, on the other hand, is defined as the ratio of two shadings obtained for the same object/scene pose.

Figure 3.3 illustrates the construction of shading-ratio images and the actual relighting process. Figure 3.3(a) shows an input image I , whose (registered) proxy is shown in (c). Figure 3.3(b) shows a patch p_I of I (the girl's face) selected for relighting. The selection of regions of interest for relighting is described later in the section. Figures 3.3 (b) and (c) show the color-coded correspondence segments defined by the user for features inside the face. Together with the face contour lines (hidden for simplicity), they define the 2-D Delaunay triangulations shown in Figures 3.3(e) and (f). The shading-ratio image is created by rendering the shading proxy under two lighting conditions (manually defined by the user): S_I , which tries to mimic the shading of O in I , and is shown in (c), and S_{new} , the target lighting for O , shown in (d). Although shading proxies only approximate the objects to be relit, image-model misregistrations are compensated for by the feature-correspondence mapping that defines a homeomorphism between the Delaunay

Figure 3.3: Steps of the image-relighting process. (a) Input image. (b) Face selected as the region of interest for relighting with color-coded salient features. (c) and (d) Shading proxies with original and target shading. (e) and (f) face patch and reference image with corresponding Delaunay triangulations. (g) and (h) S_I and S_{new} shading images. (i) Shading-ratio image. (j) Relit image.




triangulations. Moreover, the 2-D mesh associated with the reference image (shown in (f)) is warped to exactly match the 2-D mesh defined for p_I (shown in (e)). As a result, the shading images S_I and S_{new} have the same shape as the patch p_I , and are shown in (g) and (h), respectively. The shading-ratio image (i) is then obtained by dividing the *new shading image* (Figure 3.3 (h)) by the *original shading image* (Figure 3.3(g)). The resulting shading-ratio image is then used to perform pixel-wise modulation by p_I , producing the relit image showing in (j). Note how it gracefully incorporates the target shading.

3.1.4 Shading-Ratio Inpainting

The feature-correspondence mapping avoids image-model misregistration artifacts. However, the occurrence of high frequencies in the shading-ratio image may introduce artifacts when mapped to smooth regions of the input image. Such situations tend to result from proxy features not found in the input image or from topological issues in the 3-D model's construction. One should note that high-frequencies in the input image are not a problem, and such features will be preserved by the relighting process.


We avoid these artifacts by removing/smoothing gradients in S_{ratio} that are perceived as undesirable when transferred to the relit image. We do so by performing a *shading-ratio inpainting*. Given a relit image, the user indicates the presence of such artifacts using scribbles. Our system then performs inpainting on the corresponding regions of the shading-ratio image using the technique described in (BERTALMIO; BERTOZZI; SAPIRO, 2001). Figure 3.4 illustrates this process. Figure 3.4 (left) shows small artifacts in the girl's mouth, eyes, and nose. By marking these artifacts directly on the relit image (green scribbles in Figure 3.4 (right)), the system inpaints the corresponding regions of

Figure 3.4: Shading-ratio inpainting. Artifacts in the shading-ratio image (gray inset on the bottom left) are transferred to the relit image, as seen in the girl’s mouth, eyes, and nose (left). Using scribbles (shown in green), one marks these artifacts directly on the relit image. The system then performs inpainting on the shading-ratio image, fixing the problem and avoiding the artifacts (right). The final rendering is shown in Figure 3.5(c).



S_{ratio} (Figure 3.4 (right)), avoiding the artifacts. Figures 3.5 (b) and (c) show examples of relit images obtained from (a). In (b), the image was relit from below. In (c), light comes from the left. These results were obtained using the inpainted shading-ratio image shown in Figure 3.4 (right).


Figure 3.5: Relighting of a girl’s face (color pencil drawing) from the below and from the left.



3.1.5 Selecting Regions of Interest for Relighting

As our technique focus on relighting objects of an image, one can specify one or more regions of the input image to be relit. For this, we use a scribble-based interface based on (CLIPPINGMAGIC, 2014) and illustrated in Figure 3.6. Green scribbles indicate

Figure 3.6: Interface for selecting regions of interest for relighting. Green scribbles indicate regions of interest and red scribbles delimit them. In the end of *Image-Model Registration*: Newton’s face overlaid on the shading proxy (d).




regions of interest, while red scribbles delimit them (Figure 3.6(a)). The result is the selected patch and associated alpha mask (Figure 3.6 (b), (c)). The alpha mask is used as a stencil for limiting the relighting region on the original image and for smoothly transitioning between the new and old shadings at the borders of the patch. Figure 3.6 (d) shows the extracted patch of Newton’s face and neck overlaid on the shading proxy.

3.2 Artistic Relighting


Photographs and other photorealistic images can be successfully relit with Eq. (3.1), as shown in Figures 1.1, 3.2, 4.4, and 4.5. In the case of paintings and drawings, however, artists often use color, as opposed to lightness, to encode shading. For instance, in the *Portrait de l’Artiste sans Barbe* shown in Figure 3.7 (a), van Gogh used proportionally more red in in-shade areas than in lit ones, which are yellowish. Thus, directly applying Eq. (3.1) to this image would produce less pleasing results, as it would tend to over-

Figure 3.7: Regular vs artistic relighting. (a) Original painting with regions used for artistic relighting (lit region in blue and shadow region in green). (b) painting relit from the left using Eq. (3.1). Red shades, originally used to represent dark appear very bright in lit areas. (c) Artistic relighting using Eq. (3.2) . (d) Shading proxy under desired illumination.



stress the use of some colors and mix two different shading techniques (van Gogh’s and Eq. (3.1)). This situation is illustrated in Figure 3.7 (b), which shows the painting relit from the left using Eq. (3.1). The red shades that have been used to represent dark, now appear very bright in the lit areas. Figure 3.7 (c) shows the result obtained using our artistic-relighting technique with the same illumination. Note the red shades are back to darker regions. The corresponding shading proxy under the target illumination is shown on the right.

Figure 3.8: Artistic shading-color interpolation. Colors of source (r_β) and target (t_β) shadings are interpolated between the in-shade (S_β) and lit (L_β) colors, according to the values in S_I and S_{new} . Note how each color channel scales differently given a shading intensity (x -axis).



To mimic the artist’s color-usage intention, we correlate shading independently with each color channels, *i.e.*, the *shading-ratio image* will not be multiplied equally for all color channels. For this, one indicates an in-shade and a lit region by directly pointing on the painting image at p_s and p_l , respectively (Figure 3.7(a)). Let S_R , S_G , and S_B be the average R, G, and B values computed at a 10×10 pixel neighborhood N_s around p_s . Likewise, let L_R , L_G , and L_B be the R, G, and B average values computed at a similar neighborhood N_l around p_l . Recall that the *original shading image* S_I approximates the lighting of the object in the input image I (see Figure 3.3 (g)). The relit image I_r is computed separately for each color channel $\beta \in \{R, G, B\}$ as

$$I_{r_\beta}(u, v) = I_\beta(u, v) \left(\frac{t_\beta(u, v)}{r_\beta(u, v)} \right), \quad (3.2)$$

where

$$t_\beta(u, v) = S_\beta + (S_{new}(u, v) - \mathcal{I}_s) \text{slope}_\beta, \quad (3.3)$$


$$r_\beta(u, v) = S_\beta + (S_I(u, v) - \mathcal{I}_s) \text{slope}_\beta, \quad (3.4)$$

and

$$\text{slope}_\beta = \frac{L_\beta - S_\beta}{\mathcal{I}_l - \mathcal{I}_s}. \quad (3.5)$$

$\mathcal{I}_s$ and $\mathcal{I}_l$ are the average shading intensities in S_I at the neighborhoods corresponding to N_s and N_l , respectively. Equations 3.2 to 3.5 are explained in Figure 3.8. Starting from

Figure 3.9: Relighting of a painting of Isaac Newton by Godfrey Kneller. (left) Original painting. (center and right) Results produced by our artistic relighting technique. (bottom) The shading proxies and corresponding illumination.



S_β and L_β , and from $\mathcal{I}_s$ and $\mathcal{I}_l$, t_β and r_β are interpolated along the lines defined by S_β and L_β according to the values of $S_{new}(u, v)$ and $S_I(u, v)$, respectively. The ratio t_β/r_β is then used to relight the β channel of the region of interest in I . Figures 1.2, 3.5, 3.9, and 4.3 illustrate the use of our artistic relighting technique.


3.3 Selecting 3-D Models for Shading Proxies

Our shading proxies can be selected from a variety of sources, and many repositories are available on the Internet. Recently, web-based search engines have been developed to help users find 3-D models online (FUNKHOUSER et al., 2003; MIN, 2004). Such engines take 2-D and 3-D sketches, 3-D models, and keywords as input, and could help our users find models suitable for shading proxies. In addition, we have developed a simple and effective solution to help one finds appropriate models in local 3-D repositories.

Let $M = \{m_1, m_2, \dots, m_k\}$ be a set of 3-D models, which have been organized into classes (e.g., cars, airplanes, houses, etc.). The repository is prepared for future queries using a pre-processing step. For each model m_i , we render a set of views covering a full hemisphere around m_i . We render these views as binary images, so that image $b_{i\theta\phi}$ defines a visibility mask for model m_i as seen from direction (θ, ϕ) on the hemisphere. For each image $b_{i\theta\phi}$, we compute its seven Hu moments (HU, 1962) and use them as descriptors instead of the actual images, which are discarded. We have chosen the Hu image moments because these moments are invariant to scale, translation, and orientation. [These moments have been largely used for classifying objects in images because of their robustness over](#)

image affine transformations.

Figure 3.10: Suggesting models for shading proxies. Given an input image with an object to be relit (top left), the object’s mask (top right) is extracted using a scribble-based interface (see Figure 3.6 (left)). The system computes the mask’s Hu moments, compares them against the moments stored in the database using Eq. (3.6), and returns an array of thumbnail images with the eight most similar 3-D models (bottom). For this Lamborghini example, the query returned three Lamborghini models.



Given an input image I with an object to be relit (Figure 3.10 (top left)), a binary mask b_I (top right) is extracted using our scribble-based interface. The user selects the class of objects to be queried (in this example, *cars*), and the system then computes the mask’s Hu moments, uses them to query the repository, and returns the eight best matches (Figure 3.10 (bottom)). Similarity is measured using (OPENCV, 2014)

$$Sim_{index} = \sum_{k=1}^7 \left| \frac{1}{m_k^{b_I}} - \frac{1}{m_k^{b_{i\theta\phi}}} \right|, \quad (3.6)$$

where

$$m_k^\alpha = \text{sign}(h_k^\alpha) \log h_k^\alpha,$$

and h_k^α is the k -th Hu moment of image α (from a total of seven). Smaller values of Sim_{index} indicate higher similarity. Figure 3.10 (bottom) shows the eight models returned for the Lamborghini picture shown on the top left. Note it contains three Lamborghini models. [As these moments are invariant to affine transformations, our method retrieves the models whose masks are most similar to the input image, despite differences in scale, position or orientation.](#)

Figure 3.11: Examples of recovered normal maps from paintings and drawing. (left) Input image; (center) Normal map recovered using our method (all proxies derive from the same 3-D model); (right) Normal maps created using a commercial software (CRAZYBUMP, 2014).



3.4 Normal and Depth-Map Recovery

The feature-correspondence mapping generated by our method can also be useful for recovering a roughly geometric information of an imaged object. For instance, it makes it straightforward to transfer information from the shading proxy, via reference image R , to the input image. We exploit this possibility to retrieve normal and depth maps from pictures. Figure 3.11 shows examples of normal maps obtained using our method. Note that conventional techniques for estimating normal maps from images do so based on image gradients and, therefore, would not work for these examples (Figure 3.11 (right)). Note how commercial softwares fail to recover normal maps from paintings and flat drawings. Figure 3.12 brings additional examples for normal maps estimations.

Our approach is capable of retrieving smooth normal maps that capture the essence of each of these images, even though the proxies used to retrieve them derive from the same 3-D model (all faces in Figure 3.11). Furthermore, our method can also transfer normal maps to outline drawings, something unreachable by state-of-the-art techniques. The

Figure 3.12: Examples of recovered normal maps from drawings and photographs. (left) Input image; (center) Normal map recovered using our method; (right) Normal maps created using a commercial software (CRAZYBUMP, 2014).





Figure 3.13: Examples of recovered depth maps for the paintings and drawing shown in Figure 3.11 (top) and Figure 3.12 (bottom).




observations of normal-map recovering can also be applied to depth maps (Figure 3.13).

3.5 Intrinsic-Image Decomposition

Besides image relighting and normal and depth-map estimation, our technique is capable of intrinsic-image decomposition as well. Using the retrieved *original shading image* S_I by our technique, we can recover the albedo just by dividing the image by the shading. In fact, our technique is robust, once that the main challenge of state-of-the-art intrinsic-

Figure 3.14: Comparing intrinsic-image decomposition. (a) Input image. (b) Recovered albedo and shading using our method. (c) Recovered albedo and shading using Shen et al. method (2008). (d) Recolored and relit image using our method.



decomposition methods is to distinguish albedo edges from shading ones (SHEN; TAN; LIN, 2008; BOUSSEAU; PARIS; DURAND, 2009). As our reference images only contain shading values, our method does not suffer this problem. In fact, for simple objects, our method can achieve comparable decompositions (Figure 3.14 shows our method comparing with another (SHEN; TAN; LIN, 2008)). Furthermore, using our method for intrinsic decomposition, we are able to modify the color of an object independently from its shading (BEIGPOUR; WEIJER, 2011), and relight it as well (Figure 3.14 (d)).

For more complex shapes, the quality of the intrinsic decomposition will depend on how well the geometry of the used proxy matches the geometry of the imaged object, and on how well the shading applied to it matches the actual scene shading (illumination). Note that one can replace our simple OpenGL's point light sources with more sophisticated lighting models (*e.g.*, PBRT (PHARR; HUMPHREYS, 2010)). Even when our technique fails to directly perform intrinsic decomposition, the approximate geometry and shading might provide useful information that could be used for disambiguation between albedo and shading edges, which would be difficult to obtain otherwise (SINHA; ADELSON, 1993).

Figure 3.15: Transferring shading from proxies to 2-D drawings. (left) Original drawings. (center) Shadings transferred from a proxy, lit from the left. (right) Colored shadings from a proxy (using van Gogh's artistic interpolation information).



3.6 Shading Transferring

Another application is to transfer shading from proxies to outline drawings. By just applying the *new shading image* (not using the *ratio image*), our technique is able to give 3-D flavor to flat outline images (Figure 3.15). Furthermore, by using the information retrieved using the *Artistic Relighting* (Section 3.2), we can give color to the shading. In Figure 3.15 (right) we used the color information retrieved from Figure 3.7, note how the lit and shade regions of the colored shading follow the same color intuition used in the painting of van Gogh.

There are some techniques that propose to relight cartoons and hand drawings ((JOHNSTON, 2002; SÝKORA et al., 2010, 2014)). Although they do not need any 3-D model in priori, these methods work in a "layered" fashion, and they cannot handle effects as the ones demonstrated in Figure 3.15.

4 RESULTS

We have implemented the techniques described in this thesis using C++, OpenGL and Qt, with the help of OpenCV, armadillo and MATLAB functions. We used OpenGL to illuminate the shading proxies using up to four point light sources. We present many examples using a single light source, as well as multiple light sources. We also show that, by manipulating virtual lights, we can even change the light color, achieving interesting results.

We used our approach to re-illuminate a large number of images including photographs, paintings, and drawings. A typical relighting section takes from 2 to 10 minutes (in the case of more complex shapes), even for novice users (Section 4.1). All measurements were made on a 3.4 GHz i7 CPU with 16 GB of RAM.

Figure 4.1: Relighting of Steve McCurry’s photograph *Afghan Girl* using artistic and regular relighting. (left) Photograph relit from above. (right) Photograph relit from below. Results on top were generated using artistic relighting (Section 3.2) and bottom using our regular method. Note that the different approaches generate distinct results.



Figures 1.1 (b) and (c) show two relit versions of Steve McCurry’s photograph *Afghan Girl*, shown on the left. In (b), the girl’s face has been relit from above, while in (c) the light comes from the bottom right. The resulting images convey very distinct impressions, creating new portraits with unique moods. They were produced using Eq. (3.1), a single point light source, and no shading-ratio inpainting. Additional relighting examples are shown in Figure 4.1. The results on the top row were produced using Eq. (3.2). At the bottom, one sees another examples of image relit with Eq. (3.1). In (left) the results are relit from above and in (right) they are relit from below. Note that our artistic relighting can also be applied to photographs, not being limited to paintings.

Figures 1.2 (b) and (c) show versions of van Gogh’s *Portrait de l’Artiste sans Barbe* relit from the left, from the bottom left, and from above. Note how the red and yellow shades used by the artist to indicate dark and light areas are properly preserved in the new renditions. For these results, we used our artistic-relighting technique with a single light source, and no shading-ratio inpainting. Figure 3.7 compares the result of re-illuminating a painting using regular relighting (Eq. (3.1)) and artistic-relighting (Eq. (3.2)). Figure 3.7 (d) shows the shading proxy used to relight van Gogh’s face in both examples. Figures 3.5, 3.9, and 4.3 provide additional examples of artistic relighting.

Figure 4.2: Relighting of a pencil drawing. (left) Original. (center) Relighting from above. (right) Relighting using two light sources.




Figures 3.5 (b) and (c) show relit versions of the color pencil drawing of a girl’s face shown in (a). In (b), the light is coming from the bottom, while in (c) it comes from the left. For these exquisite renderings, a single light source was used with the shading proxy shown in Figure 3.3(c). Shading-ratio inpainting was used in parts of the mouth, eyes, and nose, as shown in Figure 3.4 (right). Figures 4.2 show the same pencil drawing, but relighting from above (center) and using two light sources (right), from the left and from bottom right.

Figure 3.9 shows the relighting of a Newton’s portrait from the right and from the front, using a single light source. In the bottom: the shading proxy and the corresponding illuminations. The region of interest used for Newton’s portrait relighting is shown in Figure 3.6 (center) and we used the Artistic Relighting (Section 3.2) to achieve coherent results.

Figure 4.3 shows the relighting of Tim Benson’s portrait *Dad Looking Down*, a painting with a strong painting style. The inpainting areas are shown in Figure 4.3 (b). The image (c) shows the relighting from the bottom left using a bright light source. On (d) and (e), the image has been relit, respectively, from the top right and from above using a

Figure 4.3: Relighting of Tim Benson’s portrait “Dad Looking Down”. (a) Original. (b) Areas for shading-ratio inpainting. (c) Intense relighting from bottom left. (d) Relighting from top right. (e) Relighting from top left. (f) Relighting using two light sources.



single light source. Figure 4.3 (f) shows the result of relighting using two light sources (one from bottom right and another from left). The insets in images show the relit shading proxys used. Note it is the same 3-D model, used for the relighting of van Gogh’s and Newton’s portraits. In each case, however, the model has been transformed and warped to fit the individual characters. By transforming and warping 3-D models to create shading proxies, our approach provides great flexibility, allowing us to handle a large number of images using a relatively small number of models.

Since our shading proxies can be any 3-D model, our technique can relight images of objects with arbitrary geometries (Figures 1.1, 3.2, 4.4, and 4.5). Figure 3.2 shows the relighting of a hat.

Figure 4.4 illustrates the relighting of a Lamborghini *Aventador*. The original photograph is shown on the top left. Its inset shows the shading proxy: the model of a Lamborghini *Murciélag*o – note the differences in the shapes of the front bumpers. The top-right image shows the Lamborghini relit from the left, the bottom-left relit from the

Figure 4.4: Relighting a car photograph from different directions. The original is on the top left, with the shading proxy as an inset.




Figure 4.5: Relighting of two colored bottles. (left) Photograph. (center) Bottles relit from the left. (right) Bottles relit from the right. The insets show the shading proxies and target illuminations.




bottom and the bottom-right show the result of relighting from the right; all using a single light source. Some shading-ratio inpainting was required in the front bumper region. The resulting images look quite realistic despite the strong highlights on the image.

Figure 4.5 shows the relighting of two bottles with non-uniform albedo. For this example we used two independent shading proxies, but a single light source, and no inpainting. The proxies are shown as insets and illustrate the lighting conditions. Figure 4.5 (left) is the original photograph, while in (center) and (right) the painted bottles were relit from the left, and from the right, respectively. For simple models like these, the automatic silhouette correspondences are sufficient to generate the correspondence map.


Figure 4.6 shows the relighting of a woman's face in a Botticelli's painting shown in (a). Some face-relighting techniques are restricted to frontal faces (CHEN et al., 2011, 2012). Our technique can produce expressive and realistic relighting of face profiles, as shown in Figure 4.6 (b), (c) and (d). Note the shadings at the cheek and near the eye, as the light comes from below (b) and from above (c), drawing attention to different areas of the face. These delicate renderings were obtained using Eq. (3.1), a single light source, and no inpainting. By using an additional light source, different results can be achievable Figure 4.6 (d).

Figure 4.6: Relighting of a Botticelli's painting. (a) Original. (b) Face relit from below. (c) Face relit from above. (d) Face relit with two light sources.



As already mentioned, high-frequencies are preserved by our relighting process. Figure 4.7 shows the relighting of Eric Daigh's painting "Allyson". Figure 4.7 (a) shows the original painting, while (b) and (c) are results relit from left and right respectively, and the proxy is shown in (d). The results were produced using a single light source and no-inpainting. Note how the dots are preserved by the relighting procedure.


Figure 4.7: Relighting of an Eric Daigh's painting. (a) Original. (b) Face relit from left. (c) Face relit from right. (d) Proxy. Note how the painting's high-frequency content is preserved by the relighting process.



Other example containing high-frequencies is shown in Figure 4.8. The original image in Figure 4.8(a) shows a work made by Andrew Myers using screws. Figure 4.8b shows the result relit from the left and in (c) relit from the right. We used artistic relighting to


achieve realistic relightings and, despite the use of inpainting in some areas, preserve the details in the input image.

Figure 4.8: Relighting of a screw work of Andrew Myers. (a) Original. (b) Portrait relit from left. (c) Portrait relit from right. Note how our technique is robust to images with high-frequencies details, being able to generate pleasing relit results.



Many methods for face relighting use morphable models, based on base face models, to approximate the geometry of the face (BLANZ; VETTER, 1999; WANG et al., 2007). However, poor results are generated when the imaged face is very different from the basis. Figure 4.9 shows a painting by Modigliani, who is known for portraits depicting elongated faces. Figure 4.9 (a) shows the original painting. The image in (b) shows painting relit from right, while (c) shows the painting relit from below. Figure 4.8 (d) shows the same painting relit using two light sources (above and on the right). The proxy is shown in (e). No inpainting was used to generate these results.

Figure 4.9: Relighting of Modigliani's painting *Portrait de Jeanne Hébuterne*. (a) Original painting. (b) Face relit from right. (c) Face relit from below. (d) Face relit using two light sources. (e) Proxy.



Caricature relighting is also difficult for techniques based on morphable models. Figure 4.10 shows a caricature of Locke (Lost) made by Mark Hammermeister. Although we have used the same male head 3-D model than others examples (Figure 3.11), our feature-correspondence map is capable of warping the proxy to fit the image. Figure 4.10 (a) shows the original caricature. The results in (b) is the caricature relit from right, while

(c) and (d) show the image relit from above from bottom left, respectively. Note how the correspondence mapping is capable of deforming the proxy to fit the image.

Figure 4.10: Relighting of a caricature of Locke made by Mark Hammermeister. (a) Original caricature. (b) Caricature relit from right. (c) Caricature relit from above. (d) Caricature relit from bottom left.




Figure 4.11 shows our method relighting a black-and-white photograph of a juggler taken by Mike Cipriano. The image on (top left) shows the original photograph, the one on (top right) is the result relit from above, while results in (bottom left) and (bottom right) are relit from right using a low intensity light source and relit from right, respectively.


Figure 4.11: Relighting of a black-and-white photograph of a juggler taken by Mike Cipriano. (top left) Original. (top right) Photograph relit from above. (bottom left) Photograph relit from right with a low intensity light source. (bottom right) Photograph relit from right.



Figure 4.12 shows a black-and-white photograph of Albert Einstein. We show that, similar to transferring shading and color (Section 3.6), it is possible to encode color while


relighting. Figure 4.12 (a) shows the original photograph, while (b) and (d) are results relit from right and from below, respectively. Figure 4.12(c) and (e) are relit from the same position that (b) and (d), respectively, but using the shading-color interpolation function from van Gogh’s painting (Figure 3.7).

Figure 4.12: Relighting of a black-and-white photograph of Albert Einstein. (a) Original. (b) Photograph relit from right. (c) Colored face relit from right. (d) Photograph relit from below. (e) Colored face relit from below. The colored results were generated using artistic interpolation function from van Gogh’s painting.



Our method can also be used to achieve artistic results to enhance images. Figure 4.13 shows a nautilus shell painting, by Tracey Costescu, relit from many angles. By relighting this painting, it is possible to enhance geometric features. Figure 4.13(a) shows the original painting, while (b),(c) and (d) are results relit from above, bottom-right, and from below, respectively. Note that the proxy (Figure 4.13(d)) does not exactly fit the input image, yet, the results look nice and pleasing.

Figure 4.13: Relighting of a Tracey Costescu’s painting of a Nautilus shell. (a) Original. (b) Painting relit from above. (c) Painting relit from bottom right with a low-intensity light. (d) Painting relit from below. (e) Proxy. In these examples, relighting seems to enhance depth perception.



We have shown how our method can add virtual light sources to generate distinct results. Additionally, one can change the color of each light source. Figure 4.14 shows the relighting of a skull, (a) is the original image (with proxy shown as an inset). The skull relit from right using a white light source is shown in (b), while (c) shows the result relit from the same position but using a blue light source. Figure 4.14(d) combines two light sources (blue from right and red from left) in the relighting. As the imaged skull and the proxy have several differences, some inpainting was used to avoid artifacts and generate pleasing relit results.

Figure 4.14: Relighting of a skull. (a) Original image. (b) Skull relit from right with a white light source. (c) Skull relit from same position as (b) but using a blue light source. (d) Skull relit by two light sources with different colors (blue from right and red from left).




Figure 4.15: Relighting of a black vase by manipulating independently each channel of the shading-ratio image. (a) Original photograph. (b) Vase relit from top right. (c) Vase relit from bottom right. (d) Vase relit from front left. (e) Vase relit from top right. Results (d) and (e) were generated by manipulating each channel of the shading-ratio image.



Furthermore, by manipulating each channel of the shading-ratio image we can generate different effects. Figure 4.15 shows a black vase being relit. The original photograph is shown in (a), (b) and (c) show the vase relit from top right and bottom right, respectively. Figure 4.15(d) and (e) are results where each channel in the shading-ratio image was scaled by a different factor. We tried to simulate a change in the material's spectral properties. In this case, the vase appears as made of cooper, in (d) relit from front left and in (e) relit from top right.


We can manipulate the *new shading image* S_{new} using image processing filters to achieve peculiar effects. For example, Figure 4.16 shows an effect produced by using color quantization on the *new shading image*. Note that this is not the same as performing the quantization direct on the relit result (Figure 4.16 (right)). Figure 4.16 (top) shows this effect using our regular relighting technique, Figure 4.16 (bottom) shows the effect of using artistic relight. The left column shows the relit images. The middle column shows the results obtained by performing quantization on the *new shading image* using four shading levels. Note how the relit images acquire a cartoon-like appearance. In contrast, the right column of Figure 4.16 shows the results obtained when applying quantization directly to the relit images shown on the left column.

We can filter the *new shading image* S_{new} to obtain other effects. Figure 4.17 shows

Figure 4.16: Relighting effect by performing color quantization in *new shading image*. (left) Relit version. (middle) Color quantization on the shading. (right) Color quantization the relit version (left). Results in (top) using the regular relighting and (bottom) with the artistic relighting.



Figure 4.17: Relighting effect by applying edge-preserving filters on the *new shading image* (shadings as insets). (a) Regular relighting. (b) Applying detail enhance filter on the shading. (c) Applying sketching filter. (d) Applying stylization filter.



the result of applying the domain transform's edge-preserving filters (GASTAL; OLIVEIRA, 2011) to achieve detail enhancement (Figure 4.17(b)), pencil sketching (Figure 4.17(c)), and stylization (Figure 4.17(d)). Figure 4.17 shows the relit images. The modified *new shading images* are shown as insets. Figure 4.18 shows the same effects applied in con-

Figure 4.18: Relighting effect by applying edge-preserving filters on the *new shading image* in our artistic relighting (modified shading as insets). (a) Artistic relighting. (b) Applying detail enhance filter on the shading. (c) Applying sketching filter. (d) Applying stylization filter.



junction with our artistic relighting. Similar to color quantization, applying these filters on the shading image is not the same as applying them on the relit image. Thus, our technique provides a way to apply specific filters directly on the image's shading, generating effects not achievable by other techniques.

4.1 User Evaluation

We performed a simple experiment to test the usability of our technique by novice users. For this, we asked 10 subjects (all graduate students from our lab, 8 males, 2 females, ages ranging from 23 to 29 years old) to perform two image relighting tasks. First, the subjects were given a brief explanation of how our system works, and were asked to practice it by relighting a cap (Figure 4.19). The average practice time was eight minutes and thirty-six seconds. During this period, the subjects were allowed to ask questions and request assistance. Each subject decided on his/her own when to stop practicing. The subjects were then asked to perform two relighting tasks all by themselves. The first task involved relighting the hat shown in Figure 3.1(a) using the 3-D model in Figure 3.1(b). The second consisted of relighting a human face (Figure 4.21) using the 3-D model shown in (b). For this experiment, we decided not to specify a target result. Its goal was to allow the volunteers to try distinct lighting effects on images, producing results they would find pleasing. Table 4.1 shows the average times taken by the subjects on each part of the relighting process.

Figure 4.19: Image and 3-D model used for novice training.



Figures 4.20 (center) and (right) show examples of relit hat images created by two of

Figure 4.20: Relighting examples produced by novice users. (left) Original image. (center) Relit by one subject in 172 seconds. (right) Relit by another subject in 137 seconds.



Table 4.1: Average times taken by the novices in the training.

Task	Average times (in minutes)			
	Image Registration	Feature-Correspondence	Relighting	Total
Cap	02:36	02:57	03:03	08:36
Hat	02:52	02:29	02:25	07:46
Face	03:11	04:24	02:14	09:49

our subjects. The original is shown on the left. The time required by these novice users to create the images were 2 minutes and 52 seconds (0:00:59 for image-model registration, 0:01:10 for establishing feature correspondences, and 0:00:43 for light positioning and inpainting), and 2 minutes and 17 seconds (0:01:06 for image-model registration, 0:00:33 for establishing feature correspondences, and 0:00:38 for light positioning and inpainting), respectively. These numbers show that our technique is quite intuitive and can be easily learned by novices.

Figures 4.21 (b) and (c) show examples of relit face images created by two of our subjects. The model given is shown in (d). The time required by the users to relit the images were 6 minutes and 30 seconds (0:00:43 for image-model registration, 0:04:45 for establishing feature correspondences, and 0:01:02 for light positioning and inpainting), and 10 minutes and 50 seconds (0:03:12 for image-model registration, 0:03:36 for establishing feature correspondences, and 0:04:02 for light positioning and inpainting), respectively. All subjects reported that the process will get more efficient and effective the more you familiarize with the interface.

We have also performed a second experiment to evaluate the quality of the results produced by our techniques. This experiment involved a group of 42 individuals consisting of professional photographers, artists, and general volunteers, ages ranging from 20 to 39 years old. They were asked to perform two on-line evaluation sessions, each consisting of ten questions. The volunteers performed the evaluations at their convenience, with no time limit, on their own computers. For the first session, volunteers were shown a series of ten randomly selected images from a pool containing both original images and images relit with our technique. For each image, they were asked whether they think the lighting in the image is plausible or not. For the second session, volunteers were shown groups of four images at a time, arranged in a 2 by 2 grid. Each group consisted of an original image and three relit versions of it. The order of the images in the grid was randomly defined. Volunteers were then asked to: (a) choose their favorite image in the group (i.e., the most aesthetically pleasing one); and (b) identify the original image. The results from


the first session indicate that the relit images were considered to have plausible illumination in 70% of the cases, whereas for the original images the illumination was considered plausible in 88% of the cases. Differences between object and scene illumination were the main reason for considering lighting non-plausible in the relit images. Moreover, besides photographs, the set of test images also contained several paintings and drawings, whose non-photorealistic depictions make relighting harder and should have contributed to lower the rate of perceived plausible lighting among them. The results of the second session show that the volunteers preferred a relit version over the original one in 51.15% of the cases. Even when volunteers correctly guessed the original image (which happened in 72.14% of the cases), they still preferred our relighting results in 51.32% of the cases. In almost 30% of the cases, volunteers were not able to identify the original images. These numbers indicate that our approach does indeed produce pleasing re-illuminations, which are considered as good as the original ones.

4.2 Discussion and Limitations

Our work tackles the difficult problem of relighting objects with arbitrary shapes taking as input only the image itself and the target illumination. No explicit information about the object/scene geometry and its original shading is provided. Our solution bridges this information gap offering a practical solution to this hard problem. Although Eq. (3.1) suggests that one needs to match S_I , the input image’s shading, on the proxy, in practice this is not really required and excellent results can be obtained even with poor approximations. This is so because the relighting process is driven by real-time feedback on the relit image. Thus, one can obtain the desired effect by compensating on the target shading, S_{new} . In this sense, our technique is quite robust.

Despite the need for 3-D models, the use of shading proxies makes the previously-mentioned difficult problems tractable. Currently, there is a wide range of both freely and commercially available 3-D models, and it is relatively easy to find models for essentially anything. We emphasize that the availability of large 3-D databases is not necessary. As we have shown, the warping and correspondence mapping provide enough power to create different proxies from single models. For instance, a single 3-D model was used to create proxies for all male faces shown in this thesis, regardless of their representa-

Figure 4.21: Relighting face examples produced by novice users. (a) Original face. (b) Relit by one subject in 390 seconds. (c) Relit by another subject in 650 seconds. (d) 3-D model given for the task.



tions (photographs, paintings, drawings). Thus, a small set of 3-D models should suffice for most cases. The user could incrementally build such a database according to his/her needs. Moreover, as the availability of 3-D models continues to increase and specialized web-based search engines (e.g., (FUNKHOUSER et al., 2003)) evolve, finding the desired/approximate models is expected to become even easier.

Although many techniques can interactively estimate geometry without the need of a model, they are not robust enough, failing when dealing with images containing non-uniform albedo, specular highlights, and brush and pencil strokes (in case of paintings and drawings). The use of proxies and feature-correspondence allows our technique to overcome these limitations. For instance, we have shown that despite the existence of brush strokes in Figure 1.2, strong painting style in Figure 4.3, specular effects on Figure 4.4, non-uniform albedo in Figure 4.5, and high-frequency effects in Figure 4.7 and Figure 4.8, our technique can relight the images coherently.

We have used OpenGL’s point light sources because they lend to real-time rendering and pleasant results. Our system also supports OpenGL’s ambient term. Even though we have not included a specular term in our simple lighting model, nothing precludes its use. Nevertheless, smooth highlights are still achievable depending on the object’s reflectance (Figure 4.22). While we have based our decision on performance considerations, more sophisticated lighting models can be used with our technique.


Figure 4.22: Relighting Oprah’s portrait. (left) Original. (center) Relit with two light sources: one at each side of the face. (right) Relit from right with a brighter light. Note the smooth highlights.



Since a proxy is used to transfer its shading to the relit image, ideally the proxy itself should not include (high-frequency) details that have no correspondence in the input image. Otherwise, such details might affect the shading ratio and be transferred to the relit image as artifacts. In this case, one should use the shading-ratio inpainting feature to remove the associated gradients from S_{ratio} , thus avoiding the artifacts. Smooth shading ratios never introduce artifacts in the relit images.

Defining correspondences among features in the input and in the reference images gives the user additional freedom to achieve specific relighting effects. For instance, consider the situation depicted in Figure 4.23 (left). On top, we show the Afghan girl’s face overlaid on its shading proxy. At the bottom, one sees the proxy with two color lines: a red one following the top-of-the-head contour, and a blue one crossing the middle of the forehead. By establishing the correspondence of the girl’s forehead edge with each one of the color lines, one obtains distinct results for the same light position, as shown in Figure 4.23 (right). The image on top uses the red contour for correspondence, while the one at the bottom uses the blue line. Since the red contour includes the top of the proxy’s

Figure 4.23: Defining feature correspondence provides additional freedom to specify re-lighting effects. (top right) Assigning the girl’s forehead edge to the red contour on the proxy causes her forehead to look brighter. (bottom right) Assigning the girl’s forehead edge to the blue line causes her forehead to look darker. The light source is at the same position (above) in both cases.



head, for a light source placed above, it allows for brighter shading on the top of the girl’s forehead (top right). In contrast, the flatter proxy’s forehead lends to a darker shading (bottom right).

Our technique has some **limitations**. If the object to be relit has unusual geometric features, a suitable 3-D model for use as a shading proxy might not be immediately available. Also, since our artistic relighting uses colors to represent dark and bright shades, it is sensitive to the choice of such colors. Our current implementation does not remove shadows (*e.g.*, shadow cast by the Lamborghini on the floor, in Figure 4.4). This might be addressed using shadow mapping (WILLIAMS, 1978), and we would like to explore it. Furthermore, objects that are hard to model (*e.g.*, clothes, hair) and materials with distinct BRDFs (*e.g.*, fur) are hard to relight, but this might be addressed using normal maps in the rendering of the proxy.

5 CONCLUSIONS AND FUTURE WORK

This thesis presented a practical solution to the problem of single-image relighting of objects with arbitrary shapes. Our approach uses shading proxies and user interaction to guide the relighting process. It works by computing a shading-ratio image which is used to map the input lighting condition to a target illumination. Our solution is flexible and robust, being applicable to arbitrary objects composed of non-uniform albedos. Our method is the first one to perform artistic relighting, supporting the specification of colors to represent dark and bright shades. It can relight paintings while preserving the artists' original color intention, producing exquisite images.

In addition, the thesis has demonstrated a simple and effective solution for identifying suitable 3-D models in local repositories for use as shading proxies. We demonstrate how our technique can recover smooth normal and depth maps from single images. The resulting maps capture the essence of the original images, even in the case of non-photorealistic paintings and drawings. Furthermore, our technique can be used for intrinsic-image decomposition, being able to disambiguate between albedo and shading gradients. Our technique can also be used to transfer shading to flat drawings (*e.g.*, outline drawings) and even given color to this.

The effectiveness of our techniques was demonstrated by performing real-time relighting and normal and depth-map estimation on a large number of photographs, paintings, and drawings. An evaluation performed with novice users has shown that our approach is intuitive and simple to use. After a brief explanation and practice session, subjects were able to relight objects on their own in just a couple of minutes.

Given its flexibility, robustness, and ease of use, our technique can help artists, photographers, and casual users to experiment with new lighting effects on existing images. For instance, images can be *re-targeted* to draw the user's attention to specific details. This can be observed in Figure 4.6, where the relit images draw the viewer's attention to different portions of the woman's face. In Figures 1.1 (a) to (c), the various lighting effects induce different perceptions. Thus, we believe that our technique will stimulate new and creative image-relighting applications.

5.1 Future Work

Handling shadows cast by scene objects is important aspect of image relighting, which has not been covered by this thesis. This is an interesting direction for future exploration.

Furthermore, the idea of using synthetic images to transfer information to real images can be used in a more general way. As image processing filters only works in 2-D, we can create more complex filters, as synthetic images can give us information about depth, normals and other geometric informations. Similar to the work of Liu et al. that maps

facial expressions using ratio images (LIU; SHAN; ZHANG, 2001), we would like to map geometric modifications in a consistent way.

REFERENCES

- AGARWALA, A. et al. Interactive Digital Photomontage. **ACM Transactions on Graphics**, New York, NY, USA, v.23, n.3, p.294–302, 2004.
- AKERS, D. et al. Conveying Shape and Features with Image-Based Relighting. In: IEEE VISUALIZATION, 14., 2003, Washington, DC, USA. **Proceedings...** IEEE Computer Society, 2003. p.46–51. (VIS '03).
- ANRYS, F.; DUTRÉ, P. Image-based lighting design. In: IASTED INTERNATIONAL CONFERENCE ON VISUALIZATION, IMAGING, AND IMAGE PROCESSING, 4., 2004, Marbella, Spain. **Proceedings...** [S.l.: s.n.], 2004. p.78–89.
- BARRON, J. T. Shape, Albedo, and Illumination from a Single Image of an Unknown Object. In: IEEE CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION, 2012, Washington, DC, USA. **Proceedings...** IEEE Computer Society, 2012. p.334–341. (CVPR '12).
- BARRON, J. T.; MALIK, J. Color Constancy, Intrinsic Images, and Shape Estimation. In: EUROPEAN CONFERENCE ON COMPUTER VISION, 2012, Berlin, Heidelberg. **Proceedings...** Springer Berlin Heidelberg, 2012. p.57–70. (ECCV '12, v.7575).
- BARROW, H. G.; TENENBAUM, J. M. **Recovering Intrinsic Scene Characteristics from Images**. [S.l.]: Academic Press, 1978. 3–26p.
- BASSO, A. et al. Virtual light: digitally-generated lighting for video conferencing applications. In: IEEE INTERNATIONAL CONFERENCE ON IMAGE PROCESSING, 2001. **Proceedings...** IEEE, 2001. p.1085–1088. (ICIP '01, v.2).
- BEIGPOUR, S.; WEIJER, J. van de. Object Recoloring Based on Intrinsic Image Estimation. In: IEEE INTERNATIONAL CONFERENCE ON COMPUTER VISION, 2011. **Proceedings...** IEEE, 2011. p.327–334. (ICCV '11).
- BERTALMIO, M.; BERTOZZI, A.; SAPIRO, G. Navier-stokes, fluid dynamics, and image and video inpainting. In: IEEE COMPUTER VISION AND PATTERN RECOGNITION, 2001. **Proceedings...** IEEE, 2001. p.355–362. (CVPR '01, v.1).
- BLANZ, V.; VETTER, T. A Morphable Model for the Synthesis of 3D Faces. In: ACM SIGGRAPH, 1999, New York, NY, USA. **Proceedings...** ACM Press/Addison-Wesley Publishing Co., 1999. p.187–194. (SIGGRAPH '99).

BOIVIN, S.; GAGALOWICZ, A. Inverse rendering from a single image. In: IEEE COMPUTER GRAPHICS, IMAGING AND VISUALIZATION, 2002. **Proceedings...** IEEE, 2002. p.268–277. (CGIV '02).

BOUSSEAU, A.; PARIS, S.; DURAND, F. User-assisted Intrinsic Images. **ACM Transactions on Graphics**, New York, NY, USA, v.28, n.5, p.130:1–130:10, 2009.

BOYADZHIEV, I.; PARIS, S.; BALA, K. User-assisted Image Compositing for Photographic Lighting. **ACM Transactions on Graphics**, New York, NY, USA, v.32, n.4, p.36:1–36:12, July 2013.

CHEN, T. et al. 3Sweep: extracting editable objects from a single photo. **ACM Transactions on Graphics**, New York, NY, USA, v.32, n.6, p.195:1–195:10, 2013.

CHEN, X. et al. Face Illumination Transfer Through Edge-preserving Filters. In: IEEE COMPUTER VISION AND PATTERN RECOGNITION, 2011, Providence, RI, USA. **Proceedings...** IEEE, 2011. p.281–287. (CVPR '11).

CHEN, X. et al. Artistic Illumination Transfer for Portraits. **Computer Graphics Forum**, New York, NY, USA, v.31, n.4, p.1425–1434, 2012.

CLIPPINGMAGIC. <http://clippingmagic.com/>.

COSTA, A.; SOUSA, A.; NUNES FERREIRA, F. Lighting Design: a goal based approach using optimisation. In: EUROGRAPHICS WORKSHOP ON RENDERING, 1999, Aire-la-Ville, Switzerland, Switzerland. **Proceedings...** Eurographics Association, 1999. p.317–328. (EGWR '99).

CRAZYBUMP. **Demo Version 1.2**. www.crazybump.com/.

DEBEVEC, P. et al. Acquiring the Reflectance Field of a Human Face. In: ACM SIGGRAPH, 2000, New York, NY, USA. **Proceedings...** ACM Press/Addison-Wesley Publishing Co., 2000. p.145–156. (SIGGRAPH '00).

FUNKHOUSER, T. et al. A search engine for 3d models. **ACM Transactions on Graphics**, New York, NY, USA, v.22, n.1, p.83–105, 2003.

GASTAL, E. S. L.; OLIVEIRA, M. M. Domain Transform for Edge-Aware Image and Video Processing. **ACM Transactions on Graphics**, New York, NY, USA, v.30, n.4, p.69:1–69:12, 2011.

HU, M.-K. Visual pattern recognition by moment invariants. **IRE Transactions on Information Theory**, [S.l.], v.8, n.2, p.179–187, 1962.

JIA, J. et al. Drag-and-drop pasting. **ACM Transactions on Graphics**, New York, NY, USA, v.25, n.3, p.631–637, July 2006.

JOHNSTON, S. F. Lumo: illumination for cel animation. In: ACM NON-PHOTOREALISTIC ANIMATION AND RENDERING, 2002, New York, NY, USA. **Proceedings...** ACM, 2002. p.45–ff. (NPAR '02).

KARSCH, K. et al. Rendering synthetic objects into legacy photographs. In: ACM SIGGRAPH ASIA CONFERENCE, 2011., 2011, New York, NY, USA. **Proceedings...** ACM, 2011. p.157:1–157:12. (SA '11).

KHAN, E. A. et al. Image-based Material Editing. In: ACM SIGGRAPH, 2006, New York, NY, USA. **Proceedings...** ACM, 2006. p.654–663. (SIGGRAPH '06).

LALONDE, J.-F.; EFROS, A. A.; NARASIMHAN, S. G. Estimating the Natural Illumination Conditions from a Single Outdoor Image. **International Journal of Computer Vision**, [S.l.], v.98, n.2, p.123–145, 2012.

LALONDE, J.-F. et al. Photo Clip Art. **ACM Transactions on Graphics (SIGGRAPH 2007)**, New York, NY, USA, v.26, n.3, p.3, August 2007.

LI, Q.; YIN, W.; DENG, Z. Image-based Face Illumination Transferring Using Logarithmic Total Variation Models. **The Visual Computer**, [S.l.], v.26, n.1, p.41–49, 2009.

LIPMAN, Y.; LEVIN, D.; COHEN-OR, D. Green Coordinates. **ACM Transactions on Graphics**, New York, NY, USA, v.27, n.3, p.78:1–78:10, 2008.

LIU, Z.; SHAN, Y.; ZHANG, Z. Expressive Expression Mapping with Ratio Images. In: ACM SIGGRAPH, 2001, New York, NY, USA. **Proceedings...** ACM, 2001. p.271–276. (SIGGRAPH '01).

LOPEZ-MORENO, J. et al. Compositing images through light source detection. **Computers and Graphics**, [S.l.], v.34, n.6, p.698–707, 2010.

LOPEZ-MORENO, J. et al. Non-photorealistic, depth-based image editing. **Computers & Graphics**, [S.l.], v.35, n.1, p.99 – 111, 2011.

LOSCOS, C.; DRETTAKIS, G.; ROBERT, L. Interactive Virtual Relighting of Real Scenes. **IEEE TVCG**, [S.l.], v.6, n.3, p.289–305, 2000.

MALZBENDER, T.; GELB, D.; WOLTERS, H. J. Polynomial texture maps. In: ACM SIGGRAPH, 2001, New York, NY, USA. **Proceedings...** ACM, 2001. p.519–528. (SIGGRAPH'01).

MARSCHNER, S. R. **Inverse Rendering for Computer Graphics**. 1998. Tese (Doutorado em Ciência da Computação) — Cornell University.

MIN, P. **A 3D Model Search Engine**. 2004. Tese (Doutorado em Ciência da Computação) — Princeton University.

MORTENSEN, E. N.; BARRETT, W. A. Intelligent Scissors for Image Composition. In: ACM SIGGRAPH, 1995, New York, NY, USA. **Proceedings...** ACM, 1995. p.191–198. (SIGGRAPH'95).

OKABE, M. et al. Single-view relighting with normal map painting. In: PACIFIC GRAPHICS, 2006, Taipei, Taiwan. **Proceedings...** [S.l.: s.n.], 2006. p.27–34.

OPENCV. **Structural Analysis and Shape Descriptor - matchShapes**. OpenCV 2.4.8.0 Documentation.

PATOW, G.; PUEYO, X. A Survey of Inverse Rendering Problems. **Computer Graphics Forum**, [S.l.], v.22, n.4, p.663–688, 2003.

PEERS, P. et al. Post-production Facial Performance Relighting Using Reflectance Transfer. **ACM Transactions on Graphics**, New York, NY, USA, v.26, n.3, p.52:1–52:10, 2007.

PÉREZ, P.; GANGNET, M.; BLAKE, A. Poisson image editing. **ACM Transactions on Graphics**, New York, NY, USA, v.22, n.3, p.313–318, July 2003.

PHARR, M.; HUMPHREYS, G. **Physically Based Rendering, Second Edition**: from theory to implementation. 2nd ed. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2010.

SHASHUA, A.; RIKLIN-RAVIV, T. The Quotient Image: class-based re-rendering and recognition with varying illuminations. **IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)**, [S.l.], v.23, n.2, p.129–139, 2001.

SHEN, L.; TAN, P.; LIN, S. Intrinsic image decomposition with non-local texture cues. In: IEEE COMPUTER VISION AND PATTERN RECOGNITION, 2008, Anchorage, AK, USA. **Proceedings...** IEEE, 2008. p.1–7. (CVPR'08).

SHI, J.; TOMASI, C. Good features to track. In: IEEE COMPUTER VISION AND PATTERN RECOGNITION, 1994, Seattle, WA, USA. **Proceedings...** IEEE, 1994. p.593–600. (CVPR'94).

SINHA, P.; ADELSON, E. H. Recovering reflectance and illumination in a world of painted polyhedra. In: IEEE INTERNATIONAL CONFERENCE ON COMPUTER VISION, 1993, Berlin, DE. **Proceedings...** IEEE, 1993. p.156–163. (ICCV'93).

SUNKAVALLI, K. et al. Multi-scale image harmonization. **ACM Transactions on Graphics**, New York, NY, USA, v.29, p.125:1–125:10, July 2010.

SUZUKI, S.; BE, K. Topological structural analysis of digitized binary images by border following. **Computer Vision, Graphics, and Image Processing**, [S.l.], v.30, n.1, p.32 – 46, 1985.

SÝKORA, D. et al. Adding Depth to Cartoons Using Sparse Depth (In)equalities. **Computer Graphics Forum**, [S.l.], v.29, n.2, p.615–623, 2010.

SÝKORA, D. et al. Ink-and-ray: bas-relief meshes for adding global illumination effects to hand-drawn characters. **ACM Transactions on Graphics**, New York, NY, USA, v.33, n.2, p.16:1–16:15, Apr. 2014.

TENG, X.; CHAM, T.-J. Image Relighting by Analogy. In: INTERNATIONAL SYMPOSIUM ON VISUAL COMPUTING, 2011. **Proceedings...** Springer, 2011. p.78–89. (ISVC'11, v.6938).

TUNWATTANAPONG, B.; DEBEVEC, P. **Interactive Image-Based Relighting with Spatially-Varying Lights**. ACM SIGGRAPH'09 Poster.

TUNWATTANAPONG, B.; GHOSH, A.; DEBEVEC, P. Practical Image-Based Relighting and Editing with Spherical-Harmonics and Local Lights. In: IEEE CONFERENCE FOR VISUAL MEDIA PRODUCTION, 2011, London, UK. **Proceedings...** IEEE, 2011. p.138–147. (CVMP'11).

WANG, O. et al. Video Relighting Using Infrared Illumination. **Computer Graphics Forum**, Crete, Greece, v.27, n.2, p.271–279, 2008.

WANG, Y. et al. Face Re-Lighting from a Single Image under Harsh Lighting Conditions. In: IEEE COMPUTER VISION AND PATTERN RECOGNITION, 2007, Minneapolis, MN, USA. **Proceedings...** IEEE, 2007. p.1–8. (CVPR'07).

WEN, Z.; LIU, Z.; HUANG, T. S. Face Relighting with Radiance Environment Maps. In: IEEE COMPUTER VISION AND PATTERN RECOGNITION, 2003, Madison, WI, USA. **Proceedings...** IEEE, 2003. p.158–165. (CVPR'03).

WILLIAMS, L. Casting Curved Shadows on Curved Surfaces. In: ACM SIGGRAPH, 1978, New York, NY, USA. **Proceedings...** ACM, 1978. p.270–274. (SIGGRAPH'78).

WU, T.-P. et al. ShapePalettes: interactive normal transfer via sketching. **ACM Transactions on Graphics**, New York, NY, USA, v.26, n.3, p.44:1–44:5, 2007.

WU, T.-P. et al. Interactive Normal Reconstruction from a Single Image. **ACM Transaction on Graphics**, New York, NY, USA, v.27, n.5, p.119:1–119:9, 2008.

XUE, S. et al. Understanding and improving the realism of image composites. **ACM Transactions on Graphics**, New York, NY, USA, v.31, n.4, p.84:1–84:10, July 2012.

YANG, W. et al. Natural and seamless image composition with color control. **IEEE Transactions on Image Processing**, Piscataway, NJ, USA, v.18, n.11, p.2584–2592, Nov. 2009.

ZHANG, Y.; TONG, R. Environment-Sensitive cloning in images. **The Visual Computer**: international journal of computer graphics, Secaucus, NJ, USA, v.27, n.6-8, p.739–748, June 2011.

APÊNDICE A REILUMINAÇÃO DE IMAGENS UTILIZANDO *SHADING PROXIES*

Resumo da Dissertação em Português

Reiluminação de imagens tenta recriar a aparência de uma imagem sob novas condições de iluminação. Isto é útil quando uma foto não pode ser (facilmente) recapturada sob as condições de iluminação desejadas, quando deseja-se chamar a atenção em determinada área da imagem, ou para alcançar diferentes percepções. Entretanto, reiluminar uma única imagem é um problema complexo, uma vez que esta não possui informação explícita sobre sua iluminação original, nem sobre a geometria dos objetos ou propriedades de seus materiais. Apesar de algumas técnicas estimarem iluminação de cenas sob determinadas circunstâncias (LOPEZ-MORENO et al., 2010; LALONDE; EFROS; NARASIMHAN, 2012), não existe uma solução genérica para este problema. Além disso, recuperar a geometria de objetos dado uma única imagem é um problema mal condicionado, o qual não existe solução com resultados satisfatórios para formas arbitrárias (WU et al., 2008; BARRON; MALIK, 2012; CHEN et al., 2013).

Esta dissertação apresenta uma solução prática para o problema de reiluminação de imagens com objetos de geometria arbitrária. Ao invés de recuperar a geometria do objeto a ser reiluminado, nós utilizamos o que chamamos de *shading proxy* (uma aproximação geométrica do objeto, a qual é interativamente deformada para aproximar o objeto a ser reiluminado). Nós estabelecemos correspondências entre a imagem do objeto e a imagem do *proxy*. Dessa forma, ao reiluminarmos o *proxy* com a iluminação original e a iluminação desejada, podemos calcular a *shading-ratio image*, a qual é utilizada para uma reiluminação pixel a pixel do objeto. Na prática, bons resultados são obtidos mesmo com aproximações grosseiras da iluminação original. Isto só é possível devido ao feedback em tempo real proporcionado por nossa técnica, fazendo com que o usuário possa modificar as condições de iluminação até obter o efeito desejado.

Nossa técnica se beneficia de uma grande gama de modelos 3D disponíveis, os quais podem ser utilizados como *proxies*. Nós mostramos que, enquanto a maioria das técnicas de reiluminação de imagens são específicas para faces (BLANZ; VETTER, 1999; WEN; LIU; HUANG, 2003; WANG et al., 2007, 2008), nosso método pode reiluminar objetos de geometria arbitrária, bem como representações não-fotorrealísticas, como pinturas e desenhos. Além disso, nosso método pode ser utilizado para estimar mapas de normais e profundidade, mesmo que as imagens sejam pinturas e desenhos. Neste caso, nossa técnica é capaz de inferir informação 3D a partir de *features* 2D correlacionadas.

Nós mostramos como nossa técnica é robusta e flexível, podendo reiluminar obje-

tos com formas distintas, bem como representações não-fotorrealísticas. Ademais, mostramos como nossa técnica pode ser utilizada para decomposição intrínseca de imagens, transferência de iluminação, e recuperação de mapas de normais e de profundidade.

As **contribuições** desta dissertação incluem:

- *Uma solução prática para o problema de reiluminação de imagens de objetos com geometria arbitrária.* A técnica utiliza *shading proxies* e interação com o usuário para guiar o processo de reiluminação. Nosso método é robusto e flexível, produzindo resultados de alta qualidade;
- *Uma técnica para realizar reiluminação artística de imagens não-fotorrealísticas.* Ela é capaz de preservar o uso original de cores do artista para representar regiões escuras e claras;
- *Uma técnica para estimar mapas de normais e de profundidade suaves.* A técnica pode ser aplicada em fotografias, pinturas, e desenhos. No caso de desenhos achatados (delineados), nosso método é capaz de transferir informações de normais e de profundidade do *proxy* para a imagem;
- *Uma técnica para decomposição intrínseca de imagens.* Ela pode, de maneira robusta, distinguir gradientes de iluminação e gradientes de albedo;
- *Uma técnica para transferir iluminação para desenhos planos.* A técnica pode atribuir uma impressão 3D para esboços, mesmo quando a imagem não contém informação explícita da geometria.

A.1 Trabalhos Relacionados

Técnicas de reiluminação de imagens podem ser basicamente classificadas em: técnicas baseadas em geometria (as quais utilizam representações e aproximações 3D para guiar a reiluminação) e técnicas baseadas em imagens (as quais buscam combinar duas ou mais imagens para gerar o resultado reiluminado). Nesta seção discutimos as técnicas existentes, mostrando como estas não são robustas, nem flexíveis, o suficiente para serem utilizadas na maioria das imagens.

Métodos baseados em geometria

Técnicas de *inverse lighting* buscam, dado um descritor geométrico da cena, encontrar e modificar sua iluminação. Para adquirir esta geometria, alguns métodos utilizam scanners 3D (MARSCHNER, 1998; COSTA; SOUSA; NUNES FERREIRA, 1999), múltiplas fotos capturadas sob iluminação controlada (LOSCOS; DRETTAKIS; ROBERT, 2000), ou criação manual de uma cena virtual (BOIVIN; GAGALOWICZ, 2002). Estas técnicas necessitam de acesso à cena original e/ou informação adicional da geometria e reflectância da cena, o que é difícil de se obter dada uma única imagem.

Para reiluminação de faces, alguns métodos utilizam modelos simples, como elipsoides (BASSO et al., 2001) ou modelos genérico de faces (WEN; LIU; HUANG, 2003). Outros utilizam modelos que se deformam para obter uma aproximação mais refinada (BLANZ; VETTER, 1999; WANG et al., 2007). Entretanto, estas técnicas são específicas para faces.

Algumas técnicas permitem que o usuário especifique ou reconstrua um mapa de normais. Estes mapas podem ser construídos utilizando interação do usuário (OKABE et al.,

2006), ou atribuindo, de maneira automática, normais específicas às arestas (JOHNS-TON, 2002; SÝKORA et al., 2010, 2014). Métodos baseados em recuperação de mapa de normais são altamente dependentes da suposição de que os objetos possuem um albedo uniforme, ou seja, não podem ser utilizados em objetos com albedos peculiares.

Métodos baseados em imagens

Debevec et al. utiliza um dispositivo chamado *Light-stage* para capturar um conjunto esparsos de pontos de vista de uma face sob um conjunto denso de direções de iluminação (DEBEVEC et al., 2000). Estas imagens são utilizadas para renderizar a face sob qualquer combinação das direções de iluminação capturadas. Tunwattanapong et al. apresenta um método que reduz o número de fotografias necessário da técnica previamente descrita (TUNWATTANAPONG; GHOSH; DEBEVEC, 2011). Malzbender et al. gera imagens reiluminadas através da interpolação de dados adquiridos de um conjunto de fotos (MALZBENDER; GELB; WOLTERS, 2001). Todas estas técnicas necessitam de um grande conjunto de imagens capturadas sob condições específicas de iluminação.

Algumas técnicas de reiluminação criam um mosaico combinando diferentes partes de várias imagens (AKERS et al., 2003; AGARWALA et al., 2004). Outras, calculam uma combinação linear de imagens capturadas para obter a iluminação desejada (ANRYS; DUTRÉ, 2004; TUNWATTANAPONG; DEBEVEC, 2009). Estas técnicas necessitam de imagens capturas sob o mesmo ponto de vista, e a variedade de reiluminações possíveis é limitada pelas imagens capturadas.

Um conceito bastante conhecido em reiluminação de imagens é utilizar razões de imagens (SHASHUA; RIKLIN-RAVIV, 2001) para mapear mudanças na iluminação. Peers et al. mostra como utilizar este conceito juntamente com o *light-stage* (PEERS et al., 2007). Wang et al. mostra como utilizar razão de imagens infra-vermelhas para reiluminar vídeo conferências (WANG et al., 2008). O método de Liu et al. utiliza razões de imagens para mapear expressões faciais, utilizando o que chamam de *expression ratio images* (ERI) (LIU; SHAN; ZHANG, 2001). O principal problema destas técnicas, além da necessidade de imagens sob condições específicas de iluminação, é o alinhamento entre elas.

Das técnicas específica para faces, Li et al. propõem um modelo logarítmico para recuperar a componente de iluminação e transferi-la entre faces (LI; YIN; DENG, 2009). O método de Chen et al. utiliza filtros que preservam arestas para decompor a imagem em camadas de detalhes, para então transferir a camada de iluminação de uma face a outra (CHEN et al., 2011). Em outro trabalho, Chen et al. utiliza *templates* criados por artistas para realizar a reiluminação (CHEN et al., 2012). As duas últimas técnicas somente funcionam em faces frontais.

A.2 Reiluminação de Imagens

Estimar as condições de iluminação e geometria dado uma única imagem é um problema difícil, especialmente quando lidamos com objetos arbitrários com albedo não uniforme. As técnicas apresentadas simplificam o problema reiluminando classes específicas (e.g., faces), enquanto outras fazem suposições muito fortes acerca das propriedades dos materiais e condições de iluminação (superfícies Lambertianas, albedo uniforme, condições arbitrárias de iluminação). Nós apresentamos uma técnica interativa que pode ser usada para modificar a iluminação de imagens, obtendo efeitos específicos. Nossa técnica não necessita de nenhuma informação prévia e bons resultados podem ser alcançados mesmo por usuários novatos.

Nossa técnica utiliza *shading proxies* e interação com o usuário para alcançar uma iluminação plausível de objetos com formas arbitrárias. Nossa técnica consiste em 3 etapas: (i) registro do modelo, (ii) criação de um mapa de correspondência, e (iii) reiluminação. A seguir detalhamos cada um destes.

Registro do modelo: dado uma imagem contendo o objeto que deseja-se reiluminar, encontra-se um modelo 3D para criar o nosso *shading proxy*. Este modelo é transformado (utilizando transformações como escala, translação e rotação) e deformado (utilizando Green Coordinates (LIPMAN; LEVIN; COHEN-OR, 2008)) de maneira que o resultado renderizado aproxime o melhor possível o visualizado na imagem. O modelo deformado final é o que chamamos de *shading proxy*, e é utilizado para guiar o processo de iluminação.

Criação de um mapa de correspondência: após o registro do modelo, é necessário estabelecer uma correspondências pixel a pixel entre a imagem do objeto e a imagem do *proxy*. Para isso, nossa técnica utiliza pontos chave e interpola as posições para criar um mapeamento coerente. Nós começamos estabelecendo relações entre as silhuetas das duas imagens, as quais são automaticamente obtidas (SUZUKI; BE, 1985) e segmentadas nos cantos (SHI; TOMASI, 1994). Além disso, o usuário pode interativamente definir pares de segmentos utilizando uma ferramenta similar ao *intelligent scissors* (MORTENSEN; BARRETT, 1995). Estes segmentos são então utilizados para obter uma triangulação de Delaunay das imagens, definindo um mapeamento entre elas; nós chamamos este mapeamento de *feature-correspondence mapping*.

Reiluminação: "uma imagem pode ser expressa como o produto da iluminação (*shading*) e seu albedo (*reflectance*)" (BARROW; TENENBAUM, 1978). Dessa forma, reiluminar uma imagem nada mais é do que recuperar o albedo do objeto e aplicar a iluminação desejada. Isto é equivalente a multiplicar a imagem por uma razão entre a iluminação desejada, sobre a iluminação original da imagem. Dessa forma, utilizando o mapeamento definido anteriormente, é possível calcular essa razão utilizando imagens sintéticas do *proxy* sob as duas iluminações (definidas manualmente pelo usuário). Nós chamamos esta razão de *shading-ratio image* e, multiplicando ela pela imagem de entrada, é possível obter o resultado reiluminado. Note que a multiplicação pode ser realizada ponto-a-ponto, uma vez que o *feature-correspondence mapping* define um homeomorfismo entre as triangulações de Delaunay. É importante notar que o feedback em tempo real proporcionado pela técnica possibilita ao usuário manipular as condições de iluminação até obter o efeito desejado, fazendo com que não seja necessário uma aproximação extremamente precisa da iluminação original.

A.2.0.1 *Shading-Ratio Inpainting*

Ao reiluminar uma imagem, artefatos podem aparecer quando altas frequências na *shading-ratio image* são mapeadas para regiões suaves da imagem. Tais situações ocorrem quando características do *proxy* não estão presentes na imagem do objeto, ou de problemas topológicos na construção do modelo 3D. Para solucionar este problema, nossa técnica possibilita a realização de *inpaint* sobre a *shading-ratio image*, de forma a eliminar/suavizar os gradientes indesejáveis. Dessa forma, o usuário pode indicar manualmente quais regiões apresentam artefatos indesejados, para que nossa técnica possa eliminá-los, aumentando consideravelmente a qualidade dos resultados gerados.

A.2.1 Reiluminação Artística

Enquanto fotografias e imagens fotorrealísticas podem ser reiluminadas de maneira eficaz utilizando a técnica descrita acima; pinturas e desenhos são um problema, uma vez que artistas normalmente utilizam cores, ao invés de luminosidade, para codificar iluminação. Dessa forma, aplicando diretamente o método descrito anteriormente, os resultados não seriam coerentes. Para simular a intenção do uso de cores do artista, nós propomos que cada canal de cor seja tratado de maneira diferente, fazendo com que exista uma *shading-ratio image* independente para cada um. Para realizar isso, nossa técnica necessita que o usuário indique uma região de sombra e uma região iluminada. Com isso, nossa técnica é capaz de criar uma relação para cada canal de cor que mimetiza o padrão de iluminação do artista, ou seja, pode-se identificar como cada canal de cor vai se comportar ao aumentarmos ou diminuirmos a iluminação. Isso faz com que ao reiluminarmos a pintura/desenho, a coerência de cores seja mantida, fazendo com que os resultados sejam mais coerentes. Pelo que sabemos, esta é a primeira técnica a tratar este tipo de problema em reiluminação de imagens.

A.2.2 Decomposição Intrínseca de Imagens

Além de reiluminação, nosso método ainda pode ser utilizado para decomposição intrínseca de imagens. Utilizando o mapa de correspondências e a iluminação original determinada pelo usuário, é possível dividir a imagem de entrada por esta iluminação e recuperar seu albedo. Apesar desta técnica não ser extremamente precisa, nós mostramos como ela pode ser utilizada para decompor objetos simples, enquanto que, para objetos mais complexos, a informação de iluminação dada por nosso método pode ser usada para desambiguar gradientes de albedo de gradientes de iluminação.

A.2.3 Transferência de Iluminação

Outra aplicação é a transferência de iluminação do *shading proxy* para esboços. Se ignorarmos a iluminação original, e só utilizarmos a iluminação desejada, é possível aplicá-la diretamente na imagem utilizando o mapa de correspondências, fazendo com que desenhos, que antes apenas possuíam contornos, agora possam dar uma impressão de profundidade. Da mesma forma, ao utilizarmos informação de relação de cores da reiluminação artística, é até possível atribuir cor a cada valor de iluminação.

A.2.4 Reconstrução de Mapas de Normais e Profundidade

Além destas aplicações, o *feature-correspondence mapping* pode ser utilizado para transferir outras informações do espaço 3D para imagem. Por exemplo, este mapeamento pode ser utilizado para transferir informações de normais e de profundidade de maneira robusta, gerando mapas suaves que capturam a essência da imagem, mesmo quando utilizamos o mesmo modelo 3D para diferentes imagens. Ao aplicarmos isso em esboços, é possível atribuir uma informação a imagens que previamente não a possuem, algo não realizável por técnicas atuais do estado da arte.

A.3 Resultados

Nós implementamos nossa técnica em C++ e Matlab com o auxílio de bibliotecas como Qt, OpenCV, OpenGL e armadillo. Utilizamos a iluminação de OpenGL com até 4 fontes de luz pontuais devido ao fato de proporcionar um feedback visual em tempo real,

mas nada impede que modelos mais complexos de iluminação possam ser utilizados. Nos resultados mostramos a robustez de nossa técnica, bem como sua eficácia em reiluminar objetos distintos sob diferentes condições de iluminação. Mostramos como nossa técnica pode reiluminar fotografias (Figure A.1), pinturas (Figure A.2) e desenhos (Figure A.3). Também mostramos como podemos adicionar mais fontes de luz, modificar as cores destas, bem como utilizar filtros de processamento de imagens diretamente na iluminação da imagem (Figure A.2(direita)), gerando efeitos peculiares. Testes com usuários novatos foram realizados, de forma a constatar que usuários que nunca utilizaram o sistema previamente foram capazes de gerar bons resultados de reiluminação em pouco tempo (menos de 10 minutos). Além disso, realizamos um experimento e constatamos que os resultados gerados pelo nosso método são convincentes e, em várias ocasiões, preferíveis às imagens originais.

Figura A.1: Reiluminação da fotografia da menina do Afeganistão, de Steve McCurry utilizando nosso Método Regular. (Esquerda) Fotografia original. (Direita) Fotografia reiluminada por baixo e pela direita.



Figura A.2: Reiluminação da pintura de van Gogh. (Esquerda) Pintura original. (Centro) Imagem reiluminada artisticamente frontalmente. (Direita) Imagem também iluminada frontalmente, mas aplicando um filtro de estilização sobre a iluminação. Note que, utilizando a iluminação artística, a intenção de cores utilizada pelo artista é preservada.



Figura A.3: Reiluminação artística de um desenho a lápis feito por Sunnyrays. (Esquerda) Desenho original. (Centro) Desenho reiluminado por baixo. (Direita) Desenho reiluminado utilizando duas fontes de luz.



A.4 Conclusão

Nós apresentamos uma solução prática ao problema de reiluminação de imagens com objetos de geometria arbitrária. Nossa técnica utiliza *shading proxies* e interação com o usuário para guiar o processo de reiluminação. Nosso método calcula uma *shading-ratio image* que é usada para mapear a iluminação original para a iluminação desejada. Nossa solução é flexível e robusta, podendo ser utilizada em objetos de diferentes geometrias com albedos não uniformes. Além disso, nosso método é o primeiro a realizar reiluminação artística, considerando o uso de diferentes cores para especificar regiões escuras e iluminadas. Adicionalmente, nós mostramos como nossa técnica pode ser utilizada para decomposição intrínseca de imagens, bem como para transferir iluminação para esboços e para recuperar mapas de normais e de profundidade.

Nós demonstramos a eficácia do nosso método reiluminando uma grande quantidade de fotografias, pinturas e desenhos; mostrando que a técnica pode ser usada em diversas situações para gerar efeitos não alcançáveis por técnicas atuais do estado da arte. Também mostramos como o sistema é fácil e prático de usar, mesmo para usuários novatos que não são especializados na área de reiluminação de imagens. Dado a flexibilidade, robustez e facilidade de uso, nossa técnica pode ajudar artistas, fotógrafos e usuários casuais a experimentar efeitos de iluminação em imagens existentes. Dessa forma, acreditamos que nosso método irá estimular novas e criativas aplicações de reiluminação de imagens.