



HAL
open science

Toward a Perceptually-relevant Theory of Appearance

Pascal Barla

► **To cite this version:**

Pascal Barla. Toward a Perceptually-relevant Theory of Appearance. Graphics [cs.GR]. Université Bordeaux, 2017. tel-01759272

HAL Id: tel-01759272

<https://inria.hal.science/tel-01759272>

Submitted on 10 Apr 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Toward a Perceptually-relevant Theory of Appearance

By

PASCAL BARLA



A dissertation submitted to Bordeaux University in accordance with the requirements of the HABILITATION A DIRIGER LES RECHERCHES in the Faculty of Computer Sciences.

OCTOBER 9TH, 2017

Jury members:

Pierre Poulin (Université de Montréal): rapporteur

Fabrice Neyret (Laboratoire Jean Kuntzmann): rapporteur

Tamy Boubekeur (Telecom ParisTech): rapporteur

Daniel Meneveaux (Université de Poitiers): examinateur

Mathias Paulin (Université Paul Sabatier): examinateur

Xavier Granier (Institut d'Optique): directeur

ACKNOWLEDGEMENTS / REMERCIEMENTS

I'd like to thank a lot of people for providing me with personal and scientific support, but since a good deal of them do not easily read English, I'll quickly switch to French for a little "aparté".

J'aimerais tout d'abord remercier Deborá pour avoir su subtilement me rappeler qu'il y a d'autres choses dans la vie au moins aussi importantes que la recherche, de sorte que je devrais parfois arrêter temporairement de me prendre la tête et m'occuper de la vaisselle par exemple. Je serai éternellement reconnaissant envers mon père et ma mère pour m'avoir fait croire — à tort — qu'il était tout à fait normal de suivre ma passion précoce pour l'informatique graphique; je n'ai réalisé que sur le tard que cela n'était pas chose courante, naïf et chanceux comme je suis. Je suis également reconnaissant de la clairvoyance de mes grand-parents qui, bien que ne parlant presque pas un mot d'Anglais, ont parfaitement compris le sens du mot "deadline", comme pouvaient en témoigner leurs sourires empreints de compassion. J'aimerais aussi saluer le talent d'actrices de ma tante et de ma cousine, qui ont vraiment sû donner une impression de profond intérêt lorsque je leur exposait ce qui pour moi devait constituer une théorie perceptuellement pertinente de l'apparence; vous m'avez presque berné ! Je vais bien être obligé de saluer ici mes amis d'enfance, qui ont pour leur part uniquement retenu le tout premier de mes résultats de recherche où un cube jaune était utilisé à des fins de tests; même si de leur point de vue, mon boulot consiste toujours à m'amuser avec des cubes jaunes, et à être payé pour ça !

Je vais maintenant passer à ma famille scientifique, en commençant par Joëlle Thollot, qui durant ma thèse a su contenir le Marseillais en moi, en bonne encadrante/psychologue qu'elle est. Elle a eu de nombreux "petits" depuis, et on forme une grande famille heureuse. Je me dois ici de remercier Pascal Guitton, pas uniquement car il est chauve comme moi (8-), mais aussi parce qu'il a suffisamment insisté pour que je tente d'obtenir un poste de chercheur en un temps lointain où il n'était pas nécessaire d'avoir un prix Nobel. De nombreuses excuses/remerciements vont sans surprise aux membres des équipes Iparla puis Manao, pour avoir patiemment et péniblement écouté au moins la moitié du flot de paroles que je peux produire sans discontinuer lors de réunions d'équipes, de brainstormings, de déjeuners, voire de pauses café ou même WC. Ma famille scientifique s'est agrandie lorsque j'ai pu participer au réseau PRISM dès 2013, et je suis extrêmement reconnaissant envers Roland Fleming de m'avoir enrôlé dans ce projet passionnant; je suis par contre moins reconnaissant envers Edward Snowden pour avoir révélé l'existence d'un autre projet du même nom en juin 2013, ce que j'ai trouvé un tantinet égoïste de sa part au vu des complications que cela a pu entraîner de notre côté.

Enfin, j'aimerais remercier naturellement les membres du Jury d'HDR, pour tout le temps que vous aurez investi dans la lecture de ce document, dans les transports en vue de la soutenance, et dans les conséquences qui vont s'en suivre sous la forme de longs emails de discussions interminables dont vous serez les innocentes victimes.

TABLE OF CONTENTS

	Page
1 Introduction	1
1.1 A closer look at depiction	2
1.2 The graphics pipeline(s)	4
1.3 Four stages of appearance	7
1.4 Stretching toward Optics and Vision	8
2 Image-space techniques	11
2.1 Image creation & manipulation	11
2.1.1 Diffusion curves	12
2.1.2 Image flows	16
2.2 Expressive shading & compositing	26
2.2.1 Surface flows	26
2.2.2 MatCap manipulations	31
2.3 Discussion	38
3 Object-space techniques	43
3.1 Radiance tweaking	43
3.1.1 Local warping and scaling	44
3.1.2 Multi-scale enhancement	49
3.2 Material design	56
3.2.1 Glossy appearance	58
3.2.2 Anisotropic reflections	65
3.3 Discussion	74
4 Conclusions and perspectives	79
4.1 Blueprint for a theory of appearance	81
4.1.1 Below the object surface	81
4.1.2 Between the object and the image	83
4.1.3 Above the image plane	86
4.2 Beyond appearance: Ecological Optics	89
Author's bibliography	91
Bibliography	93

INTRODUCTION

Computer Graphics techniques have historically been driven by a variety of active application domains, from the production of animations, movies, visual effects and video games, to scientific visualization or product design for instance. As a result, the field has evolved conjointly with new technologies (such as Graphics Processing Units, or GPUs), bringing innovations that often mix theoretical with practical considerations. One may argue that the first virtues of Computer Graphics are automation and interactivity. *Automation* tools have for instance lead to a progressive move from 2D hand-drawn to 3D animated characters in feature-length movies¹, as well as the near disappearance of animatronics in special effects, being replaced by virtual creatures. *Interactive* techniques have permitted video games to become the leading entertainment industry, from casual games on mobile devices, to independent or mainstream games, and more recently serious games and the generalized gamification of different types of activities. Different artistic practices have appeared along the way, corresponding to different types of tools used for specific production stages. For instance, while some artists are expert at modeling 3D shapes (e.g., with software such as ZBrush or Mudbox), others will focus on material or textures (e.g., using Substance or Quixel tools), whereas artists working in post production will deal with integration with live footage (e.g., using Nuke or AfterEffects). Different applications also place different constraints, and artists must incorporate them to produce effective images, such as in scientific illustrations for Medicine or Archeology where legibility is key.

Aside from being a versatile creation medium, Computer Graphics also constitute an unprecedented toolbox for probing how we (humans) perceive the world around us. Visual perception has been one of the main motivations for the research projects I have participated to, and it will thus be a recurring theme in the following pages. In particular, I will focus on *appearance*, which we may at this stage loosely define as “the result of the interaction between light, shape and matter”. Besides automation and interaction, one may ask: What can now be done in terms of appearance manipulation that was not possible before the advent of Computer Graphics? In principle, by offering control over each and every ingredient that make up an image, physically-impossible tweaking and wild appearance variations should become possible. However, all of this is not

¹Fortunately, frame-by-frame 2D animation remains very much alive in other application domains, from short movies, to music videos, advertisement and even documentaries.


really new, as testified by the work of traditional artists as discussed in Section 1.1, and we should at the very least take inspiration from them.

What is new is that Computer Graphics provide a *systematic* means to tinker with the image formation process, which not only opens the way to automation and interactivity, but also gives us the opportunity to better *understand* the visual ingredients that constitute object appearance. But where should we start from? After a quick presentation of the classical image production pipelines in Section 1.2, I lay down a plan into which I dive more progressively in Section 1.3, introducing the techniques that will be described in details in the following chapters. Since the long-term goal is to assemble these disparate elements into a coherent theory of appearance, I will focus on providing insights, not proofs or implementation details. Human Vision and Physical Optics will naturally play a fundamental role in this endeavor, but as discussed in Section 1.4, these domains still remain far apart from each other, so much that it will be necessary to weave links between them if one wants to understand appearance fully.

I have been more and more fascinated by both the entertaining/creative aspects of applied computer graphics techniques, and the eye-opening insights gained from dissecting the image formation process in a more fundamental manner. I have tried to share this dual passion with the colleagues and students (many of them now colleagues) I have had the chance to work with, and I hope to share it with readers of this dissertation as well.

1.1 A closer look at depiction

The manipulation of appearance in images is an old human activity, and even a quick look at the history of Art reveals these skills have been around for a very long time. For instance, paintings from 100 BC found among the ruins of Pompei exhibit effective “trompe-l’oeil” that cleverly manipulate composition and perspective. The artwork of Jerome Bosch, in the 16th century, shows a tremendous variety of material appearances, ranging from minerals, to metals, skins, different types of woods, of glass, etc. The rather recent advent of photography in the 19th century has added a new way to create images with its own codes, practices, and (relative) ease of use. However, photography should still be considered as image production, not as a simple record of reality (see Figure 1.1); indeed, the choice of viewpoint, time, parameters all participate in the quality of the final picture (photography literally means “drawing with light”!).




(a) “Les mains”
by Benoit Courti




(b) “The Tetons and the Snake River ”
by Ansell Adams

Figure 1.1: Skillfully-executed photography is more than a mere record of reality.

Painting, drawing and photography have the power to elicit an infinite range of impressions: they depict the shape of objects, plants or animals, their materials and textures, their relative depths in the scene, the time of day at which the image has been taken or is imagined to correspond, the environment in which the scene takes place such as open or close areas, the size of the different elements that may be observed, etc. Tactile or olfactory emulations may even be elicited by properly depicting familiar visual patterns that recall touch or smell by association. The artist often does not stop at producing or capturing a scene that can merely be recognized, but also works at conveying a mood, maybe tension or emotion between characters, atmospheres in landscapes, even functions or purposes when a mechanism is depicted or an action is taking place. This is already getting us a bit far from the topic of appearance, although not too far since cognitive aspects of a painting or photograph rely at least in part on simpler aspects of appearance. For instance, tension can be conveyed by wrinkles on body parts (Figure 1.1a), grandeur is enhanced by contrasts in luminosity and texture patterns (Figure 1.1b), and an atmosphere is set by a proper choice of tones, shadows and contrasts (Figure 1.2).



“Meules, effet de neige, le matin”
by Claude Monet



“Meules, fin de l'été”
by Claude Monet

Figure 1.2: Non-visual impressions may be conveyed by a proper choice of colors and shades.

Let me come back to the concept of appearance, which arises from the (imagined, observed or captured) interactions between lighting, and the shape and materials of objects present in the depicted scene. Since we only have an indirect, remote visual access to these objects, appearance must then manifest itself as a variety of *image patterns* that are diagnostic of specific shapes, materials or lighting configurations. Even if we somehow make use of these patterns to give sense to an image, we do so without even giving it a thought. It is thus difficult to pinpoint what effective combination of image patterns makes a plausible depiction of a scene, as any student of art quickly realizes: what is effortlessly perceived reveals to be challenging to reproduce.

Painters must learn how to create such patterns with various tools and media, and depending on the latter the appropriate technique will vary. As a simple example, consider surface highlights. They may be directly painted in white on top of existing dried paint layers when using oil or acrylic; whereas with watercolor, they are rather conveyed by leaving the white surface of the paper untouched. Students of photography are faced with the challenge of finding and capturing the right patterns at the right time and viewpoint. They may additionally make use of adequate lenses, focal lengths or even color filters when working in black and white. Studio photography goes one step further by providing control over the lighting environment.

I hope to show that despite the wide variety of artistic styles or photographic practices, the sought-for patterns are essentially the same. The work of many artists suggest that this is indeed the case by blurring the distinction between photography and painting. In photography,

more or less pronounced makeup will directly affect the apparent material (Figure 1.3a) and shape (Figure 1.3b), up to the point where the 3D volume may simply disappear. In painting, the artist may accurately reproduce the appearance of a person or object up to the point where it becomes nearly indistinguishable from a photograph (Figure 1.3c). It has even been suggested that some artists such as Johannes Vermeer made use of lenses as hinted by the occurrence of optical artifacts (e.g. bokeh, aberrations) in their paintings. Even if we acknowledge that the same varieties of image patterns are to be found in both paintings and photographs, this does not make them any easier to characterize. They may occur at a specific range of scales such as the texture indicating the rough peel of an orange, which is usually resolved only up to a few meters. The patterns may be constrained by each others, as with highlights that appear to bend toward the occluding contours of objects. They may depend on specific conditions, as is the case of the conspicuous glow of translucent objects that chiefly appears when the light comes from behind.




Figure 1.3: Artists that blur the boundaries between painting and photography: (a) using makeup makes the skin look metallic; (b) painting directly on people and decors turns a real scene into a flat 2D world; (c) painting with an extreme attention to details brings photorealistic impressions.

In spite of this seemingly daunting complexity, we know at least one system that effortlessly makes sense of these patterns in both paintings and photographs: the human visual system. In principle, it is thus possible to dissect any figurative image in elements that can be related to constituents of appearance (i.e., *perceived* lighting, shape, and material), even though they will not necessarily correspond to individual physical properties. There is actually much more to study: layout and composition, spatial relationships and balance, dynamics and motion; all of them crucial to the making of interesting, legible, engaging compositions. However, as promised, I will concentrate on the appearance of individual objects in static or moving environments, a topic dense and central enough to merit a first focus.

1.2 The graphics pipeline(s)

Computer Graphics tools are not fundamentally different from those of painting or photography: their purpose is to depict a captured or imaginary scene.

One way to achieve a photographic look is through physical simulation: complex light sources may be modeled by artists or captured from real-world luminaires, then light emitted by these sources might be transported throughout the scene, scattered by surfaces or through volumes, and finally gathered at a virtual sensor, where accurate modeling of lens systems will reproduce

familiar optical effects (i.e., motion blur, depth of field, etc). There are actually several methods to perform this simulation of light transport, which correspond to different strategies to simulate the *same* physical phenomena; hence, with enough time, any accurate-enough approach should converge to the same image, and thus the same appearance.



(a) 1958 Impala rendering
made with ©Maxwell




(b) Screenshot from Uncharted 4
by ©Naughty Dog

Figure 1.4: Photorealistic rendering (a) relies on light transport and detailed geometric and material models to produce confounding images. In games (b), performance constraints push designers to find the right tradeoff between visual plausibility and technical feasibility.


This suggests that with this approach, control over appearance is done before rendering, in the modeling of shapes, materials and light sources (see Figure 1.4a). In most cases though, a rendered image is only the first stage of a production, yielding a “raw material” that is most often tweaked in post-process, in particular for visual effects. This is usually referred as compositing, since it is common practice to mix synthetic imagery with live footage to give the illusion of a coherent appearance. Among the tweaks used by compositing artists, effects such as color grading, the imitation of optical effects (lens flares, motion blurs, etc), and even the addition of localized highlights are frequent. Video games take even more liberties compared to the physical ground truth, being constrained by real-time performance as they are. In this case, the balancing of efficiency requirements and of the final desired appearance becomes both a technical and an artistic achievement (see Figure 1.4b). In this respect, the recent notion of physically-based rendering is rather subjective: material models in particular are often chosen to be a good trade-off between physical plausibility and mathematical convenience.

A diametrically opposed image synthesis approach consists in imitating the style of artists who use traditional media, hence producing digital imitations of paintings or drawings. This not only includes the virtual deposition of paint or chalk for instance, but also the choice of how brush strokes and pencil lines should be applied. In this case, there is no reference solution to compare or converge to, since there might be as many styles as there are artists; the main challenge thus becomes one of achieving a desired stylized appearance, while still proposing automation and interactivity when possible.

A first solution consists in providing virtual tools that closely imitate pencils or brushes, hence placing all the burden of reproducing a compelling appearance in the hands of the digital artist. This naturally appeals to many traditional artists, since the computer is used minimally in the artistic process; but this is limited in the context of interactive content where drawing sequences must be replayed depending on the action (Figure 1.5a). Another solution consists in relying on some elements of a 3D scene, and introduce stylized rendering techniques that



(a) Screenshot from Jotun
by ©Thunder Lotus Games



(b) 3D medical illustration
by ©Beatmap

Figure 1.5: A stylized appearance may be achieved by (a) 2D drawing and animation using digital tools, or (b) expressive shading techniques applied to a 3D model.

produce appearances similar to those obtained with traditional media. This was initially called Non-Photorealistic Rendering to contrast with the simulation of light transport, with an emphasis on *how* to reproduce a stylized look while offering automation and interactivity whenever possible. Later work has started putting emphasis on *what* is represented by pencil lines, brush strokes or color gradients; they have been gathered under the term Expressive Rendering, which clearly states that the goal of the imaging process is to convey a specific message (see Figure 1.5b). Such considerations are more in line with the study of appearance.

In the techniques presented in the following chapters, the central objective is always to provide the artist (be it a product designer, a scientist, a film director, or an animator) with some form of control over appearance. However, the tools will differ depending on the type of image formation process: physically-based rendering, compositing, expressive rendering or 2D drawing.

This is clarified by the diagram of Figure 1.6, which I'll use throughout this document. On the one hand lies the *object surface*, where shape, material, and emission properties may be designed to control appearance in physically-based rendering. On the other hand lies the *image plane*, where colors, but also depth or opacity may be directly drawn or manipulated with 2D techniques to elicit some appearance percepts. In-between these two extremes, expressive rendering techniques permit to alter the image creation process to manipulate appearance in a potentially non-physical (yet plausible) manner.

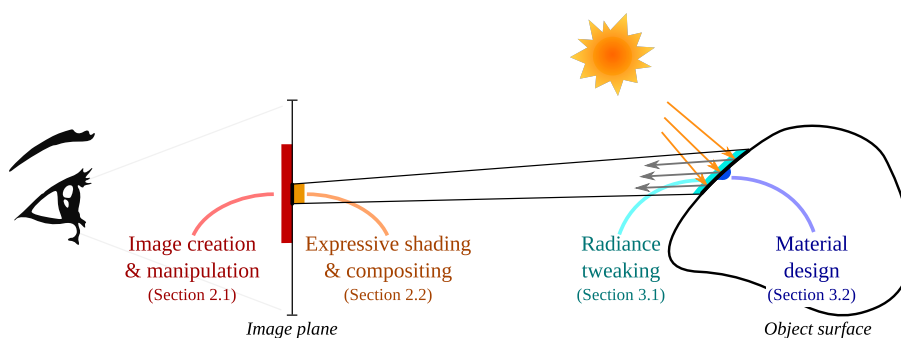



Figure 1.6: The computer graphics techniques presented in the upcoming chapters are distinguished depending on whether they work in the image plane, at the object surface, or in-between.


1.3 Four stages of appearance

Let's dive into the different stages shown in the diagram of Figure 1.6, starting from the image plane and progressing down to the object surface. This will serve as an overview of the two main chapters, and at the same time will give the opportunity to introduce the relevance of human visual perception to each stage of appearance.


At the image level (Chapter 2), the 3D scene is not directly available; still, appearance can be created or modified. In the minimal case, the digital artist may start with a blank canvas and fill it in, or use a photograph and modify it (Section 2.1). Two different approaches are usually made available to artists: vector or raster graphics tools (but digital painting software are tending to blur the distinction). We will first focus on image *gradients*, which makes sense to visual perception: the human visual system (HVS) is much more sensitive to contrasts and variations than absolute intensities. However, some of these variations will deserve a particular attention: *occluding contours*, as they are discontinuities of visibility, not mere strong gradients. In regions bounded by these discontinuities, analysis of gradient fields is informative of shape, due to the *compression and stretching* of textures or reflections.




At the shading and compositing stage (Section 2.2), we not only have access to colors but also normals, depths, material properties in so-called auxiliary buffers. They may either be output by a classic renderer, estimated from images, or painted. We will then show that shape cues conveyed by *compression and stretching* can directly be produced from available auxiliary data. Similarly, material cues may be manipulated, but care should be taken since a material may appear differently in different lighting conditions as shown in recent perceptual studies. Our solution will be to directly manipulate the *coupling of material and lighting*, bringing more direct artistic control for shading or compositing purposes.



The object level (Chapter 3) deals essentially with rays reflecting off and transmitting through object surfaces. Here, the interaction between light, shape and material is considered atomic: radiance coming from surrounding light sources is scattered according to the material at a surface point. The simulation of light transport does not leave much space for expressiveness, unless scattered radiance is tweaked in some non-physical way (Section 3.1). We will consider the immediate neighborhood of a surface point to compute *local curvatures*, and explore perceptually-relevant techniques to exaggerate their depiction, either modifying the light ray directions or intensities. From a perceptual standpoint, using view-dependent curvatures is the most logical choice, with scale automatically depending on the distance between the view and the surface. But in other instances such as scientific illustration, a manual control over the *pertinent scales* to focus on will be preferred for exploration purposes.



We will have by then reached the object surface, but not the end of the journey as we have not yet studied what materials are "made of" (Section 3.2). A material encapsulates light transport at micron (and smaller) scales, which is determined by the micro-structure of object surfaces visible (at best) in extreme close-up views. What is fascinating here is that slight variations in the configuration of micro-structures can have profound visual impacts on the final appearance of objects by producing distinctive image patterns. We will first study the appearance of gloss, which may be due to a variety of micro-structures of different scales and distributions. Among the dimensions involved in the perception of gloss, we will focus on *haze*, which has received only little attention so far. We will then consider more intricate reflections and highlights due to *elongated micro-structures* such as scratches or fibers, and show in particular how their orientations, profiles and distributions impact material appearance.




1.4 Stretching toward Optics and Vision

By going deeper down the scales we have already reached a point where Geometric Optics (i.e. the use of rays) as is most common in Computer Graphics, may not be a valid approximation: we shall then turn to Wave Optics (see Figure 1.7). Fortunately, in most cases of interest, visual effects due to the wave nature of light (e.g., iridescence, sheen, colored absorption) can be confined to the material function. This does not mean they do not have an impact on the final image; only that they can safely be decoupled from the macroscopic properties of light and shape at a human scale. The material thus acts as a *black box* that defines the minimum scale under which individual details essentially become invisible to the naked eye, no matter how close the surface.



(a) Microscopic pits on a Compact Disc



(b) Nanostructures on the wings of a Morpho

Figure 1.7: Microscopic surface irregularities with sizes on the order of a micron create complex interference patterns between reflected light waves in a way that depends on their wavelength. These yield structural colors that may be observed on artificial objects (a) or natural specimen (b).

What is actually visible to the naked eye depends, in the real-world, on the distance to objects; while in images it also depends on resolution. The minimum scale is thus not absolute but is rather a sliding scale that gets enlarged the more the viewpoint gets farther away from the surface of interest. Not only the material, but also the surface shape then gets modified — typically smoothed out. The configuration pictured in the diagram of Figure 1.6 remains valid, but is *instantiated* differently: part of the geometric complexity becomes incorporated in the material

function. Different points in the final image might actually correspond to different instantiations of surface/material, even for points on a same object if it is sufficiently elongated in depth. Some objects might entirely project below the pixel size, such as fur or hair, even branches and leaves seen from afar; these represent other challenges that we will not address here.

The light emitted from the image displayed on a screen or reflected off a page when it is printed yet has to reach the eye and be absorbed by photo-receptors. This is not a simple one-to-one mapping, since eyes make saccades every few hundred milliseconds, but also because even the very first stages of retinal processing involve complex spatio-temporal filters that vary from central (foveal) parts toward peripheral parts of the eye. Psychophysics is a mid-level approach to Human Vision that is more adapted to our purpose, as it is concerned with the identification of *which* image patterns, or cues, are used by our brain, leaving the question of *how* it is wired to Neurophysiology. The study of perceived shapes and materials is frequently called 3D Perception. In recent years, it has increasingly made use of Computer Graphics models to study how Vision works (Figure 1.8). The research presented in this dissertation not only takes inspiration from these studies, but also begins to be used in perceptual experiments to more accurately probe visual perception by providing control over intermediate representations of appearance.




Figure 1.8: A same image may have different physical interpretations; yet human vision somehow deals with this ambiguity even without a direct access to either shape or lighting. In this example reproduced from the work of Fleming et al. [1], the first interpretation is favored over the second: image distortions are more readily perceived to be due to shape than lighting.

A fascinating and exciting implication is that Computer Graphics techniques may be used as *stepping stones* to connect the physical and perceived worlds. On the one side, they take their roots from Optics, summarizing (and most often approximating) the essential phenomena that will survive to the image. On the other side, they grant direct manipulation over image patterns, hence fiddling with the visual cues that mediate our perception of the world.

In the next two chapters, I will elaborate on the connections between Optics, Graphics and Vision, relying on several research projects carried out over the course of ten years. These projects involved several PhD and PostDoc students, and lead to inter-disciplinary collaborations that have significantly expanded my scientific field of view.² I will gather the essence of these investigations in a last chapter, and sketch out the contours of what a perceptually-relevant theory of appearance should be.

²I will provide details on involved students, funding programs and collaborations in footnotes like this one.³

³Well, not exactly like this one, since I have not provided any type of details whatsoever.⁴

⁴I'll also try to write more interesting footnotes, and no more nesting, promised.


IMAGE-SPACE TECHNIQUES

The most direct approach for controlling appearance is to work directly in the image plane, that is to say without having to create a full 3D scene. One may start from scratch as with traditional painting and drawing, or manipulate an existing color image; both approaches are explored in Section 2.1. Additional 3D-like information may be made accessible in the form of so-called auxiliary buffers, which opens the way to novel appearance manipulation techniques as described in Section 2.2. In the end, methods working in the image plane essentially manipulate the same visual cues to appearance, as discussed in Section 2.3.

2.1 Image creation & manipulation

We assume in this section that no 3D information is available to control appearance. Hence all object properties (shape, material, textures) must either be extracted from an image or created from scratch by an artist.

We may distinguish two different image creation or manipulation approaches in the following: vectorial and bitmap in Sections 2.1.1 and 2.1.2 respectively. Vector graphics, in which primitives are defined geometrically, date back to the earliest days of computer graphics. Early cathode-ray tubes, starting with the Whirlwind-I in 1950, were all vector displays, while the seminal Sketchpad system [2] allowed users to manipulate geometric primitives like points, lines, and curves. Raster graphics provides an alternative representation for describing images via a grid of pixels rather than geometry; it arose with the advent of the framebuffer in the 1970s [3] and is now commonly used for storing, displaying, and editing images. While raster graphics offers some advantages over vector graphics — primarily, a more direct mapping between the hardware devices used for acquisition and display of images, and their internal representation — vector graphics offers a more compact representation, resolution-independence (allowing scaling of images while retaining sharp edges), and geometric editability.



2.1.1 Diffusion curves

Vector-based drawing tools offer only limited support for representing complex color gradients. While graphic artists, such as Takashi Morisaki (see Figure 2.1), have been able to achieve stunning results using existing vector-graphics tools, creating these artworks requires a high degree of artistic skill, not to mention an enormous degree of patience (the reported creation times for some art pieces run into the hundreds of hours).




Figure 2.1: Impressive results (“untenseki” ©Takashi Morikasi, shown at left) may be obtained with a large number of gradient meshes (right), at the price of several tens of hours of work.

What makes creating this artwork so difficult using vector primitives is that most existing vector formats support only linear or radial gradients. A more sophisticated vector-based tool for handling complex gradients is the *gradient mesh*. A gradient mesh is a lattice with colors at each vertex that are linearly interpolated across the mesh. While more powerful than simple gradients, gradient meshes still suffer from some limitations. A significant one is that the topological constraints imposed by the mesh lattice give rise to an over-complete representation (as seen in Figure 2.1-right), which becomes increasingly inefficient and difficult to create and manipulate.

Most color variations in an image can be assumed to be caused by edges (material and depth discontinuities, texture edges, shadow borders, etc.); even subtle shading effects can be modeled as though caused by one or more edges, and it has been demonstrated that edges constitute a near-complete and natural primitive for encoding and editing images [4, 5]. The main idea behind the diffusion curves image framework [OBW⁺08] is to adapt this idea to vector graphics.¹

Diffusion curve image The basic element of a diffusion curve is a geometric curve defined as a cubic Bézier spline (left of Figure 2.2) specified by a set of control points. The geometry is augmented with additional attributes: two sets of color control points (colored discs next to the curve) corresponding to color constraints on the *right* and *left* half spaces of the curve; and a set of *blur* control points (gray-level discs on the curve) that define the smoothness of the transition between the two halves. Both types of control points are interpolated along the curve to yield continuous color and blur constraints. Intuitively, the curves diffuse color on each side as in heat diffusion, with a soft transition across a curve given by its blur (right of Figure 2.2).

¹This project has been developed with Alexandrina Orzan and Adrien Bousseau during their PhD.




Figure 2.2: Rendering diffusion curves requires (1) the rasterization of the color and blur sources, along with the color gradients $\mathbf{w} = (\mathbf{w}_x, \mathbf{w}_y)$, (2) the diffusion of colors and blur, and (3) the reblurring of the color image I guided by the blur map B .

In practice, diffusion is computed on a pixel grid of fixed resolution in the original paper; hence color and blur constraints must first be rasterized. This corresponds to step 1 in the pipeline of Figure 2.2. For perfectly sharp curves, color sources are theoretically infinitely close. However, rasterizing pixels separated by too small a distance on a discrete pixel grid leads to overlapping pixels. Our solution is to distance the color sources from the curve slightly, and to add a color gradient constraint directly on the curve. The gradient maintains the sharp color transition, while the colors, placed at a small distance in the direction normal to the curve, remain separate.

Given the color sources and gradients, we next compute the color image I resulting from the steady state diffusion of the color sources subject to the gradient constraints. The blur image B is similarly obtained by diffusion of the blur sources, with no gradient constraint this time. This corresponds to a Poisson equation in the case of color, and a Laplace equation in the case of blur:

$$(2.1) \quad \begin{array}{ll} \Delta I = \operatorname{div} \mathbf{w} & \Delta B = 0 \\ \text{s.t. } I(x, y) = C(x, y) & \text{s.t. } B(x, y) = \sigma(x, y) \end{array}$$

where Δ and div are the Laplace and divergence operators, \mathbf{w} and C are the gradient and color constraints respectively, and σ is a blur constraint. The two diffusion processes correspond to step 2 in Figure 2.2. The final image is obtained after I is blurred according to B (step 3).

Application scenarii As a drawing tool, the main advantage of diffusion curves is to allow an artist to evolve an artwork gradually and naturally, whereas gradient meshes require careful planning and a good understanding of the final composition (see Figure 2.3). However, evaluating a diffusion curves image is much more complex compared to gradient mesh evaluation, and the pixel-based solver of the original paper brought other practical issues.²

An artist might also want to add stylization and expression to an existing image, which requires to first convert it to the diffusion curve representation. Conversion is based on a Gaussian scale space method we had developed in a previous work [OBBT07] and may be summarized as follows: edges are extracted using a Canny detector [6]; then we find the scale at which an edge is best represented (the more blurred the edge, the higher the scale) [4, 7] and use it to

²The demonstration software (running on Windows) provided at http://dept-info.labri.fr/~barla/download/diffusionCurves_pack.zip permits to concretely grasp the issues raised by a pixel-based solver.




Figure 2.3: Comparison with gradient meshes: (a) Original photograph; (b,c) Manually created gradient mesh (from a tutorial at <http://lifeinvector.com/>), with 340 vertices (and as many color control points); (d,e) Our drawing created by manually tracing over the image; there are 38 diffusion curves, with 365 geometric, 176 left-color, and 156 right-color control points.




Figure 2.4: Reconstruction example: (a) original image; (b) result after conversion into our representation; (c) automatically extracted diffusion curves; (d) RGB difference between original and reconstructed image (amplified by 4); note that the most visible error occurs along edges, most probably because, through vectorization, we change their localization.

identify the degree of blur; we finally connect pixel-chains from the edge map and proceed to sample colors in the original image on each side of the edge in the direction of the edge normal. A reconstruction result is shown in Figure 2.4: even though the original image is not reproduced exactly (as shown by the difference image at right), the diffusion curve representation provides a good approximation that may then be edited curve by curve.

Follow-up work The original diffusion curves paper has fostered a number of follow-up projects in the Computer Graphics community. The main idea has been extended to different applications, such as terrain generation [8], volumetric modeling [9], texture design [10] or the modeling of surface details [11]. Other authors have proposed more accurate methods for converting a raster image into the diffusion curve representation, either focusing on colors and textures [12], organizing the curves in a hierarchical structure [13] (hence easing editing), or using shape optimization techniques to obtain fewer and more representative curves [14].

The set of diffusion curve types has also been extended to provide more control, notably on color gradients away from curves, such as with diffusion constraints [15]. A major improvement in this direction has been the extension from harmonic to bi-harmonic diffusion curves, replacing the Laplace equation with the bi-Laplace equation [16]. This introduces several new types of control curves, in particular providing artists with control over color gradients *on* curves. Most importantly, it permits to produce smoother diffusions that better mimic shading effects (see Figure 2.5a). Recently, the generalized diffusion curves approach [17] has suggested an




Figure 2.5: Bi-harmonic diffusion curves permit to obtain a plausible smooth-shaded appearance (a) while retaining a few number of curves (b). The generalized diffusion curves provide an alternative with simpler controls that produces similar-looking results (c), although not as controllable (some differences appear on the top edge of the sunglasses, and on the rounded wing).

alternative technique using a pair of Laplace diffusions blended with an appropriate weighting function. Each curve stores two sets of colors on each side: one color is assigned close to the curve, the other is assigned away from it. The resulting color gradients exhibit a smoothness similar to bi-harmonic diffusion curves (see Figure 2.5c), with simpler albeit less accurate artistic controls.

The pixel-based solver of our original paper has been first improved by a variable stencil method [11, 18], with better stability and performance; then by methods based on 2D ray tracing [19, 20]. Subsequent work has focused on obtaining fast and resolution-independent solvers. An approximate method based on triangulation [21] has been proposed first, then solutions based on boundary elements [22, 23], or fast multipoles [24].




Figure 2.6: Input primitives (a) are first discretized, triangulated and input to our FEM solver that outputs a set of elements (b), which are efficiently rendered to yield the final image (c).

Vectorial solver Our second contribution has been the development of a new solver based on Finite Element Methods (FEM) [BBG12].³ It is, to the best of our knowledge, the only solver that works for both Laplace and bi-Laplace equations while producing vectorial outputs. Its principle is outlined in Figure 2.6: the control curves are first discretized, used as constraints in a Delaunay triangulation and input to the FEM solver, which outputs one element per triangle; each element then acts as a color gradient that is rendered on the fly at runtime using a close-form evaluation. One clear benefit of this approach is that it decouples interaction from evaluation: only rendering is required when panning or zooming as in Figure 2.14, yielding real-time performances (several

³This project has been developed with Simon Boyé during his PhD.



Figure 2.7: Our solver may be used to produce vectorial normal maps (a), and the triangular representation may be accessed randomly and used as a vectorial environment lighting (b). Our formulation grants more artistic controls such as point primitives and transmission curves that let colors go through the curve in an asymmetric fashion. For instance, placing in (c) a diffusion curve (solid) inside a closed transmission curve (dashed) permits to mimic highlights (d).

hundreds of frame per second depending on the model).⁴ When a control curve is modified, the full process must be re-applied, which still provides interactive feedback in most cases. Another advantage of working with a triangulation is that it is easily adapted to work on 3D surfaces by relying on the available triangles of a mesh representation, which grants random access provided an adequate acceleration structure is put in place. As detailed in the paper, we have used the solver for color diffusion, as well as height field and normal map generation (Figure 2.7(a,b)); we have also introduced a few more types of control primitives, notably transmission curves (Figure 2.7(c,d)) that permit to easily mimic highlights. In the short-term, it would be interesting to use our solver for producing generalized diffusion curves [17].⁵

2.1.2 Image flows

We now turn to the manipulation of appearance in bitmap images. The palette of image processing techniques available in software such as Photoshop or Gimp is considerably large. We will essentially focus on methods that rely on a notion of image structure, which is often a necessary first step before being able to manipulate appearance. In Figure 2.8, we show two basic examples: the structure of the image gradient is extracted in an image to guide an oil paint filter (left); or a new structure is obtained manually using a warping tool (right). Both methods might require substantial manual refinement and corrections: the former to separate objects that have been filtered together (e.g., branches and leaves), the latter to correct an appearance that departs too much from the original (e.g., wiggly silhouettes).

I argue that with a better image structure, not only this manual task might be greatly simplified, but new "smart" image manipulation tools coming closer to *direct* appearance manipulation could be devised. The works of Ben-Shahar et al. (e.g., [25, 26]) and Fleming et al. [1, 27] have been an inspiration for this line of study: they showed in the context of Computer and Human Vision that the analysis of image gradients lead to a characterization of texture and shading patterns. We have pursued a similar direction of research with Graphics applications in mind.

Why gradient fields? When stylizing images and videos, the use of a gradient field permits to orient filters along image features for a variety of applications going from line drawing [28], to color abstraction [29] or painterly rendering [30–32]. Directly dealing with gradient vectors

⁴A WebGL implementation is available at <http://simonboye.github.io/wvg/>.

⁵Our solver is available in the Vitelotte module of the Patate library: <http://patate.gforge.inria.fr>




Figure 2.8: Two image processing tools available in ©Adobe Photoshop . Left: the oil paint tool gives a photograph a painterly look. It is semi-automatic and is guided by the input image gradients. Right: the liquify tool lets users warp image pixels manually, with optionally frozen areas (in pink). In the bottom row, some parts of the image are unwarped and refined manually.

requires them to be consistently oriented. Unfortunately, these orientations are most of the time too incoherent for smoothing purpose: for instance, a thin line feature is expected to yield a unique tangent direction while its sides define opposite gradients. Attempting to locally re-orient them as Kang et al. [29] is bound to fail if the reference direction is not reliable and/or when considering large neighborhoods.

In order to overcome these difficulties, the most common approach consists in encoding the *non-oriented* vector information into a *structure tensor* [33]. It is defined as the integral of the outer product of the image gradient ∇I with itself on a support of given size: $\int_{\mathcal{X}} \nabla I(\mathbf{x}) \nabla I(\mathbf{x})^T d\mathbf{x}$. This representation has been employed to interpolate very sparse non-oriented vector data for synthesis [34] or abstraction [35] purposes. Filtering structure tensors regularizes noise and/or merges gradients coming from multiple image channels [36]. Unfortunately, a structure tensor can only encode a locally *constant* gradient field. This constitutes a major limitation when used to average information over curved neighborhoods. In particular, it quickly fails to preserve the original structure and tends to smooth out important surface or image details (see Figure 2.9b-left). This problem is amplified with increasing neighborhood sizes and/or sparsity.

We have introduced a new local approximation method [CGBG13] that better preserves image features (Figure 2.9b-right) even with large neighborhood sizes or sparse inputs (Figure 2.9(c,d)). The key idea is to extend the structure tensor to a *higher-order* basis, more specifically an isotropic linear basis that provides the best trade-off between accuracy and smoothness.⁶

Non-oriented MLS gradient fields Our algorithm takes as input a set of non-oriented unit vectors $\mathbf{g}_i \in \mathbb{R}^2$ specified at sample positions $\mathbf{x}_i \in \mathbb{R}^2$. The \mathbf{x}_i and \mathbf{g}_i usually correspond to pixel coordinates and *normalized* image gradients; however, as seen in Figure 2.9(c,d), the method might also be applied to vector input, in which case the positions are not restricted to lie on a pixel grid, and the gradients are the vectors orthogonal to local tangents. As illustrated in Figure 2.10a, given an arbitrary evaluation point $\mathbf{x} \in \mathbb{R}^2$, we assume the neighboring unit gradients \mathbf{g}_i may be well approximated by a low degree polynomial gradient field defined in matrix form as

⁶This project has been developed when Jiazhou Chen spent one year of his PhD in our research team.

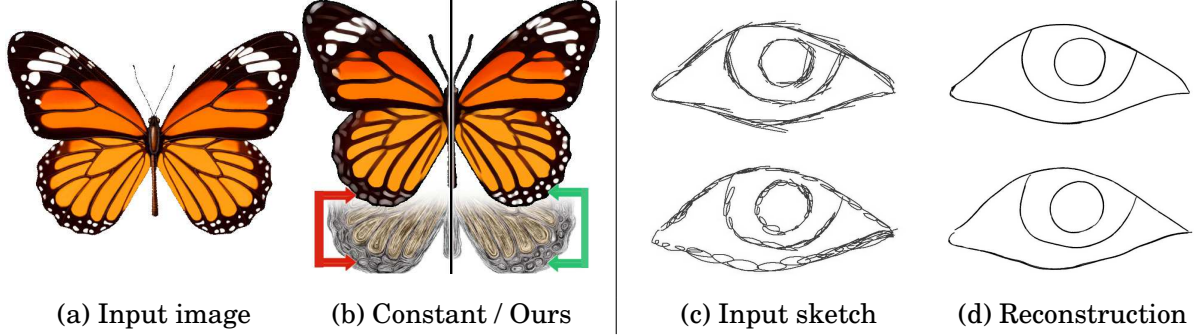


Figure 2.9: Left: our isotropic linear approximation (b-right) outperforms on the constant approximation (b-left) classically used in flow-guided smoothing [36]; in particular, round structures are better preserved. Right: our approach for reconstruction of a 2D curve (d) from an input sketch (c) is quite robust to the style: straight lines on the top, scribbling on the bottom.

$\mathbf{g}_{\mathbf{x}}(\cdot) = \mathbf{B}(\cdot)^T \mathbf{u}(\mathbf{x})$, where \mathbf{B} is a polynomial basis matrix with 2 columns, and \mathbf{u} is the vector of unknown coefficients.

To find \mathbf{u} , we use a strategy based on Moving Least Squares (MLS): the coefficients \mathbf{u} are obtained by maximizing the alignment between the *local* gradient field $\mathbf{g}_{\mathbf{x}}$ and the prescribed input vectors \mathbf{g}_i independently of their orientation, using a dot product:

$$(2.2) \quad \mathbf{u}(\mathbf{x}) = \arg \max_{\mathbf{u}} \sum_i W_i(\mathbf{x}) \left(\mathbf{g}_i \cdot \left(\mathbf{B}(\mathbf{x}_i - \mathbf{x})^T \mathbf{u} \right) \right)^2,$$

where $W_i(\mathbf{x}) = w\left(\frac{\|\mathbf{x}_i - \mathbf{x}\|}{s_i}\right) \|\nabla I(\mathbf{x}_i)\|^2$ weights samples according to a smooth decreasing function w over a support of size s_i centered around \mathbf{x} , with $\|\nabla I\|$ accounting for differences in gradient magnitude ($\|\nabla I\| = 1$ for line drawings). We use $w(t) = (1 - t^2)^2$ if $t < 1$, $w(t) = 0$ otherwise [37]. In order to avoid the trivial solution of a vector field with infinite magnitude, we impose that the norm of $\mathbf{g}_{\mathbf{x}}$ near the considered samples should be equal to 1 on average, using a quadratic normalization constraint. We show in the paper that this amounts to a *generalized* eigenvalue problem. If we use a constant basis, then the problem simplifies to a classic eigen value problem, where \mathbf{u} is the maximum eigenvector of the structure tensor $\mathbf{T}(\mathbf{x}) = \sum_i w\left(\frac{\|\mathbf{x}_i - \mathbf{x}\|}{s_i}\right) \nabla I(\mathbf{x}_i) \nabla I(\mathbf{x}_i)^T$.

From these local approximations, we reconstruct a continuous non-oriented vector field $\mathbf{G} : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ that globally approximates the discrete vectors \mathbf{g}_i (Figure 2.10b). Since we are using a centered basis, \mathbf{G} is defined as $\mathbf{G}(\mathbf{x}) = \mathbf{g}_{\mathbf{x}}(0)$. This MLS reconstruction of the gradient field can be used directly for filtering an input image I . When the input samples are supposed to lie on an unknown curve, each local gradient approximation $\mathbf{g}_{\mathbf{x}}$ can be integrated to recover a local scalar potential $f_{\mathbf{x}} = c + h_{\mathbf{x}}$. The constant offset c is recovered such that the 0-isosurface of $f_{\mathbf{x}}$ best approximates the input sample positions \mathbf{x}_i nearby \mathbf{x} , in a moving least squares sense. The sought-for manifold is then implicitly defined as the 0-isosurface of a continuous unsigned scalar potential $F : \mathbb{R}^2 \rightarrow \mathbb{R}$, which is computed by $F(\mathbf{x}) = f_{\mathbf{x}}(0)$ (Figure 2.10c).

Choice of basis The choice of basis for $\mathbf{g}_{\mathbf{x}}$ is critical since it characterizes the quality of the approximation. We consider constant \mathbf{g}^0 , linear isotropic \mathbf{g}^* , and general linear \mathbf{g}^1 gradient fields

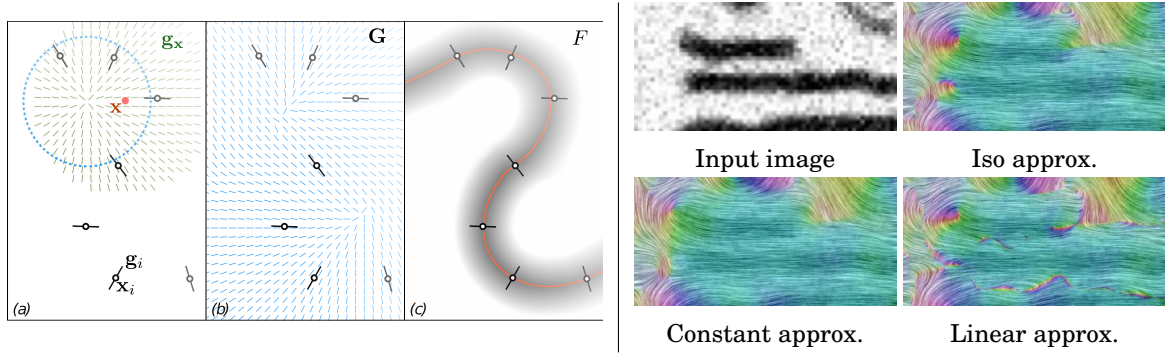


Figure 2.10: Left: overview of our MLS approach. (a) In green, the local approximation \mathbf{g}_x of the input vectors \mathbf{g}_i around the evaluation point \mathbf{x} . (b) The global reconstruction of the continuous non-oriented vector field \mathbf{G} . (c) The reconstructed continuous and unsigned scalar potential F , and its 0-isovalue. Right: comparison of choices of basis on a zoomed test image. The gradient fields are visualized using line-integral convolution [38], and a color coding of the angle. The constant approximation tends to oversmooth (e.g., at line extremities). The full linear approximation leads to over-fitted issues. Our isotropic linear solution provides the best trade-off.

whose expression and respective scalar potential are given below:

$$(2.3) \quad \begin{aligned} \mathbf{g}^0(\mathbf{y}) &= \mathbf{r} & \rightarrow & f^0(\mathbf{y}) = c + \mathbf{r}^T \mathbf{y} \\ \mathbf{g}^*(\mathbf{y}) &= \mathbf{r} + l\mathbf{y} & \rightarrow & f^*(\mathbf{y}) = c + \mathbf{r}^T \mathbf{y} + \frac{1}{2} l \mathbf{y}^T \mathbf{y} \\ \mathbf{g}^1(\mathbf{y}) &= \mathbf{r} + \mathbf{L} \cdot \mathbf{y} & \rightarrow & f^1(\mathbf{y}) = c + \mathbf{r}^T \mathbf{y} + \frac{1}{2} \mathbf{y}^T \mathbf{L} \mathbf{y} \end{aligned}$$

where \mathbf{L} is a symmetric matrix to ensure that \mathbf{g}^1 is a gradient field. Once integrated, they represent a line, a circle, and a general quadric respectively. The different choices of basis are visualized on a test image at the right of Figure 2.10. The constant basis cannot approximate well highly-curved regions, yielding over-smoothed features. On the other hand, the general linear basis already presents too many degrees of freedom, and leads to over-fitting issues in the form of oscillations and details which are not present in the input data. The isotropic basis clearly appears to be the best trade-off: compared to a constant basis it only adds one degree of freedom, which permits to properly capture round or curved image features (see also Figure 2.11b).




Figure 2.11: Applying flow-guided smoothing using our isotropic approximation of the gradient field (b) preserves round features well, but is agnostic to occluding contours and feature scales. Adjusting and clipping of the neighborhood using a hand-drawn scale image (c-top) and hand-selected segmentation (c-bottom) is equivalent to using a non-uniform weighting, which locally improves the preservation of small features (e.g., on the peacock body and head) and the preservation of original discontinuities (e.g., between the body and the tail feathers).

Advantages and limitations Our approach brings a number of other interesting functionalities.⁷ First, since it is based on MLS, it always provides a continuous solution at arbitrary levels of zooms since the evaluation location \mathbf{x} is not constrained to lie on a grid of pixels (see Figure 2.10-right for instance). Second, the isotropic basis provides with a simple way to compute approximate tangent field curvatures $\kappa_t \approx l/\|\mathbf{G}\|$, which may be put to good use for filtering purposes. Third, the support size of the weighting function may be varied locally, or even clipped to discontinuities according to user-drawn masks. This is illustrated in Figure 2.11: our method (here used to guide a coherence-enhancing filter [35]) permits to properly identify the round feather shapes of the peacock, but it fails to capture the proper scales and discontinuities (e.g., between body and feathers). With the addition of a pair of scale and segmentation masks, our method is easily adapted to better convey the structure of the image after filtering. A challenging direction of future work is to identify object boundaries and feature scales automatically.

The ultimate goal of appearance manipulation in images is to modify the perceived attributes of depicted objects, such as their shape or material. One approach consists in separating an image into reflectance and shading components, as in the seminal intrinsic decomposition technique of Barrow et al. [39]. Modern solutions to the intrinsic decomposition problem [40–42] target complex components (transmission, indirect vs direct lighting, etc) but often require user annotations. An alternative approach to manipulate image content is to reconstruct 3D information, then modify it and finally re-render the image. This is either done by making simplifying assumptions on shape and/or lighting [43–45], or by relying on databases of 3D objects [46].

In contrast to previous work, we have proposed a method [VBBF16] that does not require user intervention and manipulates *directly* apparent shape without requiring any 3D reconstruction.⁸ It relies on recent work in visual perception [1, 27], which suggests that image properties may be indicative of object shape properties, simply because they are sufficiently *correlated* together. In particular, orientation patterns due to compression and stretching of the reflected environment are mainly due to surface curvature, as is illustrated in Figure 2.12-left.

Properties of the structure tensor In order to better understand these correlations, we consider a simplified model of image formation in the neighborhood of a pixel of interest \mathbf{x} . We assume the image intensity is given by $I(\mathbf{x}) \approx S(\mathbf{m}(\mathbf{x}))$, where $\mathbf{m} = [m_0 \ m_1]^T = [\frac{-n_0}{n_2} \ \frac{-n_1}{n_2}]^T$ is the slope corresponding to the surface normal $\mathbf{n} = [n_0 \ n_1 \ n_2]^T$ expressed in screen-space; and S is a *local* shading function that accounts for both reflected radiance and tone mapping. We may now compute the structure tensor as:⁹

$$(2.4) \quad \mathbf{T}_I(\mathbf{x}) = \sum_i w_i(\mathbf{x}) \nabla I(\mathbf{x}_i) \nabla I(\mathbf{x}_i)^T = \sum_i w_i(\mathbf{x}) \nabla \mathbf{m}(\mathbf{x}) \nabla S(\mathbf{m}) \nabla S(\mathbf{m})^T \nabla \mathbf{m}(\mathbf{x})^T,$$

where $w_i = w\left(\frac{\|\mathbf{x}-\mathbf{x}_i\|}{s_i}\right)$ as before and we have used $\nabla I(\mathbf{x}) = \nabla \mathbf{m}(\mathbf{x}) \nabla S(\mathbf{m})$. The surface slope gradient $\nabla \mathbf{m} = [\nabla m_0 \ \nabla m_1]$ is by definition equivalent to the Hessian of the surface depth. It is related to surface curvature, but also depends on the positioning of the surface with respect to the viewpoint. The shading gradient ∇S incorporates angular variations due to both lighting and material; hence its magnitude will be higher for specular materials and sharp lighting environments.

⁷The method is available in the Grenaille module of the Patate library: <http://patate.gforge.inria.fr>

⁸This work has been conducted in the context of the PRISM Marie Curie ITN project.

⁹We use the constant basis of our MLS-based technique instead of the isotropic one to simplify derivations.




Figure 2.12: Left: perception studies suggest that orientation patterns (here due to the compression and stretching of reflections on a shiny 3D object) provide strong visual cues to surface shape. Right: structure tensors (shown at bottom with LIC and the same color code as before) capture these compression and stretching patterns. Indeed, tensors computed from geometric data (normal maps) or renderings of an object made of different materials are globally similar.

In order to relate \mathbf{T}_I to surface shape, we compute the structure tensor \mathbf{T}_m on surface slopes:

$$(2.5) \quad \mathbf{T}_m(\mathbf{x}) = \sum_i w_i(\mathbf{x}) \nabla \mathbf{m}(\mathbf{x}) \nabla \mathbf{m}(\mathbf{x})^T,$$

The difference between Equations 2.4 and 2.5 lies in the term $\nabla S \nabla S^T = \|\nabla S\|^2 \mathbf{R}_{\phi_S} \mathbf{P}_0 \mathbf{R}_{-\phi_S}$, where \mathbf{R} is a 2D rotation matrix, ϕ_S is the angle between ∇S and $\mathbf{e}_0 = [1, 0]^T$, and $\mathbf{P}_0 = \mathbf{I} - \mathbf{e}_1 \mathbf{e}_1^T$ is the projection onto \mathbf{e}_0 . This additional term has several implications on apparent shape:

- In regions where S is nearly uniform (e.g., a dim diffuse reflection), we get $\|\nabla S\| \approx 0$ and $\mathbf{T}_I \approx \mathbf{0}$. Shape features are not apparent, which is called masking in human vision [47].
- When ∇S is strongly directional (e.g., a collimated lighting), all surface slope gradients are projected in roughly the same direction. For instance, a horizontal ∇S implies $\mathbf{T}_I \propto \mathbf{T}_{m_0}$: only horizontal shape features are apparent, which is related to directional masking [1].
- Shape will tend to be exaggerated when $\|\nabla S\|$ is high and ∇S aligned with $\nabla \mathbf{m}$ (e.g., shiny elongated reflections). This is related to distortion effects in shape perception [48, 49].

In spite of the discrepancies between apparent and actual surface shape, \mathbf{T}_I remains strikingly similar to \mathbf{T}_m as shown in Figure 2.12-right. As confirmed in the original paper through a correlation analysis, the structure tensor indeed contains reliable visual cues to surface shape across a wide range of material and lighting.

Flow-guided warping Our image deformation algorithm relies on the aforementioned property of the structure tensor to exaggerate the orientation patterns found in an input image. To this end we introduce a flow-guided warping technique (see Figure 2.13) that exaggerates compression and stretching patterns so as to give the illusion of shape change. More specifically, we use the maximum eigenvalue λ_I^{\max} of \mathbf{T}_I to guide warping since it provides the magnitude of compression or stretching along the dominant local orientation. The method is detailed in Algorithm 1, which takes as input an image I and is controlled by two parameters: a scale bias γ to select fine or large features, and a warping parameter α that controls the signed magnitude of warping. We optionally take a gray-scale mask to restrict warping to a single object.




Figure 2.13: Our warping algorithm starts with an image I in (a) and computes structure tensors \mathbf{T}_I shown in (b) before and after continuation. Their maximum eigen values λ_I^{\max} are shown in (c) before and after local adaptation. The displaced coordinates ($\alpha = -1$, guided by $\nabla\lambda_I^{\max}$) are visualized using a distorted grid in (d). They are used to produce the warped image $I \circ D$ in (e).

Algorithm 1 Flow-guided warping

```

1: procedure WARP( $I, \alpha, \gamma$ )
2:    $s \leftarrow \text{scaleSelection}(I, \gamma)$  ▷ Bias warping toward small or large features
3:    $\mathbf{T}_I \leftarrow \text{structureTensor}(I, s)$ 
4:    $\mathbf{T}_I \leftarrow \text{enforceContinuation}(\mathbf{T}_I, s)$  ▷ Reduce variations due to shading
5:    $\lambda_I^{\max} \leftarrow \text{maxEigenValue}(\mathbf{T}_I)$ 
6:    $\lambda_I^{\max} \leftarrow \text{localAdaptation}(\lambda_I^{\max}, s)$  ▷ Give equal importance to all features
7:    $\nabla\lambda_I^{\max} \leftarrow \text{guidingFlow}(\lambda_I^{\max}, s)$ 
8:   for each pixel  $\mathbf{p}$  do
9:      $D(\mathbf{p}) \leftarrow \text{displace}(\mathbf{p}, \nabla\lambda_I^{\max}, s, \alpha)$  ▷ Move a distance  $\alpha s$  along  $\nabla\lambda_I^{\max}$ 
10:   $D \leftarrow D * G_s$  ▷ Smooth-out potential warping discontinuities
11:  return  $I \circ W$ 

```

A detailed description of the algorithm is provided in the original paper; I only present the main steps here. The warping method starts by computing a scale image s controlled by the bias parameter γ , which is used to compute a tensor image \mathbf{T}_I . The latter is smoothed in the direction orthogonal to the maximum eigen direction using flow-guided smoothing [50] to reduce variations due to shading. Next, the maximum eigenvalues λ_I^{\max} are computed, then centered and rescaled in a way similar to local lightness adaptation. Pixel coordinates are displaced along the gradient direction $\nabla\lambda_I^{\max}$ until they have moved a (signed) distance αs . The resulting coordinates are then smoothed out using a Gaussian G_s at the local scale s to get rid of potential clumping of pixels in the same locations; they are then used to produce the final warped image.

Results As shown in Figure 2.15, the method works for a variety of materials and lighting environments. The figure also illustrates the effect of the scale bias γ , and warping magnitude α ; the latter makes shape appear either sharper or rounder using negative or positive values respectively. The method may be used as a brush tool similar to that of Figure 2.8-right, with much less manual labor since it relies on the apparent shape of the depicted object. The size of the brush then directly controls the scale bias γ , while its intensity controls the absolute warping magnitude $|\alpha|$, its sign being chosen using a control key.

We have validated our warping technique with perceptual experiments that show that subjects could readily perceive the change in shape. More precisely, we asked 12 subjects (excluding

authors) to compare original and warped images for different values of α and assess “In which of the two images the object shape appears sharper”. Participants gave correct results in 86.6% of the cases (with a std of 26.5%) when the original images were synthetic, and in 85.8% of the cases (with a std of 28%) when they were photographs. This confirms that our flow-guided warping algorithm effectively alters apparent shape in a controllable way on a variety of input images.

Remaining limitations The method has a number of limitations that remain to be addressed. Images containing many occluding contours (e.g., natural landscapes) or optical effects (e.g., depth of field) clearly violate our hypothesized image formation process. These problems are circumvented in the current approach by providing an alpha mask identifying an object in focus occupying a large region of the image, hence relying on the user to provide an element of image structure.

High-frequency variations due to textures or shadows are not part of our hypothesis either. Nevertheless, our approach provides reasonably plausible results in these cases, suggesting that these image patterns could be analyzed in terms of distortions as well. However, multiple *layered* orientation patterns of different types might occur as well. Examples include shading patterns overlaid on textures, the occurrence of shadows, or background seen through transparent objects. One potential line of research would be to study the differences in structures of each such pattern. This might make it possible to better decompose an image without requiring user annotations.

We have focused on local apparent shape as we only consider the immediate neighborhood of an image pixel. A potential improvement would be to analyze and manipulate object shape with a more global approach. This is necessary not only to disambiguate concave shape features from convex ones, but also to preserve object symmetries or regularities through manipulation.

Finally, our method is not temporally coherent: when applied to a video, the shape seems to wobble. The manipulation of shape in videos will likely require to consider a kinematic description of apparent shape in terms of spatio-temporal image properties.



Figure 2.14: Our vectorial solver for diffusion curves permits to create complex vector gradients. It is based on a triangular representation that is output-insensitive; hence it does not need to be updated when global transformations (e.g., panning, zooming in, instancing) are applied.




Figure 2.15: Flow-guided warping applied to photographs exhibiting different materials and lighting conditions. Each result shows the initial image in the middle, a sharpened result ($\alpha = -1$) on the left and a rounded result ($\alpha = 1$) on the right. The scale γ is indicated at left. With increasing values of γ , shape features of larger sizes are altered.

2.2 Expressive shading & compositing

Images may not only store colors, but also additional information such as geometry (normals, depth, etc) or shading terms (diffuse, specular, etc). The simplest way to obtain these so-called *auxiliary buffers* is to render them, which is commonly done in expressive shading for games and interactive modeling tools, or in compositing for visual effects and design applications (Figure 2.16).

However, it is also possible to create them from scratch using painting tools (as we have shown in Section 2.1.1, see Figure 2.7b), or to estimate them from images as is common in texturing software such as Substance or Quixel. The main challenge in using auxiliary buffers for appearance creation and manipulation is that only visible surfaces are available, and no direct information about lighting is provided.¹⁰




Figure 2.16: Top: a matcap is an image of a sphere depicting a material that has been either directly painted by an artist or captured from a photograph. The variety of material appearances achievable with matcaps may be transferred to a 3D object by assigning a color to each surface point based on its *local normal*: the color is fetched at the matcap location corresponding to the same normal on the sphere. Bottom: most rendering engines output multiple *render buffers*, which may be edited in post-process before being added together to produce the final composite.

2.2.1 Surface flows

In section 2.1.2 we have described how patterns of compression and stretching due to shading and reflections are diagnostic of surface properties, and may be manipulated directly in the image to give the illusion of shape changes. We now work the other way around, assuming that surface shape information is available in the form of auxiliary buffers that store depths and normals, with the goal of *creating* an image with a desired appearance. This is again inspired by the work of Fleming et al. on shape from reflections [1], but also on their work on shape from smeared textures [27] where they suggest that *different* patterns of texture compression and stretching also convey shape properties. We investigate these patterns from a different vantage point, by analyzing how they may emerge as a product of the image formation process.

¹⁰All the projects presented in this section have been developed in Gratin [VB15], a node-based GPU-friendly software adapted to this intermediate level of manipulation. It is free and open source, so there is no reason for not trying it out (visit <http://gratin.gforge.inria.fr/>).

In computer graphics, images are most often synthesized using the rendering equation [51]:

$$(2.6) \quad I(\mathbf{x}) = L_e(\mathbf{p}, \boldsymbol{\omega}_o) + L_r(\mathbf{p}, \boldsymbol{\omega}_o), \quad \text{with} \quad L_r(\mathbf{p}, \boldsymbol{\omega}_o) = \int_{\Omega_{\mathbf{n}}} \rho_{\mathbf{n}}(\mathbf{p}, \boldsymbol{\omega}_o, \boldsymbol{\omega}_i) L_i(\mathbf{p}, \boldsymbol{\omega}_i) (\mathbf{n} \cdot \boldsymbol{\omega}_i) d\boldsymbol{\omega}_i,$$

which states that the image intensity I is the sum of emitted L_e and reflected L_r radiances leaving a surface point \mathbf{p} in the direction $\boldsymbol{\omega}_o$ toward the center of projection and intersecting the image plane at \mathbf{x} .¹¹ The reflected radiance L_r is obtained by integrating the incoming radiance L_i at \mathbf{p} on a hemisphere $\Omega_{\mathbf{n}}$ around the surface normal, with each light direction $\boldsymbol{\omega}_i$ having a potentially different contribution given by the material $\rho_{\mathbf{n}}$. We will treat the material itself as a black box in this section as we cover it in more details in Chapter 3. The reflected radiance equation may be applied recursively by expressing the incoming radiance as $L_i = L_r(\mathbf{q}, -\boldsymbol{\omega}_i)$ with \mathbf{q} the closest intersection point with a surface in the $\boldsymbol{\omega}_i$ direction. However, we will only concentrate on directly visible surfaces, and thus consider L_i merely as a light field function.

As shown in the previous sections, the human visual system is particularly sensitive to discontinuities and distortion patterns in the image, which are related to *spatial variations* of luminance. Since our goal is to reproduce these patterns from auxiliary buffers holding shape information (normals and depths), we first analyze how they occur using a *differential analysis* of reflected radiance. This then permits to identify how images may be deformed to reproduce similar discontinuities and distortion patterns in an artistically-controllable manner [VBFG12].¹²

Differential analysis The reflected radiance integral in Equation 2.6 is not directly amenable to differentiation since the domain of integration $\Omega_{\mathbf{n}}$ itself varies in the picture plane (i.e., \mathbf{n} is a function of \mathbf{x}). We thus perform a change of variable $\boldsymbol{\omega}'_i \leftarrow R\boldsymbol{\omega}_i$ where R is the smallest rotation that aligns the normal \mathbf{n} with $\mathbf{e}_2 = (0, 0, 1)$ (similarly, $\boldsymbol{\omega}'_o \leftarrow R\boldsymbol{\omega}_o$). Assuming $L_e = 0$, this yields:

$$(2.7) \quad I(\mathbf{x}) = \int_{\Omega} \rho(\mathbf{p}, R\boldsymbol{\omega}_o, \boldsymbol{\omega}'_i) L_i(\mathbf{p}, R^{-1}\boldsymbol{\omega}'_i) (\mathbf{z} \cdot \boldsymbol{\omega}'_i) d\boldsymbol{\omega}'_i,$$

where the material ρ is now defined in canonical space, which assumes that the surface normal is aligned with \mathbf{e}_2 , and Ω is the hemisphere of directions around that axis. For ease of notation, we omit function variables and write Equation 2.7 as $I = \int_{\Omega} \rho L_i d\Omega$, with $d\Omega = (\mathbf{z} \cdot \boldsymbol{\omega}'_i) d\boldsymbol{\omega}'_i$. The derivative of I along the \mathbf{x} axis in the picture plane is then given by:

$$(2.8) \quad \partial_{\mathbf{x}} I = \int_{\Omega} L_i \partial_{\mathbf{x}} \rho d\Omega + \int_{\Omega} \rho \partial_{\mathbf{x}} L_i d\Omega, \quad \text{with} \quad \begin{cases} \partial_{\mathbf{x}} \rho &= (\partial_{\mathbf{x}} \mathbf{p} \cdot \partial_{\mathbf{p}} \rho) + (\partial_{\mathbf{x}} R \boldsymbol{\omega}_o \cdot \partial_{\boldsymbol{\omega}'_i} \rho) \\ \partial_{\mathbf{x}} L_i &= (\partial_{\mathbf{x}} \mathbf{p} \cdot \partial_{\mathbf{p}} L_i) + (\partial_{\mathbf{x}} R^{-1} \boldsymbol{\omega}'_i \cdot \partial_{\boldsymbol{\omega}'_i} L_i) \end{cases}$$

where $\partial_{\mathbf{x}}$ stands for $\partial/\partial\mathbf{x}$, whereas $\partial_{\mathbf{x}} \rho$ and $\partial_{\mathbf{x}} L_i$ denote material and lighting variations respectively, each decomposed into spatial and angular gradient terms. The main difference with previous differential analysis of the rendering equation (e.g., [52]) is that we express it directly in the picture plane. As detailed in the original paper, the derivatives of R and its inverse may be expressed in terms of cross products between the normal and its derivative, yielding:

$$(2.9) \quad \nabla I = J_{\mathbf{p}} V_{\mathbf{p}} + J_{\mathbf{n}} V_{\mathbf{n}} \quad \text{with} \quad \begin{cases} J_{\mathbf{p}} &= [\partial_{\mathbf{x}} \mathbf{p}, \partial_{\mathbf{y}} \mathbf{p}]^T \\ J_{\mathbf{n}} &= [\partial_{\mathbf{x}} \mathbf{n} \times \mathbf{n}, \partial_{\mathbf{y}} \mathbf{n} \times \mathbf{n}]^T \end{cases}$$

¹¹The case of transmitted radiance is very similar and is not included for brevity. Likewise lens and aperture effects are not included for simplicity and we consider Equation 2.6 to act per wavelength or color channel.

¹²This project has been developed with Romain Vergne during his PostDoc.

where $V_{\mathbf{p}}$ (resp. $V_{\mathbf{n}}$) encapsulates spatial (resp. angular) variations of both material and lighting¹³, and $J_{\mathbf{p}}$ (resp. $J_{\mathbf{n}}$) is a Jacobian that characterizes the effects of surface slant (resp. curvature) on the image gradient. We call these special Jacobians *surface flows*.

The first-order flow $J_{\mathbf{p}}$ affects 3D *spatial* variations of either the material (e.g., due to textures) or of the lighting (e.g., due to shadowing or inter-reflections). In contrast, the second-order flow $J_{\mathbf{n}}$ affects 3D *angular* variations of the material (e.g., due to gloss) or of the lighting (e.g., due to light falloffs). We visualize both flows in the left of Figure 2.17, using their corresponding structure tensors $T_{\mathbf{p}} = J_{\mathbf{p}}J_{\mathbf{p}}^T$ and $T_{\mathbf{n}} = J_{\mathbf{n}}J_{\mathbf{n}}^T$, and the same color code as before.

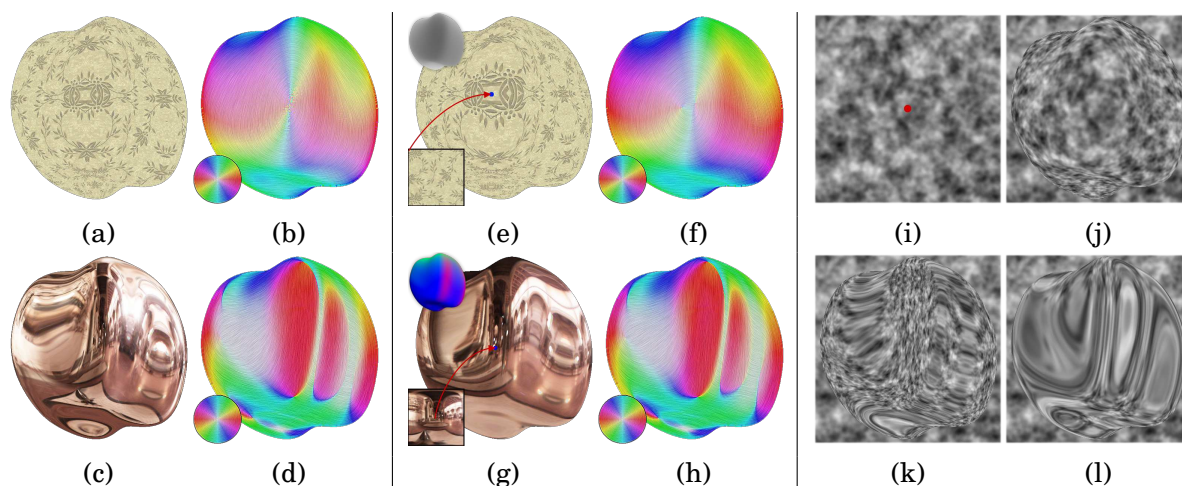


Figure 2.17: Conventional texture and environment mappings (a,c) produce variations related to first- and second-order surface flows. They are visualized in (b,d) with a LIC and color code where hue and saturation indicate direction and magnitude of the flow respectively. Our technique mimics texture and shading effects (e,g) using image processing alone. It takes depth and normal images as input (see insets) and uses them to deform images in ways that closely approximate surface flows (f,h). Deforming an input noise image (i) yields different appearances when the first-order operator (i) or the second-order operator (k,l) is used. The second-order operator permits to imitate reflection (k) or transmission (l) effects depending on the sign of the α parameter.

Deformation operators The key idea of our appearance design technique is to mimic the effect of surface flows using image deformation alone. We assume that the final image I results from an arbitrary combination of one or more texture images $I_{\mathbf{T}}$ and reflection images $I_{\mathbf{R}}$. A texture image is obtained by deforming an input texture pattern \mathbb{T} via $I_{\mathbf{T}}(\mathbf{x}) = \mathbb{T} \circ D_{\mathbf{p}}(\mathbf{x})$, where $D_{\mathbf{p}}$ is a first-order deformation operator. Similarly, a reflection image is obtained by $I_{\mathbf{R}}(\mathbf{x}) = \mathbb{R} \circ D_{\mathbf{n}}(\mathbf{x})$, where \mathbb{R} holds a reflection pattern and $D_{\mathbf{n}}$ is a second-order deformation operator. Both deformation operators are considered invariant at one or more *anchor points* \mathbf{x}_i (i.e., $D(\mathbf{x}_i) = \mathbf{x}_i \forall \mathbf{x}_i$).

The connection with Equation 2.9 becomes apparent when we differentiate $I_{\mathbf{T}}$ and $I_{\mathbf{R}}$, yielding $\nabla I_{\mathbf{T}} = J_{\mathbf{p}}^D \nabla \mathbb{T}$ and $\nabla I_{\mathbf{R}} = J_{\mathbf{n}}^D \nabla \mathbb{R}$ respectively, where J^D denotes the Jacobian of a deformation operator. The problem now amounts to find deformation operators $D_{\mathbf{p}}$ and $D_{\mathbf{n}}$ such that $J_{\mathbf{p}}^D$ and $J_{\mathbf{n}}^D$ approximate $J_{\mathbf{p}}$ and $J_{\mathbf{n}}$. This is illustrated in the middle panel of Figure 2.17: a texture or a reflection pattern is deformed given an input depth or normal image respectively, yielding surface

¹³We will replace $V_{\mathbf{p}}$ and $V_{\mathbf{n}}$ in the following, hence we do not need to provide explicit formula in this section.

flows that closely approximate those obtained with rendering; the deformation process is here controlled via a single anchor point (in blue).

The first-order deformation operator $D_{\mathbf{p}}$ is based on depths z and empirically defined by:

$$(2.10) \quad D_{\mathbf{p}}(\mathbf{x}) = \mathbf{x} + \sum_i w_i(\mathbf{x}) f_{\mathbf{p}}(z - z_i)(\mathbf{x} - \mathbf{x}_i) \quad \text{with} \quad f_{\mathbf{p}}(\Delta z) = \frac{2\Delta z(e^\gamma - 1)}{(e^\gamma + 1) - \Delta z(e^\gamma - 1)},$$

where the w_i balance the influence of anchor points and $f_{\mathbf{p}}$ is controlled by a magnitude parameter $\gamma \in \mathbb{R}^+$ (the deformation has no effect when $\gamma = 0$). The choice of an exponential function for $f_{\mathbf{p}}$ instead of a linear one has been made to prevent excessive distortions when $\Delta z < 0$ (i.e., for points in front of the anchor point), hence mimicking the non-linearity of depth in rendered 3D scenes. The effect of $D_{\mathbf{p}}$ is shown in Figure 2.17(i,j) on an input noise using a single anchor point: deformation gives the illusion that the noise acts as a texture on a 3D object (we use $\gamma = 5$).

The second-order deformation operator $D_{\mathbf{n}}$ is based on normals \mathbf{n} and empirically defined by:

$$(2.11) \quad D_{\mathbf{n}}(\mathbf{x}) = \mathbf{x} + \sum_i w_i(\mathbf{x}) f_{\mathbf{n}}(\|\bar{\mathbf{n}} - \bar{\mathbf{n}}_i\|)(\bar{\mathbf{n}} - \bar{\mathbf{n}}_i) \quad \text{with} \quad f_{\mathbf{n}}(\Delta \bar{\mathbf{n}}) = \frac{\alpha \Delta \bar{\mathbf{n}}}{e^{-\beta} + \Delta \bar{\mathbf{n}}(1 - e^{-\beta})},$$

where $\bar{\mathbf{n}}$ denotes the 2D projection of a screen-space normal onto the image plane, and $f_{\mathbf{n}}$ is a function of relative normal variations controlled via two parameters: $\alpha \in \mathbb{R}$ controls the signed magnitude of the deformation (positive and negative values mimic reflection and refraction patterns respectively), while $\beta \in \mathbb{R}^+$ controls its rate (how much distortion is increased toward silhouettes, the function being linear when $\beta = 0$). The effect of $D_{\mathbf{n}}$ is shown on the same input noise as before using one anchor point in Figure 2.17(k,l): deformation gives the illusion that the noise is either reflected (in (k), using $\alpha = 1, \beta = 0$) or transmitted (in (l), using $\alpha = 1, \beta = 0$).

Drawing tools The pair of first- and second-order deformation operators $D_{\mathbf{p}}$ and $D_{\mathbf{n}}$ may now be used in drawing tools acting respectively on input texture T or reflection R patterns.

The *brush tool* allows users to paint shading in exactly the same way they would paint brush strokes in raster graphics software. It consists in repeatedly applying a brush pattern with a single anchor point in its center; the only difference is that the pattern is deformed before being applied. When used with $D_{\mathbf{p}}$, it permits to create decal textures as shown in Figure 2.18(a,b), with the pattern T being manually rotated by the artist. When applied with $D_{\mathbf{n}}$, it permits to create plausible highlights as shown in Figure 2.18(c,d), with roughness controlled by varying the blur applied to the pattern R .

The *gradient tool* directly takes its inspiration from the classic linear gradient tool, except that the color gradient is deformed prior to its application. It makes use of two anchor points \mathbf{x}_0 and \mathbf{x}_1 , one for each extremity of the gradient, and requires two weighting functions w_0 and $w_1 = 1 - w_0$. The weight w_0 for a point \mathbf{x} is obtained by projecting it perpendicularly onto the $[\mathbf{x}_0, \mathbf{x}_1]$ segment, and performing a linear interpolation (clamped at extremities). When applied with $D_{\mathbf{n}}$, it permits to create shading gradients that suggest a directional light source as shown in Figure 2.18(e) where a smooth gradient pattern R is used to mimic a diffuse material.

The *image tool* allows users to distort an input texture using multiple anchor points, with weights controlled by Shepard interpolants [53]. This is particularly useful to mimic refraction effects as shown in Figure 2.18(f), with anchor points precisely controlling where specific points of the background pattern R appear by transmission.

The different tools may be combined to compose diverse and complex appearances, as shown in Figure 2.22. Multiple brush strokes of different sizes and smoothness may be applied to give a




Figure 2.18: The brush tool permits to deformed textures (a,b) or shading patterns (c,d) with controllable sizes (red contours) and locations (blue dots). Textures are oriented by the user to position decals; shading patterns are blurred to convey different materials. The gradient tool (e) permits to create shading patterns by deforming a linear gradient (red curve) controlled by 2 anchor points. The image tool (f) gives control over transmission effects: the background image is deformed using 3 anchor points. Color insets visualize the weight functions for each anchor point.

cartoonish look as in Figure 2.22a: here we use normals estimated from drawn contours output by the LUMO system [54]. First- and second-order deformations may be combined to achieve a plausible appearance such as in Figure 2.22b, using texture images to either mimic reflectance colors or additional surface details (i.e., normal maps). The input depth or normal buffers may also be altered to give the illusion of complex material appearances like translucency as shown in Figure 2.22c: the normal buffer is here perturbed with a noise prior to deformation only for the mimicking of subsurface effects, as suggested by the work of Motoyoshi [55].

Comparison to previous work Our approach is not the only solution for creating or manipulating appearance: another strategy consists in starting from existing 3D data or estimating them from images. For instance, inverse rendering techniques [56–58] rely on known 3D shape and material characteristics along with user scribbles to infer plausible illumination environments through a constrained optimization process. These methods considerably reduce the amount of trial and error required to reach a satisfying result, but still offer a limited control on final appearance and are restricted to physically-realistic results. Ritschel et al. [59, 60] provide solutions to depart from the laws of physics by deforming shading effects such as reflections, shadows or textures with user-controlled handles. However, their methods are still bound to the specification of a 3D scene (objects, lights and materials); hence shading effects cannot be created as with drawing applications, but only modified. The method of Khan et al. [43] estimates depth and normal buffers from photographs using simple heuristics, but then mostly rely on a re-rendering of the 3D scene using a modified material for instance, which shares the aforementioned limitations.

I would argue that for the design of appearance in still images, a more direct and less constrained artistic control is most often desired. The method of Yeung et al. [42] is the one closest to ours in this respect: the background image gets distorted to mimic the effect of refraction without having to estimate 3D scene data. The solution remains limited to the matting of transparent objects; it would be interesting to adapt it to more general types of appearance manipulations. One downside of such direct manipulation techniques is that they are not adapted to the design of appearance in dynamic scenes: the image deformation tools should not only take into account the distortions due to shape, but also those due to optic flow generated by motion.

2.2.2 MatCap manipulations

A simple technique for retaining artistic style in shading while permitting *dynamic* and arbitrary object modifications is the LitSphere [61]. It consists in the image of a sphere that has been either painted or captured from an image, and is more commonly called a 'MatCap' (for 'material capture') in computer graphics software (see the top row of Figure 2.16). It has found widespread adoption in sculpting software thanks to its simplicity¹⁴: the color at a surface point is obtained by looking up a color in the MatCap at a location given by the surface normal in screen-space. However, material or lighting cannot be modified *independently* with this approach; for instance, manipulations as simple as rotating lighting with respect to the view are not possible.

How could we keep the simplicity of MatCaps while permitting dynamic appearance manipulations? Assuming that an input MatCap conveys plausible material properties (Figure 2.19a), our approach is to decompose it into a pair of spherical image-based representations (Figure 2.19b). Our approach [ZMB⁺15] takes inspiration from perceptual experiments [62–64] that suggest that material perception likely depends on lighting and might be based on image cues such as the blurriness or contrast of highlights. Thanks to our decomposition, common appearance manipulations such as rotating lighting, or changing apparent material color and roughness are performed through simple image operators (Figure 2.19(c-e)).¹⁵

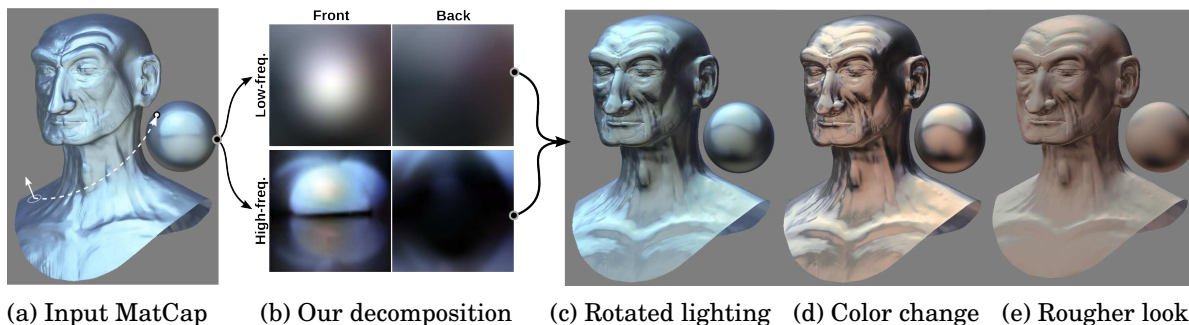


Figure 2.19: A MatCap is decomposed into a representation that permits dynamic appearance manipulation via image filters and transforms. (a) An input MatCap applied to a sculpted head model (with a lookup based on screen-space normals). (b) The low- & high-frequency (akin to diffuse & specular) components of our representation stored in dual paraboloid maps [65]. (c) A rotation of our representation orients lighting toward the top-left direction. (d) Color changes applied to each component. (e) A rougher-looking material obtained by blurring, warping and decreasing the intensity of the high-frequency component.

Low-/high-frequency separation Most MatCaps are given as a single image where diffuse and specular shading effects are blended together. Separating an image into these components without additional knowledge is inherently ambiguous. We provide a simple heuristic separation into low- and high-frequency components, which we find sufficient for our purpose. Our solution is based on a gray-scale morphological opening directly inspired by the work of Sternberg [66].


¹⁴Matcaps have been first introduced in ZBrush to visualize shape during the 3D sculpting process.

¹⁵This project has been developed with Carlos Zubiaga when he was a PhD student enrolled in the PRISM-ITN, with the help of Laurent Belcour when he was a PostDoc in our research team.

As shown in the top left of Figure 2.20, the input Matcap is first eroded, then dilated to yield the low-frequency component; the high-frequency component is then obtained by subtraction. This solution has the advantage of outputting *positive* components without requiring any parameter tuning, which we found in no other technique. On the downside, the method may attribute smooth specular effects (such as haziness, see section 3.2.1) to the low-frequency “diffuse” component.

Material as a filter In the following we consider low- and high-frequency components independently. If we assume that the sphere is seen in orthographic view, then the color at each pixel may be interpreted as radiance L_r reflected from a point on the sphere toward the view axis. However, neither the material $\rho_{\mathbf{n}}$ nor the incoming lighting L_i from Equation 2.6 are known in this case; nothing actually guarantees that the input MatCap may be decomposed into physically-plausible material and lighting functions. Nevertheless, we still need to make minimal assumptions on lighting and material in order to analyze shading in the input MatCap. We consider the simplified case where the lighting L_i is infinitely distant and $\rho_{\mathbf{n}}$ acts as a *locally* homogeneous radially-symmetric filter. In other words, the reflected radiance is obtained by convolution — $L_r = \rho_{\mathbf{n}} \otimes L_i$ — an assumption made in previous work for the analysis of light transport [67, 68]. Our goal is instead to use this simplified formulation to estimate a set of material properties that is sufficient for later MatCap manipulation.

With the aforementioned assumptions, a pixel inside the matcap image may be identified by angular coordinates (θ, ϕ) in a screen-space spherical parametrization where the poles are aligned with the camera axis. Since $\rho_{\mathbf{n}}$ is assumed to be radially-symmetric, our analysis may be confined to a single dimension, yielding a 1D convolution. We choose the θ dimension (corresponding to surface slant) since $\rho_{\mathbf{n}}$ may still vary toward grazing angles. As detailed in another work [ZBB⁺15], analyzing this 1D convolution in Fourier space leads to a 1D product, which results in formula relating 1D statistical moments of reflected radiance, to moments of lighting and material. Omitting the dependence on (θ, ϕ) , we obtain the following moment relations:



$$(2.12) \quad \begin{aligned} K[L_r] &= K[L_i^\phi] \times K_\rho, \\ E[\bar{L}_r] &= E[\bar{L}_i^\phi] + E[\bar{\rho}_\theta], \\ \text{Var}[\bar{L}_r] &= \text{Var}[\bar{L}_i^\phi] + \text{Var}[\bar{\rho}_\theta], \end{aligned}$$

where K denotes the energy, hat functions are normalized by energy (e.g., $\bar{L}_r = \frac{L_r}{K[L_r]}$), and E and Var stand for statistical mean and variance respectively. Moreover, we have made the assumption that the material is separable with $\rho_{\mathbf{n}} = K_\rho \times \rho_\theta \times \rho_\phi$, where K_ρ denotes the material color, and ρ_θ and ρ_ϕ represent slices in the θ and ϕ dimensions respectively. The ϕ superscript in L_i^ϕ then denotes the product of L_i with ρ_ϕ , integrated along the ϕ dimension.

Intuitively, Equation 2.12 tells us that the material acts as a filter at the (θ, ϕ) location: it *colors* the incoming lighting by multiplying its intensity by K_ρ ; it *warps* it by displacing its mean by $E[\bar{\rho}_\theta]$, and it *blurs* it by increasing its variance by $\text{Var}[\bar{\rho}_\theta]$.

MatCap decomposition Considering the material color to be constant over the sphere, we estimate it by a simple white balancing procedure: pick the pixel with highest luminance and assume that $K[L_i^\phi] = 1$, yielding $K_\rho = K[L_r]$ at that point. Such a white-balanced MatCap component is shown in the bottom of Figure 2.20; it corresponds to the MatCap of Figure 2.19a.

We next focus on the material variance, and again consider it to be constant over the sphere. Following the same approach, we assume there are patterns in the lighting where $\text{Var}[\bar{L}_i^\phi] = 0$,

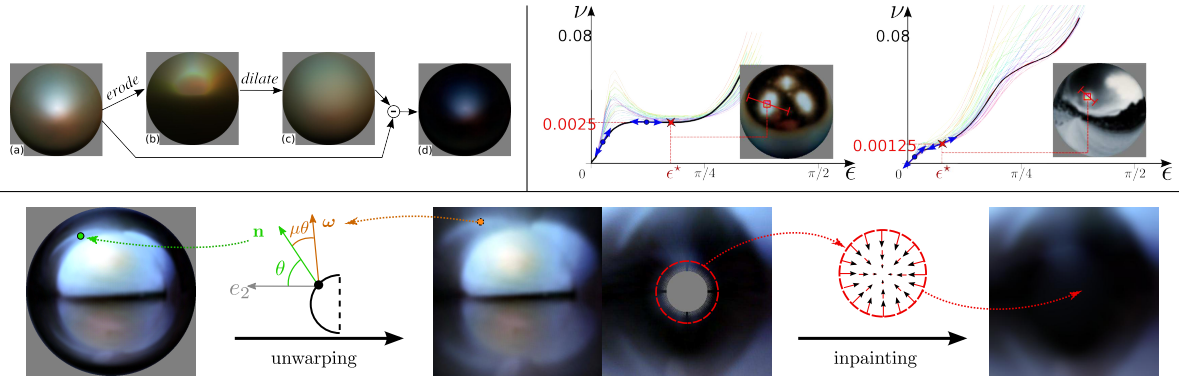


Figure 2.20: Top left: (a) An input MatCap is (b) eroded then (c) dilated to extract its low-frequency component. The high-frequency component is obtained by (d) subtracting the low-frequency component from the input MatCap. Top right: The material variance for a component is estimated (red cross) by analyzing image variance for a number of pixels at increasing scales (colored curves) and keeping the minimum variance (on the black curve). Material variance is either taken at the first minimum, or at the second inflection point (blue arrows) if no minimum is within reach. Bottom: A dual paraboloid map is filled by first warping each texel \mathbf{q} to a normal $\mathbf{n}_{\mathbf{q}}$ and looking up the color from an input, white-balanced MatCap component. This leaves an empty region in the back paraboloid map (the "blind spot") that is filled with radial inpainting.

which corresponds to point or line light sources,¹⁶ the material radiance may then be directly read off at the point of minimal variance using $\text{Var}[\bar{\rho}_\theta] = \text{Var}[\bar{L}_r]$. Unfortunately, the variance $\text{Var}[\bar{L}_r]$ depends in practice on the size ϵ of the window inside which it is computed, as shown in the top right panel of Figure 2.20; in particular, for small values of ϵ , variance is governed by the window itself. Our solution (indicated by a red star) is to vary window size until a minimum variance value is encountered; if there is none, we pick the second inflexion point in scale space.

In our analysis of the moments of measured real-world materials [ZBB⁺15], we have further observed that the material mean may be approximated as a linear function of the form $E[\bar{\rho}_\theta] = \mu\theta$, with $\mu = 0$ corresponding to the ideal diffuse case and $\mu = -1$ corresponding to the ideal specular case. Moreover, we have shown that the material mean slope μ is correlated to the material variance: $\mu \approx 1 - 0.3\text{Var}[\bar{\rho}_\theta] - 1.1\text{Var}[\bar{\rho}_\theta]^2$. We thus use our estimate of $\text{Var}[\bar{\rho}_\theta]$ to compute μ , which is used in turn to convert a white-balanced MatCap to a spherical representation. As shown at the bottom of Figure 2.20, we store the result of the conversion process in a dual paraboloid map [65]. Each texel in the paraboloid map corresponds to a direction $\boldsymbol{\omega}$, which we rotate to obtain the corresponding screen-space normal $\mathbf{n} = \text{rot}_{\mathbf{t}, -\mu\theta}(\boldsymbol{\omega})$ where $\mathbf{t} = \frac{\mathbf{e}_2 \times \boldsymbol{\omega}}{\|\mathbf{e}_2 \times \boldsymbol{\omega}\|}$, $\theta = \frac{a \cos(\mathbf{e}_2 \cdot \boldsymbol{\omega})}{1 + \mu}$ and $\mathbf{e}_2 = (0, 0, 1)$ stands for the (fixed) view vector in screen space. Since for each texel we end up with a different rotation angle, the resulting transformation is indeed an image warping. Inevitably, a disc-shaped region on the back-side of the dual paraboloid map will receive no color values. We call it the blind spot and its size increases with material variance: hence a wider blind spot corresponds to lower frequency content. It is thus reasonable to smoothly fill the back paraboloid map using inpainting. In practice, we apply Poisson image editing [69] with a radial guiding gradient that propagates boundary colors of the blind spot toward its center.

¹⁶We show in the paper that Equation 2.12 is also valid when considering derivatives of the lighting, such that material variance may also be estimated at the edges due to sharp area light boundaries.

Dynamic MatCaps Rendering using our decomposition works the other way around: the color at a surface point is given as a function of its screen-space normal \mathbf{n} . For each component, we first map \mathbf{n} to a direction $\boldsymbol{\omega}$ in the sphere using $\boldsymbol{\omega} = \text{rot}_{\mathbf{t}, \mu\theta}(\mathbf{n})$, with $\mathbf{t} = \frac{\mathbf{e}_2 \times \mathbf{n}}{\|\mathbf{e}_2 \times \mathbf{n}\|}$ and $\theta = \text{acos}(\mathbf{e}_2 \cdot \mathbf{n})$. The shading color is then obtained by a lookup in the dual paraboloid map based on $\boldsymbol{\omega}$, followed by a multiplication by K_ρ . The low- and high-frequency components are finally added together.

The main advantage of our representation is that it may be manipulated dynamically using simple image processing techniques. This is illustrated in Figure 2.23a where different MatCaps are applied to different parts of an object. A first benefit of our approach is that the highlights in input MatCaps may now be aligned to make the lighting direction coherent (Figure 2.23b): this is done by manual rotations of our spherical representation. The material may also be edited (Figure 2.23c): material colors are modified and more importance is given to either the low- or high-frequency component; the apparent roughness may also be increased by using a higher μ and blurring the dual paraboloid map to increase the material variance; or it may be decreased by lowering μ and sharpening the spherical image.¹⁷ Finally, all MatCaps are rotated together (Figure 2.23d) to give the impression of a global rotation of the environment.

Remaining limitations We have made a number of assumptions on lighting properties for MatCap decomposition; however, these might not always be met, in which case material parameters will be over- or under-estimated. This will not prevent our approach from working, since it will be equivalent to having a slightly sharper or blurrier lighting. Interestingly, recent psychophysical studies (e.g., [63]) show that different material percepts may be elicited only by changing lighting content. This suggests that our approach could be in accordance with visual perception.

A more fundamental limitation of our approach is that it works for one object at a time. Hence it does not offer any solution to mimic local light transport, such as inter-reflections or shadowing effects. A challenging direction for future work would be to find the minimal additional user input necessary to mimic and control these effects.

Auxiliary buffers are also output by offline rendering engines, along with visibility buffers and shading buffers for different material components, as shown in the top of Figure 2.21. Artists may adjust the intensity of each shading buffer at the compositing stage, or add highlights and depth of field effects using auxiliary buffers, without having to re-render the entire scene. However, any information about existing lighting is lost past the rendering stage; hence more complex manipulations such as changes in surface shape or material roughness are not *directly* possible. We have investigated the possibility of manipulating appearance further at the compositing stage by first reconstructing MatCaps for different object parts of a rendered scene [ZGVB16].¹⁸

Normal editing in post-process Our goal is to grant the editing of normals (i.e., local shape) in real-time in post-process, with proper repercussions in diffuse and reflection (either specular or glossy) buffers. An exact reconstruction of these shading buffers would be impossible as we lack much of the necessary 3D data. We instead strive for a plausible result, ensuring that the input diffuse and reflection shading buffers are recovered when reverting to the original normals.

The key idea of our method is to reconstruct a pair of pre-filtered environments per object/material: one for the diffuse component, the other for the reflection component. This is

¹⁷Sharpening remains limited since practical solutions tend to raise image noise.

¹⁸This project has been developed by Carlos Zubiaga when he was a PhD student enrolled in the PRISM-ITN.

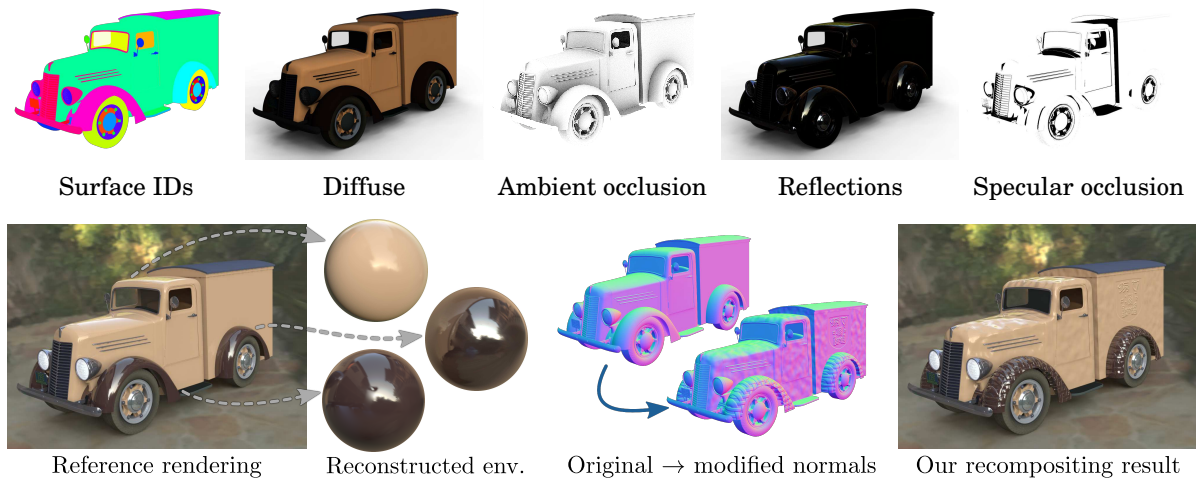


Figure 2.21: Our method permits to modify surface shape by making use of the shading and auxiliary buffers output by modern renderers (top row). We first reconstruct shading environments for each object/material combination of the Truck scene, relying on normal and shading buffers. When normals are modified by the compositing artist, the color image is recomposited in real-time, enabling interactive exploration. Our method reproduces inter-reflections between objects, as seen when comparing the reconstructed environments for rear and front mudguards.

illustrated in the bottom left of Figure 2.21 where both components are visualized together using MatCaps: observe in particular how inter-reflections affect MatCaps for a same object (here the mudguard) at different locations. As detailed in the paper, we use different reconstruction methods for each component. The diffuse environment is estimated using low-order spherical harmonics [70] whose coefficients are optimized using quadratic programming. The reflection environment is obtained by partitioning the input set of pixels into patches (squares and triangles) according to normal and depth discontinuities, then mapping them to spherical patches (quads and triangles) via reflection; overlapping patches are weighted based on area (favoring smaller patches), while holes are filled using inpainting as before.

The resulting pair of diffuse and reflection environments are then looked up to yield new shading colors from arbitrarily modified normals, as shown in the bottom right of Figure 2.21. Modifying local shape hence becomes a matter of recompositing the reconstructed shading buffers, which may be achieved in real-time even on moderately efficient graphics hardware.

Outlook on expressive compositing Our approach is a first step toward the editing of surface shape (and more generally object appearance) at the compositing stage, which we believe is an important and challenging problem. One may object that with the increasing performance of physically-accurate rendering engines, modifying shapes or materials and re-rendering might provide sufficiently interactive feedback in the very near future. Even though a solution based on render outputs has the drawback of only mimicking physically-plausible results, it possesses two key advantages: it requires much less computational power, but more importantly it opens the way to non-physical manipulations which are routinely used in post process. Moreover, the rendering and compositing stages are clearly *separate* in practice, involving different artistic skills and types of assets (i.e., 3D scenes vs render buffers). Providing compositing artists with more control over appearance will thus require specific solutions.




Figure 2.22: Our approach permits to create an image from scratch (a) by first extracting a normal map from a drawing (we use outputs from the LUMO system [54]), then painting a few highlights using the brush tool. Multiple deformations may be combined in cascade to obtain rich appearances as shown in (b), where surface are added by applying the texture tool with a normal map. The full process uses both types of operators multiple times (blue arrows correspond to the 1st order operator, orange arrows to the 2nd order operator). Complex translucency effects (c) may be imitated with a pair of layers: the background image is deformed to depict reflections (left); then blurred, colored and deformed using perturbed normals to depict translucency (right); finally both layers are added together.




Figure 2.23: Using material IDs (center), three MatCaps are assigned to a robot object (a). Our method permits to align their main highlight via individual rotations (b) and change their material properties (c). All three MatCaps are rotated together in (d).

2.3 Discussion

Moving away from edges One assumption we relied on in Section 2.1.1 is that complex appearance may be obtained by controlling what happens at image edges (i.e., strong gradient variations). Among the different types of edges, occluding contours reveal to be mandatory, as they permit to partition the image into different objects, and even separate overlapping parts of a same object (i.e., with internal contours). However, the use of a diffusion approach remains intrinsically limited to the creation of *smooth-shaded* images, with the potential addition of rather simple highlights. Other potentially useful edges are (more or less smooth) gradient discontinuities that may be used to represent cast shadows or surface creases, even though they remain somewhat difficult to control with current solutions. In my view, most other types of curves remain difficult to interpret and still require significant trial and error before the desired effects may be achieved. More importantly, the color gradients away from curves are most often difficult to control, so that when the geometry of a control curve is changed, its diffusion properties must also be adjusted.

I would thus argue that a new approach to the vectorial control of appearance is needed to achieve more varied results. It should preserve the qualities of the diffusion curves approach, namely a sparse set of control primitives, including occluding contours and crease curves. But it should also provide more local control and flexibility to the artist, and produce predictable results when the control primitives are modified. The work of 2D artists is most often iterative, as shown in Figure 2.24: the artist might go through different steps, erasing, modifying or refining her work. Taking inspiration from this iterative process, while at the same time providing accurate controls via vectorial primitives, promises to be a challenging direction of research.




Figure 2.24: Excerpt of Exercise 41 from CGCookie website (<https://cgcookie.com/exercise/exercise-41-color-shifting/>): the proposed *iterative* drawing process starts with outlines, then a solid base is added, after which color gradients are progressively established by successive refinements, adding subtle highlights, rims and shadows in the last steps.

Conveying shape and materials The compression and stretching of texture and shading patterns presented in Sections 2.1.2 and 2.2.1 revealed to be very efficient at conveying or manipulating *shape*. However, the methods we proposed could be refined in a number of ways. First, occluding contours should be explicitly and automatically dealt with. Currently, they must be manually provided when working with image flows, whereas they emerge implicitly when using surface flows. Second, multiple compression and stretching layers should be considered and correlated. For instance, the deformations due to reflected and transmitted light at a transparent object surface should be visually correlated since they are constrained by the same surface shape.

The same is true for patterns due to textures and reflections. Third, distortions should occur at different scales, and through this process preserve more global structures such as symmetries or regularities.

For the depiction or manipulation of *material* appearance, we have shown in Section 2.2.2 that a promising approach consists in studying the filtering properties of a given material component directly in images. However, we did so only for spheres, which discards complex shape variations from the onset. Extending our approach to arbitrary surfaces does not seem straightforward, since local stretching and compressions will inevitably alter the apparent filtering properties of a material. Another limitation concerns the simplicity of the filter shape, as we have only considered Gaussian functions; in particular, more heavy-tailed filters might provide for richer material appearances, as discussed in more details in Section 3.2.1.

An exciting avenue of future work would not only be to address these limitations, but also to integrate these seemingly disparate techniques together. Once texture and shading flows have been analyzed, the resulting image structure stores how neighbor pixels are distorted by the depicted 3D shape. The effect of a material as a filter could then be studied after these distortions have been discarded (in a canonical space), hence permitting to extend our approach to surfaces more complex than spheres. Interestingly, a similar idea has been proposed by Jonathan Miller in his book on reflections in art [71], as reproduced in Figure 2.25. He wonders what makes the depiction of the parabolic mirror in Campin’s painting so compelling, to the point where it appears perfectly smooth. Clipping out the mirror from its frame already considerably reduces the impression of reflection; undistorting it makes it look like a mere circular window. If the same approach were carried out on a rough reflective surface, the remaining undistorted image would appear blurry or hazy, in a way dependent on material filter characteristics. A last challenging issue would still remain: the identification of occluding contours (which have the effect of “clipping” neighborhoods) when they are not directly provided.



“Saint John the Baptist and the Franciscan Heinrich von Werl” by Robert Campin (detail)

Clipped-out mirror Undistorted image

Figure 2.25: Jonathan Miller, in his book on reflections [71], asks how a perfect mirror must be depicted to avoid mistaking it from a mere window on a different scene. To this end, he clips out the image of the curved mirror and manually discards its distortions, revealing how context and subtle deformations can drastically affect our perception.

Artistic layering In Section 2.2 we considered applications where we get access to geometric and shading information in the form of auxiliary buffers, themselves obtained in a variety of ways. Once such additional information is available, it becomes tempting to use it to separate

image in layers that correspond to different lighting or shading effects: texture may be separated from shading as in Section 2.2.1, diffuse components may be separated from specular ones as in Section 2.2.2, etc. Moreover, when multiple objects reflect each others or cast shadows onto nearby objects, the potential combinations of layers grows accordingly to include direct and indirect lighting effects. Besides the choice of layers, their ordering may also matter once they are blended in a non-commutative way, which is the case with traditional painting for instance.

Looking at demo painting sessions such as the still life of Figure 2.26 reveals that artists tend to build images in successive layers that significantly differ from what is usually considered at the compositing stage. Some objects are depicted in reasonable details (diffuse shading, inter-reflections, highlights) early on, while some other object parts appear only toward the very end. Accurate contours and highlights are usually added as (numerous) final touches that bring coherence and crispness to the artwork. Nevertheless, the process may vary tremendously from one artist to the other. There is thus much to learn from artistic practice when it comes to layering. On a perceptual level, the relation between these artistic layers and our visual understanding of complex shadowing and inter-reflection effects is very much intriguing on its own.



Figure 2.26: Screenshots from a still life oil painting demo by Jos van Riswick (https://www.youtube.com/watch?v=bitD3_YKNdM). The artist starts by a rough sketch then paints mostly diffuse object colors. However the painting process is not really organized according to lighting effects: coarse reflections appear early on the bowl and bottle, whereas entire parts of the kettle are only depicted later. Detailed highlights and reflections do appear in the last stages though, after soft shadows and diffuse gradients have been carefully crafted. Interestingly, the precise contours of most objects are only painted at the very end of the demo session.

Image structure In my view, future work on the creation or manipulation of appearance will benefit from decomposing images into perceptually-meaningful elements that are laid out in a comprehensive structure. This structure might be extracted from an input image, or created with vector-like controls including (but not limited to) boundary curves. The main idea would be to go away from the raw color manipulation approach of gradient meshes or diffusion curves, toward cascaded editing where different appearance models are nested inside each other. A good example is the perspective grid tool of Illustrator¹⁹, which lets the artist directly specify vanishing lines that distort the whole image canvas but still permits local edits of the distorted shapes. With a rich-enough structuration such as the one overlaid in Figure 2.26, one might be able to modify apparent shape, material properties and even lighting, hence augmenting an image with dynamic vectorial appearance controls.

¹⁹<https://helpx.adobe.com/illustrator/using/perspective-drawing.html>

There is a notable difference between working in image space as discussed in this chapter, and working in object space as discussed in the next chapter: in the former, image structure does not have to be accurate. This is well exemplified by countless paintings and drawings that use more or less naive perspective effects for instance; even though the perspective as a whole is not accurate²⁰, subparts of the image will clearly indicate perspective effects, unless the goal of the artist is precisely to flatten out the image as in the Cubist movement. In other words, this means that the tools provided to artists should not constrain them to produce accurate depictions, but leave them enough degrees of freedom to explore different ideas — while making sure that perceptually-relevant patterns are still generated, at least locally.



“Rippled surface” by M. C. Escher

Perspective effects

Distorsion effects

Figure 2.27: This artwork by Escher conveys the image of a reflecting pond on a rainy day with an economy of drawing primitives. Yet, they subtly convey perspective and ripple patterns through very accurate deformations.


²⁰It actually takes a lot of discipline for an artist to draw a correct perspective.

OBJECT-SPACE TECHNIQUES

When precision and automation are criteria of choice, the rendering of a full 3D scene is by far superior to image-based techniques. Appearance may then be controlled in at least two different ways. First, radiance may be tweaked in light transport as described in Section 3.1, where we assume that shapes, light sources and materials are provided. Second, materials may be designed to elicit specific appearances as detailed in Section 3.2, where we focus not only on microscopic surface properties but also on the visual patterns they produce. As discussed in Section 3.3, these seemingly different approaches may overlap in the context of multi-scale rendering.

3.1 Radiance tweaking

We now assume a full 3D scene is provided, with objects defined by their surface geometry and material properties, and light sources positioned in the environment. Rendering such a scene consists in solving Equation 2.6, which is referred to as “light transport” in the graphics literature and essentially relies on geometric optics: light is assumed to travel along rays. There are several strategies for performing such simulations, either starting from light sources or from the camera sensor, with the goal of finding light paths (concatenations of light rays) of greatest importance.



Our goal in this section is to devise methods for altering light transport, hence tweaking reflected or transmitted radiance in ways that better convey shape properties in particular. Consequently, we mostly focus on directly visible surfaces, which are identified by the intersection between a beam emanated from the viewpoint and the closest surface in the scene as shown in the inset figure. The intersected surface patch may not necessarily be flat, and we assume that its material properties are given beforehand; the design of materials is deferred to Section 3.2.

In many application domains such as games, visual effects or scientific illustrations, light transport simulation is used as a “raw material” that may then be manipulated in non-physical




Figure 3.1: Radiance tweaking in visual effects and games (©Autodesk): (a) classic shading (b) with ambient occlusion (c) stored in texture space. In scientific visualization [72]: (d) ambient occlusion reveals cavities, but (e) local curvature remains important.

ways. A common example is ambient occlusion [73]: it consists in computing the soft shadowing due to a purely ambient lighting (e.g., as in a fog) on a purely diffuse (i.e., Lambertian) object. As shown in Figure 3.1(a-c), it naturally darkens cavities and crevasses in a view-independent manner, which permits to store it in a texture. The method is prized by 3D artists for its ability to reveal surface shape details that may otherwise go unnoticed, to the point where it is used in arbitrary lighting environments with arbitrary materials, which is obviously physically wrong.

Another example is mean curvature shading [74], which consists in the brightening or darkening of surface reflectance depending on the local surface curvature. It does not have a direct physical meaning and is most often used for 3D illustrative shading. As shown in Figure 3.1(d-e) where it is compared to ambient occlusion, it better conveys shallow surface details but is less efficient at depicting deep cavities. A limitation of both methods is that they do not depend on the viewpoint, which prevents them from being used for altering light transport on surfaces made of glossy or specular materials.

3.1.1 Local warping and scaling

Skilled artists are able to convey shape through the subtle tweaking of shading and reflections. For instance in Archeology, fine surface characteristics are depicted in sharp relief as in Figure 3.2a. Medical illustrations often represent anatomical shape and surface details with great accuracy such as the lungs in Figure 3.2b. The original material and its appearance under natural lighting conditions are most often retained. Similar examples may also be found in artistic illustrations, such as in the work of Norman Rockwell or Burne Hogarth, as in Figure 3.2c: here, the folds of a garment cloth are skillfully reproduced, while reflectance characteristics are efficiently conveyed.

Recent work in human vision has shown that the perception of surface shape depends on how reflections get distorted [1, 75, 76]. This suggests a way to enhance shape depiction: in order to better reveal surface features from the current point of view, one can deform patterns of reflected lighting. However, an arbitrary deformation may also alter illumination properties; fortunately, evidence has been made that perception is relatively insensitive to local inconsistencies in




Figure 3.2: Three examples of artistic surface enhancement with complex materials and illumination: (a) A medieval shoe where every little detail has been enhanced; (b) A pair of lungs where small veins are exaggerated; (c) A character wearing shiny cloth exhibiting multiple folds.

illumination direction [77]. Our approach [VPB⁺09] is illustrated in Figure 3.3: an object is analyzed to extract surface curvatures, which are then converted into deformation parameters to guide a *light warping* algorithm that is directly incorporated inside the rendering equation.¹




Figure 3.3: The light warping technique relies on an analysis of screen-space normals to compute view-centered surface curvatures, which are then mapped to warping parameters, here visualized in grey-levels. The 3D object is then rendered using a locally warped environment lighting, which enhances surface shape depiction without compromising material appearance.

Shape descriptor Surface curvatures are extracted from a *view-centered* curvature tensor \mathbf{H} obtained by differentiating normals in screen-space: $\mathbf{H} = [\nabla_x \mathbf{g} \ \nabla_y \mathbf{g}]$, with $\mathbf{g} = \left(\frac{-n_0}{n_2}, \frac{-n_1}{n_2} \right)^T$ and $\mathbf{n} = (n_0, n_1, n_2)^T$. This simple and fast method permits to identify surface features as curvature properties, such as convexities and concavities shown in blue and orange respectively in Figure 3.3. It has the advantages of being independent of the underlying surface representation (meshes, implicit surfaces, etc) and adapted to the current viewing scale: in particular, when zooming in, new surface details are progressively identified. However, it also suffers from a few practical limitations: it requires heuristic treatments to forbid differentiation across occluding contours, and it only considers visible surface points.

¹This project has been developed in the context of the Animaré project with Romain Vergne and Romain Pacanowski during their PhD.




Figure 3.4: Left: the depicted curvature is exaggerated through reflected light patterns by expanding the region that is reflected off the surface (curving arrows), which in turn compresses reflections on the object compared to rendering without warping ('idle'). Right: The warping process modifies the light direction by (1) stereographic projection, followed by (2) scaling to a new position in stereographic space and (3) inverse stereographic projection back to the sphere of directions. The effect of principal curvature directions in step (2) is visualized in the stereographic plane: scaling maps projected lighting directions from the green circles to the orange ellipses.

Light warping The key idea of our warping approach is illustrated in Figure 3.4-left: the deformation of reflection patterns that naturally occurs on curved objects is exaggerated. In the simple case of a mirror reflection, one observes that changing the direction of incoming light rays has the effect of contracting a wider region of the environment, as if the object were more curved. This is achieved by modifying the reflected radiance equation in one of two ways: either by directly altering incoming light directions using a warping function W ; or by applying the inverse warping:

$$(3.1) \quad \begin{aligned} L_r'(\mathbf{p}, \boldsymbol{\omega}_o) &= \int_{\Omega_n} \rho_n(\mathbf{p}, \boldsymbol{\omega}_o, \boldsymbol{\omega}_i) L_i(\mathbf{p}, W(\mathbf{p}, \boldsymbol{\omega}_i)) (\mathbf{n} \cdot \boldsymbol{\omega}_i) d\boldsymbol{\omega}_i, \\ &= \int_{\Omega_n} \rho_n(\mathbf{p}, \boldsymbol{\omega}_o, W^{-1}(\mathbf{p}, \boldsymbol{\omega}_i)) L_i(\mathbf{p}, \boldsymbol{\omega}_i) (\mathbf{n} \cdot W^{-1}(\mathbf{p}, \boldsymbol{\omega}_i)) J d\boldsymbol{\omega}_i, \end{aligned}$$

where L_r' is the altered reflected radiance and J is the Jacobian of W^{-1} . The first equation is used when sampling the material ρ_n , the second when sampling the incoming lighting L_i .

The warping function W depends on the curvature tensor \mathbf{H} at \mathbf{p} . As illustrated in Figure 3.4-right, it is organized in three steps: the incoming light direction $\boldsymbol{\omega}_i$ is first projected onto the stereographic plane perpendicular to the \mathbf{e}_2 axis (i.e., the local view axis); the corresponding 2D vector is then scaled according to principal curvatures in the corresponding principal directions \mathbf{u} and \mathbf{v} ; it is finally turned back into a warped lighting direction using the inverse stereographic transformation. The resulting distorted reflection patterns convey curvature anisotropy, which is in accordance with perceptual studies on shape from shading [78].

Results Light warping is demonstrated in Figure 3.5, where the shape of an input 3D object made of a rough glossy material is enhanced in two different illumination settings. Observe how the enhancement remains coherent while the patterns of reflected lighting are completely different in each image. At the same time, material appearance is preserved.

We have investigated an alternative approach to surface enhancement where this time light *intensity* is altered instead of its direction, again in a way dependent on surface curvature. The



Figure 3.5: Our light warping approach enhances surface depiction while preserving the overall appearance of 3D objects, as exemplified with these two renderings that use drastically different illuminations. Observe how various surface features are properly enhanced in both settings: sharp features on the face, broad variations around shoulders, and rough details on the torso.

method is called *radiance scaling* [VPB⁺11]; it is not only based on surface curvature but also on material characteristics, which makes the method easy to adapt to different rendering scenarii.²

Radiance scaling The main idea of is to adjust reflected light intensity per incoming light direction using a scaling function denoted by $\sigma(\kappa, \delta)$, where $\kappa(\mathbf{p}) : \mathbb{R}^3 \rightarrow [-1, 1]$ remaps surface curvatures at \mathbf{p} so that -1 corresponds to maximum concavities, 0 to planar regions and 1 to maximum convexities; and $\delta(\omega_o, \omega_i) : \Omega^2 \rightarrow [0, 1]$ remaps surface reflectance so that 0 corresponds to minimum reflected intensity, and 1 to maximum reflected intensity. Intuitively, δ identifies the light directions that contribute the most to reflected intensity.

The scaling function is then simply introduced into the reflected radiance equation:


$$(3.2) \quad L'_r(\mathbf{p}, \omega_o) = \int_{\Omega_n} \sigma(\kappa, \delta) \rho_n(\mathbf{p}, \omega_o, \omega_i) L_i(\mathbf{p}, \omega_i) (\mathbf{n} \cdot \omega_i) d\omega_i.$$

We chose to use the following formula for the scaling function:

$$\sigma(\kappa, \delta) = \frac{\bar{\sigma}K + \delta(1 - \bar{\sigma}(1 + K))}{\bar{\sigma} + \delta(K - \bar{\sigma}(1 + K))} \text{ with } K = e^{\alpha\kappa}.$$

It is plotted in the inset figure: α controls the magnitude of the effect (here $\alpha\kappa \in \{-1, -0.5, 0, 0.5, 1\}$),

whereas $\bar{\sigma}$ is a scaling-invariant value that grants more accurate control over selected surface features ($\bar{\sigma} = 0.8$ in the figure). The visual effects of these parameters are shown in Figure 3.6.



Rendering scenarii The choice of reflectance mapping function δ depends on the rendering scenario. Taking the material function ρ_n as the sum of diffuse and specular terms, scaling may be applied independently to each term with separate magnitudes α_d and α_s . Since the diffuse term is commonly equal or close to a constant, it will tend to be dominated by the cosine term in

²This project has been developed in the context of the Animaré project with Romain Vergne and Romain Pacanowski during their PhD.




Figure 3.6: Controls provided by the scaling function. (a) The original rendering with no scaling ($\alpha = 0$). (b-d) Our enhanced result using two scaling coefficients ($\alpha = 10$ on top; $\alpha = 20$ at bottom) and three scaling-invariant points: (b) with $\bar{\sigma} = 0.5$, there is a good equilibrium between concavity and convexity enhancement; (c) with $\bar{\sigma} = 0.01$, convexities get more brightened; (d) with $\bar{\sigma} = 0.99$, shading gradients become reversed.

Equation 3.2; hence we simply use $\delta_d(\boldsymbol{\omega}_o, \boldsymbol{\omega}_i) = (\mathbf{n} \cdot \boldsymbol{\omega}_i)$. The specular term is more complex and usually specifies reflectances outside of the $[0, 1]$ range; we thus choose $\delta_s(\boldsymbol{\omega}_o, \boldsymbol{\omega}_i) = (\mathbf{h} \cdot \mathbf{n})$ where $\mathbf{h} = \frac{\boldsymbol{\omega}_o + \boldsymbol{\omega}_i}{\|\boldsymbol{\omega}_o + \boldsymbol{\omega}_i\|}$ is the half-way vector between $\boldsymbol{\omega}_o$ and $\boldsymbol{\omega}_i$.³ This approach allows artists to enhance diffuse and specular material components differently as illustrated in Figure 3.7.

Global illumination techniques are usually time-consuming processes. Various methods have thus been proposed to precompute and reuse radiance data. When the material is perfectly diffuse and the illumination environment sufficiently distant, Equation 3.2 may be approximated by:

$$(3.3) \quad L_r'(\mathbf{p}, \boldsymbol{\omega}_o) \approx \sigma(\kappa, \delta) \rho_d A(\mathbf{p}) \bar{L}(\mathbf{n}) \quad \text{where} \quad \begin{cases} A(\mathbf{p}) &= \int_{\Omega_n} V(\mathbf{p}, \boldsymbol{\omega}_i) (\mathbf{n} \cdot \boldsymbol{\omega}_i) d\boldsymbol{\omega}_i, \\ \bar{L}(\mathbf{n}) &= \int_{\Omega_n} L_i(\boldsymbol{\omega}_i) d\boldsymbol{\omega}_i, \end{cases}$$

with ρ_d the constant material term, A the ambient occlusion stored at each surface point \mathbf{p} and computed using an ambient white illumination and a local visibility function V , and \bar{L} an irradiance average stored in a prefiltered environment map and looked up using the local screen-space normal \mathbf{n} assuming no dependence on \mathbf{p} (i.e., distant lighting). For such an approximation, we use a reflectance mapping function that is independent of both light and view directions: $\delta_a = \frac{\bar{L}(\mathbf{n})}{\max_{\mathbf{m}} \bar{L}(\mathbf{m})}$. This choice is coherent with perfectly diffuse materials since in this case the light direction that contributes the most to reflected radiance is the normal direction on average. Figure 3.11a and 3.11b show results obtained with this approach⁴, where in the latter case we have used a constant lighting to visualize radiance scaling applied to ambient occlusion only.

Complex renderings involving global illumination produce inter-reflections as shown in Figure 3.11c, which depicts objects made of shiny materials. The same scaling method is easily applied not only to directly visible surfaces, but also to surfaces seen via reflection as shown in Figure 3.11d: observe in particular how dragon scales are enhanced everywhere they appear.

Remaining limitations A practical issue of local radiance tweaking techniques is that altering reflected radiance will usually break the reciprocity of light paths, and even energy conservation. Depending on the targeted application, this may become an important drawback such as when using rendering strategies that rely on such properties (e.g., bidirectional path tracing). Another limitation is that we rely on purely local surface properties. The class of reflected lighting

³We will investigate further the importance of the half-way vector in Section 3.2.

⁴This method is integrated in Meshlab as a plugin, and used to explore surface details on captured geometry.



Figure 3.7: The radiance scaling technique grants more control over the enhancement of different material terms. (a) A glossy object rendered using Ashikmin’s material model [79], with a zoomed view on the chest. (b) Applying radiance scaling only to the diffuse term enhances surface features away from highlights (e.g., it darkens concave stripes on the arms and chest). (c) Applying it only to the specular term enhances surface features in a different way (e.g., it brightens some of the concave stripes, and enhances foreshortened areas). (d) Combining both enhancements brings up all surface details (e.g., observe the alternations of bright and dark patterns on the chest).

patterns humans are able to rely on for perceiving shape is obviously much more diverse than local distortions, or alternations of bright and dark colors in convexities and concavities [80]. More complex patterns are likely to be dependent on the main illumination direction (e.g., [48, 76, 81]), material characteristics (e.g., [82, 83]) and motion (e.g., [26, 84]). In the next section, we remain focused on radiance tweaking, but this time taking larger surface neighborhoods into consideration.

3.1.2 Multi-scale enhancement

Scientific illustrations usually try to combine different objectives such as legibility and aesthetics [85]. With the advent of efficient geometric acquisition technology, automatic illustrative rendering techniques have started to become an interesting solution as they permit the interactive exploration of captured shapes. This has applications in the study of inscriptions on engraved stones or traces over human bones in Cultural Heritage for instance. The same requirement emerges with manually created geometry: artists spend a lot of time using modeling and sculpting software to work out fine shape details such as skin pores, wrinkles or grain; hence techniques that depict this geometric complexity without requiring additional manual painting are needed.

A classic approach for analyzing 3D shape is to perform a differential analysis on infinitesimal neighborhoods. We have already made use of surface curvatures — second-order properties — in Section 3.1.1. Higher-order properties are also of interest, for instance to identify surface inflections — third order properties. A variety of solutions exist for carrying out the analysis to larger neighborhoods [72, 86, 87], which may be seen as applying differential operators to smoothed-out shapes. Methods based on differential geometry excel at identifying thin surface

features such as ridges and valleys [88] or demarcating curves [89]. However, for very detailed shapes typical of acquired geometry, this often leads to a myriad of lines that hinder legibility. An alternative approach relies on a decomposition of surface normals into smooth and detail layers. Such a decomposition has first been applied for normal enhancement [90, 91], where a single relief layer is manipulated to exaggerate or attenuate surface details through shading. It has been extended to multi-scale analysis in the exaggerated shading technique [92], where different levels of smoothed normals are used to align a single light source at grazing angles and the object is rendered with a simple cosine-based shading. Even though these methods do convey information about surface features, they do not provide a *precise control* over which surface properties are depicted. For instance, normal enhancement tends to produce embossed renderings, while exaggerated shading tends to flatten the overall shape, which is problematic for shape exploration.

Methods based on differential geometry or normal decomposition do not provide enough control on the legibility and precision of identified surface features. I would argue that this is because the analyzed surface features are *centered at* a given point. In contrast, we propose to locate surface features (concavities, convexities and inflections) *in the vicinity* of a point. Since surface features may be of varying sizes, we have introduced an automatic surface relief analysis method [ABG⁺12] that analyzes 3D geometry at multiple scales. It permits to drive a variety of illustrative shading techniques using high-level controls.⁵

Surface relief analysis We consider *relief* in a standard way as a height field relative to a smooth (and not necessarily planar) base surface. This is a reasonable assumption both in scientific applications where surface details are often very shallow (as in Archeology) or naturally defined as height data (as in Cartography), and in artistic applications where most surface details are stored in displacement maps. A surface S is hence defined as a smooth base surface B displaced along its normal field \mathbf{n}_B by a height function h . Mathematically, any point $\mathbf{b} \in B$ yields a point \mathbf{p} on the surface S with: $\mathbf{p} = \mathbf{b} + h(\mathbf{b})\mathbf{n}_B(\mathbf{b})$. When such a base/relief decomposition is not provided as is the case of acquired geometry, it may be extracted automatically [93].

A relief feature is then meant to be a *region* on the base surface where the relief is prominent (such as a convexity, an inflection or a concavity); multiple features may occur and be nested at multiple scales. Conveying surface relief through illustrative shading then amounts to assigning a proper shading value to each surface point, in a way that depends on the multi-scale features it belongs to. The key idea of our feature analysis procedure is to first extract 1D neighborhoods in multiple directions around each point \mathbf{b} of the base surface B , and then analyze their relief at multiple scales. As shown in the left of Figure 3.8, we intersect B in a few directions to obtain base polygonal parametric curves \mathbf{b}_i parametrized in term of arc-length. The corresponding relief functions are given by $h \circ \mathbf{b}_i(t)$ (in red), where t denotes the arc-length.

In order to identify neighboring surface features, we approximate $h \circ \mathbf{b}_i$ by a set of cubic polynomials fitted for multiple support sizes s_j . Higher values of s_j correspond to coarser scales. Cubic-polynomials are particularly well suited for feature analysis since they exactly exhibit a pair of convex and concave regions separated by an inflection point. Each polynomial (in green in Figure 3.8) is given by $h_{i,j}(t) = u_0 + u_1 t + u_2 t^2 + u_3 t^3$ and computed by a standard least square minimization applied to a finite set of samples t_k uniformly spread in the range

⁵This project has been developed in the context of the SeARCH project by Lucas Ammann during his PostDoc.

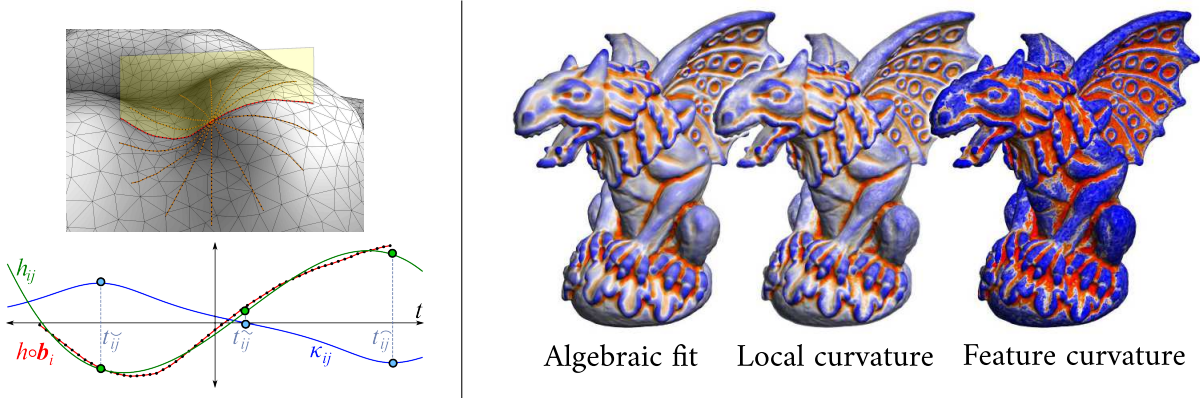


Figure 3.8: Top left: the curves \mathbf{b}_i on the detailed surface S are extracted in multiple directions by intersecting planes defined relative to the base surface B . Bottom left: the height along a curve $h \circ \mathbf{b}_i$ on a support of size s_j is fitted with a cubic polynomial h_{ij} from which we compute the curvature function κ_{ij} and feature locations t_{ij}^{\sim} (convexity), t_{ij}^{\sim} (concavity) and t_{ij}^{\sim} (inflexion). Right: the mean curvature obtained from the least square fit of an algebraic sphere yields similar results than averaging $\kappa_{ij}(0)$ at the current point. In contrast, averaging the curvature $\kappa_{ij}(t_{ij}^*)$ at the closest convexity or concavity location t_{ij}^* exhibits much more segmented curvatures.

$[-s_j/2, s_j/2]$: $h_{ij} = \operatorname{argmin}_P \sum_k (P(t_k) - h \circ \mathbf{b}_i(t_k))^2$. The fitted polynomial h_{ij} provides us with analytic expressions of the first three order derivatives; this not only permits to compute the curvature κ_{ij} but also the parametric locations of the neighboring convexity (t_{ij}^{\sim}), inflexion (t_{ij}^{\sim}) and concavity (t_{ij}^{\sim}), which is the main benefit of our analysis.

Feature combination For an arbitrary choice of surface feature that should be conveyed through shading, we combine the corresponding feature values f_{ij} at all scales and directions:

$$(3.4) \quad F = \sum_j \alpha_j \sum_i \hat{w}_{ij} f_{ij} \quad \text{with} \quad \hat{w}_{ij} = \frac{w_{ij}}{\sum_{i,j} w_{ij}},$$

using $w_{ij} = (1 - \eta_{ij})(1 - d_{ij})$ where η_{ij} measures the average deviation of surface normals from the intersection plane in the i th direction at the j th scale, and d_{ij} measures the distance to a given feature location t_{ij} , normalized by the support radius $s_j/2$. The α_j parameter will be used for user-controlled shape exploration and is set to 1 by default.

The right part of Figure 3.8 illustrates our approach at a single scale by visualizing surface curvature using a simple color code (blue for convexities, red for concavities, white for inflexions). If we choose $f_{ij} := \kappa_{ij}(0)$, it boils down to an estimation of the local mean curvature, which produces results equivalent to algebraic fitting methods [37] for instance. If instead we use $f_{ij} := \kappa_{ij}(t_{ij}^*)$ where t_{ij}^* corresponds to the location of the closest concavity or convexity, it results in a much more segmented analysis, each point displaying the curvature of the neighboring surface feature it belongs to. Figure 3.9 illustrates our approach at multiple scales using Equation 3.4. Here we use another feature function: $f_{ij} := 1 - d_{ij}^{\sim}$, which corresponds to the normalized distance to the closest inflexion. Our solution leads to more legible results compared to blending using uniform weights. Moreover, users may still explore each scale by manually tweaking the α_j parameter.⁶ In the following, we use 8 directions with 3 to 5 scales.

⁶This bears some resemblance to multi-band frequency equalizers in sound editing.



Figure 3.9: Diffuse shading is darkened proportional to the distance to curvature inflections, in (a) with only a coarse scale, in (b) using only a fine scale, and in (c,d) combining all 3 scales together using our adaptive weights or uniform weights respectively. The use of either a single scale or uniform weighting of scales fails to convey all the relief features in a legible way.

Illustrative shading Existing shading techniques may be adapted to take advantage of our surface feature analysis, as demonstrated in Figure 3.12.

We propose a variant of accessibility shading inspired by ambient occlusion. It computes a local visibility value $v_{ij} \in [0, 1]$ for each 1D polynomial h_{ij} as the integration of the cosine term over the range of visible directions. This yields $v_{ij} = 1 - \frac{1}{2} \left(\cos^2(\phi_{ij}^-) + \cos^2(\phi_{ij}^+) \right)$, where the integration bounds ϕ_{ij}^\pm correspond to the minimum and maximum visibility angles around the normal \mathbf{n}_{ij} as shown in the left of Figure 3.12 (top row). The derivation of ϕ_{ij}^\pm is given in closed form in the original paper. The combined accessibility shading value V obtained via Equation 3.4 may now be used to explore the visualization of surface relief details at multiple scales, as shown in the top row of Figure 3.12.

In order to directly visualize surface features, we define three membership functions $m_{ij}^{\widehat{}}$, $m_{ij}^{\widetilde{}}$ and m_{ij}^{\sim} that inform on proximity to each feature inside each 1D fitted polynomial; as shown in the left of Figure 3.12 (middle row), they form a partition of unity. Assuming $t_{ij}^{\widetilde{}} < t_{ij}^{\widehat{}}$, we write:

$$(3.5) \quad m_{ij}^{\widehat{}}(t) = \left[\frac{t_{ij}^{\widetilde{}} - t}{\Delta_{ij}} \right]; \quad m_{ij}^{\widetilde{}}(t) = \left[\frac{t - t_{ij}^{\widetilde{}}}{\Delta_{ij}} \right]; \quad m_{ij}^{\sim}(t) = \left[1 - \frac{|t_{ij}^{\widetilde{}} - t|}{\Delta_{ij}} \right]$$

where $\Delta_{ij} = |t_{ij}^{\widetilde{}} - t_{ij}^{\widehat{}}|/2$, and $[x]$ is a function that clamps x to the $[0, 1]$ range. If instead $t_{ij}^{\widetilde{}} > t_{ij}^{\widehat{}}$, memberships are obtained by exchanging values for $m_{ij}^{\widetilde{}}$ and $m_{ij}^{\widehat{}}$. The combined membership values M obtained by Equation 3.4 may then be used to linearly blend between three different types of reflectances, as shown in the middle row of Figure 3.12 where we have made use of three MatCaps, one for each type of feature.

Radiance scaling [VPB⁺11] also benefits from the multi-scale controls provided by our surface relief analysis technique. Instead of directly using the relief curvature κ_{ij} , we remap it to a curvature function $c_{ij} = \kappa_{ij}(0) \left(1 - \left[\frac{t - t_{ij}^*}{\Delta_{ij}} \right] \right)^\gamma$, where $[\cdot]$ clamps to $[0, 1]$ as before, and $\gamma \in \mathbb{R}^+$ is a user-specified parameter that controls the sharpness of concave and convex regions. As illustrated in the left of Figure 3.12 (bottom row), $c_{ij} = \kappa_{ij}$ at $t_{ij}^{\widetilde{}}$, $t_{ij}^{\widehat{}}$ or t_{ij}^{\sim} , but exhibits a smoother behavior at other locations in a way controlled by γ . This permits to convey surface features as either thin lines ($\gamma < 1$) or broader regions ($\gamma > 1$) as shown in the bottom row of Figure 3.12.

Outlook on shape exploration First and foremost, it is mandatory that the exploration process work in real-time, even though the analysis in preprocess might be more time consuming (the

latter takes from a few seconds up to a minute in our implementation). Second, the choice of α_j in Equation 3.4 depends on the kind of features the user wants to preserve. As mentioned previously, in its simplest form, α_j is manually controlled to globally filter out or boost the j th scale. We have also experimented with a surface-varying feature-based filter: $\alpha_j = \Lambda\left(\text{std}_i(\eta_{ij}) \sum_i \hat{w}_{ij} \kappa_{ij}(t_{ij}^*)\right)$ where Λ is a user controlled sigmoid function that smoothly transitions between the selected and rejected parts. The first factor retains anisotropic relief features defined as the standard deviation of the normal deviation η_{ij} over the set of directions. The second factor selects regions for which the closest relevant feature has a high mean curvature. The effect of this feature filter is demonstrated in Figure 3.10: surface details are hardly seen in (a); using either MatCaps on surface features (b) or a combination of radiance scaling and accessibility shading (c) not only reveals carved lines but also stone grain; using the proposed feature filter (d) permits to get rid of most of the irrelevant details and reveals a horse head seen from the side.





Figure 3.10: The original diffuse shading (a) of the scanned portion of a cave wall hardly reveals the engraved shape (a horse head seen from a side view with smoke escaping from its muzzle). After analysis and assignment of different lit-spheres to the extracted features (b) or a combination of radiance scaling (RS) and accessibility shading (AS) (c), we can provide a more legible rendering. Filtering removes most of the unwanted details (d).

Such filtering functionalities are appealing when it comes to shape exploration; however it is crucial that they do not lead the user (e.g., archaeologist) toward false conclusions by exaggerating accidental surface features. Our approach presents a limitation in this respect, since it depends on a preliminary base/relief decomposition which might erroneously assign the curvature of the base surface (e.g., the surface before it was carved) to the relief function, in which case illustrative shading will convey erroneous (or at least over-exaggerated) relief information. One interesting solution would be to devise semi-automatic base/relief separation methods where an expert identifies regions expected to belong to the base surface only.

It would also be interesting to study object shape beyond surface relief. For instance, existing shape exploration tools propose completion curves [94] to help in the reconstitution of deteriorated artefacts. In general, analyzing larger shape features and conveying their properties through shading remains challenging. To some extent, this is already achieved by accessibility shading techniques when applied to whole object shapes; yet they only characterize the depth of cavities (not protrusions) and offer no direct controls for scientific exploration. Interestingly, one simple technique used by archaeologists to study protrusions and cavities on parietal art for instance consists in rotating a hand-held light source around the object while keeping a grazing angle configuration, which creates more or less elongated shadows depending on surface shape. A particularly exciting direction of future work would be to devise automatic lighting, shading and shadowing techniques that convey at once the shape properties of large object parts without requiring tedious manual trials and errors.



(a) Scaling w/ ambient occlusion



(b) Scaling w/ a prefiltered envmap



(c) No scaling




(d) Scaling on all inter-reflections

Figure 3.11: Radiance Scaling may be made to work with prefiltered environment lighting (a) or ambient occlusion only (b), which permits to enhance surface details in real-time rendering settings. The method is easily adapted to work with global illumination (c,d): both the dragon and its reflections in the vase and on the ground are altered to better convey surface shape.

Accessibility shading



Lit sphere shading



Radiance scaling





Figure 3.12: Top row: estimated normals \mathbf{n}_{ij} are used to compute visibility angles ϕ_{ij}^\pm for accessibility shading. Starting from a standard diffuse shading, we convey coarse scales then all scales using accessibility shading. Middle row: membership functions m_{ij} (in red, green and blue for concavities, inflections and convexities) serve to identify features for direct visualization. A single lit sphere is used in the original rendering. In comparison, using one lit-sphere per surface feature using membership functions (shown in red, green and blue) at a fine scale better conveys surface details, while using all scales provides a more balanced depiction of surface features. Bottom row: the curvature function c_{ij} (in blue) is used instead of κ_{ij} (in orange) as it offers more control thanks to the sharpness parameter γ (different thin blue lines). The shape of a detailed face model is enhanced by applying radiance scaling at all scales, using different values for sharpness. The grey images show the remapped curvature function c_{ij} .

3.2 Material design

We have finally reached the limit scale of the image formation process: the small (often even assumed infinitesimal) region of a surface hit by a single ray emanated from the viewpoint. Properly capturing the radiance reflected or transmitted at a single interaction point would require simulating light transport on the microscopic structure found inside the corresponding (small or infinitesimal) surface region. This is not only extremely



time consuming, but it also provides little insights on the visual effects observable in the final image. Fortunately, since we are only concerned by the radiance eventually leaving the surface, such a light transport simulation does not have to be performed at each and every surface point: it may be stored or modeled in a material function.

Material functions For opaque materials, the function of interest is the Bidirectional Reflectance Distribution Function (BRDF) first introduced by Nicodemus [95]. We have already encountered it multiple times, in particular in the rendering equation 2.6, and wrote it $\rho_{\mathbf{n}}$ with the \mathbf{n} subscript denoting its dependency on the local surface tangent frame. Formally, it is defined as a ratio of reflected radiance over irradiance: $\rho_{\mathbf{n}}(\omega_i, \omega_o) = \frac{dL_r(\omega_o)}{L_i(\omega_i)\cos\theta_i d\omega_i}$, and has units of inverse steradians. Its relation to the reflected radiance equation is made obvious by multiplying $\rho_{\mathbf{n}}$ by its denominator and integrating, which yields Equation 2.6. A BRDF is a 4D function with two dimensions for each of its input unit vectors. In the most general case, it gives out a scalar reflectance for each choice of wavelength and polarization; however, it is common in graphics to overlook polarization⁷ and consider only three representative wavelengths for each of the red, green and blue channels.⁸ A physically-realistic BRDF must ensure three key properties: it must be positive, reciprocal (i.e., invariant to a swapping of ω_i and ω_o) and energy conserving [95].

A similar function may be introduced for transparency in materials, called the BTDF where the 'T' stands for transmission. The full response of a transparent material is called a BSDF, where the 'S' stands for scattering; it is given by two pairs of BRDF+BTDF — one pair when light comes from outside of the material, and another pair when it comes from inside. A BSDF summarizes light transport occurring at microscopic scales; however, it only retains the angular distributions of reflected and transmitted radiance, not the way light might have been transported laterally on the surface (due to inter-reflections) or below the surface (due to diffusion). Even though such lateral scattering occurs in the majority of real-world materials, it has its most vivid visual manifestation in translucent materials, for which a dedicated function has been introduced [96]: the BSSDF, where 'SS' stands for subsurface scattering.⁹

Material models I will here focus on BRDFs, which may be organized in a number of categories depending on their dimensionality and symmetries in 4D space. The simplest BRDFs are constant functions: this is the case of perfect diffuse (i.e., Lambertian) and perfect specular (i.e., mirror) materials; in the former case the function is constant for any (ω_i, ω_o) pair, in the latter case it is

⁷Illumination is most often considered randomly polarized, which simplifies light transport.

⁸This is a valid approximation as long as geometric optics effects prevail and dispersion remains limited.

⁹The BSSRF is essential for the modeling of living tissues such as skin, plants or food, due to their high water content. For instance, it has been heavily used in Pixar's *Ratatouille* to make meals look appetizing [97].




Figure 3.13: Photographs of (a) a hazy and (b) scratched materials (©Elcometer); (c) brushed metal patterns (©Phuznutt).

given by a Dirac in the specular direction. More complex functions are required for describing most real-world materials between these two ideals. One-dimensional BRDFs correspond to radially-symmetric functions, either centered around the normal or specular direction, such as the Phong model [98]. They have been superseded by more general isotropic BRDF models, which have a direct physical meaning: the microscopic surface is then assumed to be statistically invariant to rotation around the surface normal. As a result, the BRDF is itself invariant to the conjoint rotation of ω_i and ω_o around \mathbf{n} and may be represented with three dimensions or less. An example is shown in Figure 3.13a, a particular type of glossy material that we study in Section 3.2.1. When the micro-geometry does not manifest this rotational invariance, the corresponding BRDF is called anisotropic and usually fills the full 4D space. The highlights produced by anisotropic materials exhibit complex elongated shapes on the surface as seen in Figure 3.13(b,c); we study them in details in Section 3.2.2.

A wide variety of BRDF models has been introduced in the Optics and Graphics domains. In Computer Graphics, the majority of models rely on the microfacet theory [99], which assumes that the surface is composed of a distribution of facets at the microscopic scale. Their reflectance properties are usually kept simple to yield analytic BRDF models; in most cases, they are assigned either Lambertian or specular reflectance functions. A *glossy* material is for instance given by:

$$(3.6) \quad \rho_{\mathbf{n}}(\omega_i, \omega_o) = \frac{D(\mathbf{h})G(\omega_i, \omega_o)F(\mathbf{h} \cdot \omega_{i|o})}{4 \cos \theta_i \cos \theta_o} \quad \text{with} \quad \mathbf{h} = \frac{\omega_i + \omega_o}{\|\omega_i + \omega_o\|},$$

where the half-way vector \mathbf{h} is the bisector between ω_i and ω_o , D is a Normal Distribution Function (NDF), G is called the masking-shadowing term, and F is the Fresnel term. The use of the NDF in Equation 3.6 amounts to count the proportion of *specular* microfacets whose normals are exactly aligned with \mathbf{h} , since other microfacets will not contribute to reflected radiance. Only a subset of the remaining microfacets will be visible from either the ω_i or ω_o directions, which is accounted by the masking-shadowing term. This implicitly assumes that only single scattering is considered; even with this restriction, G should be dependent on the choice of distribution D [100] to be physically sound. For isotropic BRDFs, D becomes a function of θ_h since microfacets are then assumed to be oriented uniformly around the surface normal; G is then simplified accordingly. Finally, the specular microfacets are not considered as perfect mirrors, but as interfaces between air and a medium of given refractive index, with the corresponding reflectance governed by Fresnel equations [101] and encapsulated in the F term.

There exists a number of other theories adapted to different micro-structurations of matter, such as the Beckmann-Kirchoff theory [102] for diffraction due to surface irregularities having sizes comparable to visible wavelengths.¹⁰ An exhaustive survey of existing models and acquisition techniques for fitting material model parameters is outside of the scope of this document; the interested reader might consult the survey of Guarnera et al. [103]. Irrespective of the choice of model, a material may be made inhomogeneous simply by varying the model parameters across an object surface. This leads to the so-called surface-varying BRDF (SV-BRDF), which is usually written $\rho_{\mathbf{n}}(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o)$ to let explicitly appear the dependency on surface location \mathbf{x} .

3.2.1 Glossy appearance

In his seminal book on appearance measurement [104], Hunter has identified six *visual* properties of gloss: specular gloss, sheen, luster, haze, distinctness-of-image and surface-uniformity. He suggests that, except for surface-uniformity, all of these visual properties may be connected to reflectance (i.e., BRDF) properties. There exists standard test methods for measuring some of these properties (such as ASTM D523, D430 or D4039), and recent work has studied how they are related to human vision [105]. Nevertheless, the visual effects of a BRDF in arbitrary (as opposed to controlled) illumination environments is not yet properly understood.

This poses a challenge that is illustrated in Figure 3.14, which shows renderings of spheres made of five different materials in two different environment illuminations in orthographic view. Upon reflection off the sphere, the image of the environment is altered in a variety of ways: how should these material-induced visual effects be studied and described? How are they related to Hunter’s six types of gloss? How are they perceived by human observers?




Figure 3.14: Rendering of different BRDFs coming from the MERL database (from left to right: specular-black-phenolic, yellow-phenolic, color-changing-paint2, gold-paint and neoprene-rubber) under two different environment maps (upper row: galileo; lower row: uffizi). Each BRDF has a different effect on the reflected image of the environment.

One way to identify material properties consists in fitting parametric BRDF models to data, and consider parameter values and ranges as relevant properties. Ngan et al. [106] have


¹⁰Microfacet models may be viewed as a special case of diffraction, as shown by Beckmann: when roughness is much greater than visible wavelengths, scattering turns out to depend only on height field slopes and is independent of wavelengths; in other words, a geometric approach becomes valid.

conducted such an empirical study, using as input measured materials coming from the MERL database [107]. This database holds 100 measured isotropic BRDFs and displays a wide diversity of material appearances. The main issue with this approach is that parameters are limited by the choice of models that only span a subset of real-world BRDFs. In contrast, we have chosen to directly study the *filtering properties* of BRDF slices [ZBB⁺15].¹¹

Preliminary study A BRDF slice is a 2D function obtained by fixing the view direction of a BRDF, as commonly done when evaluating Equation 2.6. Four such slices are depicted in the top row of Figure 3.15 for four different viewing elevations. We propose to compute the statistical moments of BRDF slices to identify filtering properties, which raises two complications.

First, BRDF slices may exhibit a combination of multiple components, at least two in common situations: diffuse and specular. However, moments are good descriptors only for *unimodal* functions; computed statistics are not meaningful otherwise. We rely on a simple heuristic method to separate diffuse and specular components in BRDF data: we sample the BRDF at a viewing elevation of 45 degrees, retrieve the minimum reflectance value, then remove this constant from the data in order to obtain the specular component. Some materials still exhibit multi-modality after this process, and we remove them manually, leaving a total of 40 BRDFs.

Second, using a classical parametrization in terms of elevation and azimuth angles forbids the computation of odd-order moments: these are anti-symmetric, and thus cannot be easily reconciled with the periodicity of the azimuthal dimension. We thus convert a BRDF slice using a *view-centered* angular parametrization with poles orthogonal to the view direction ω_o as shown in the inset figure. This has the advantage of minimizing distortions around the incidence plane, where most of the BRDF energy is concentrated. A pair of angles (θ_i, ϕ_i) in the parametrization corresponds to a direction $\omega_i = (\sin \theta_i \cos \phi_i, \sin \phi_i, \cos \theta_i \cos \phi_i)$, assuming that the projection of the view direction on the tangent plane is aligned with $\mathbf{e}_0 = (1, 0, 0)$. The bottom row of Figure 3.15 visualizes four slices in this angular parametrization.



Statistical analysis We may now compute BRDF slice moments, which are statistical quantities that characterize the shape of a distribution. For a scalar distribution f defined over our parametrization, we have:

$$(3.7) \quad \mu_{n,m}[f] := \mathbb{E}_f[\theta_i^n \phi_i^m] = \iint \theta_i^n \phi_i^m f(\theta_i, \phi_i) d\theta_i d\phi_i,$$

where $\mu_{n,m}$ is a scalar coefficient of the moment tensor of order $n + m$. A BRDF slice for a given view direction, which we simply denote by ρ for clarity, is *not* a scalar distribution. To remedy this, we first only consider intensity instead of separate color channels or wavelengths. Then the energy of a BRDF slice is given by a scalar $\alpha = \mu_{0,0}[\rho]$, and we write $\bar{\rho} = \frac{\rho}{\alpha}$ the resulting normalized distribution. Another simplification arises from the quasi symmetry of slices with respect to the incidence plane: all $\mu_{n,m}$ with m odd are close to zero.

The 1st-order moments of $\bar{\rho}$ indicate the mean location of the slice in the angular parametrization, which corresponds to a mean direction on the hemisphere. In our case, the mean along θ is imply $\mu_{1,0}[\bar{\rho}]$. The co-variance matrix of $\bar{\rho}$ characterizes how the BRDF slice is spread around its mean, with larger values implying larger spread. It is defined as the *centralized* moment

¹¹This project has been developed with Carlos Zubiaga during his PhD as part of the PRISM ITN and Laurent Belcour during his PostDoc.




Figure 3.15: Top row: 3D polar visualization of four slices of the isotropic gold-paint BRDF at increasingly grazing viewing elevations θ_o . Bottom row: the same BRDF slices visualized in our view-centered angular parametrization where the light direction is given by (θ_i, ϕ_i) .

matrix Σ of order 2, which is written $\Sigma_{n,m}[\bar{\rho}] = E_{\bar{\rho}}[(\theta_i - \mu_{1,0})^n \phi_i^m]$ for $n + m = 2$. Co-skewness and co-kurtosis characterize the asymmetry and peakedness of BRDF slices. They are defined as the *standardized* moment tensors of order 3 and 4 respectively, obtained by centering $\bar{\rho}$ on its mean and scaling it by respective variances. In our case, we write $\gamma_{n,m}[\bar{\rho}] = E_{\bar{\rho}}\left[\left(\frac{\theta_i - \mu_{1,0}}{\sqrt{\Sigma_{2,0}}}\right)^n \left(\frac{\phi_i}{\sqrt{\Sigma_{2,0}}}\right)^m\right]$. It is common to use the modern definition of kurtosis, also called excess kurtosis, which is equal to 0 for a Normal distribution; this amounts to use $\{\gamma_{4,0} - 3, \gamma_{2,2} - 1, \gamma_{0,4} - 3\}$.

Moment profiles Each statistical property may be plotted as a function of view elevation θ_o , which we call a moment profile. Figure 3.16 visualizes the entire set of moment profiles for the materials used in our study. We have omitted statistics where profiles are close to 0 due to the near-symmetry of BRDF slices across the incidence plane, which are characterized by an odd m subscript. Also note that moment profiles of odd-order (b,d) start at a value of 0 (when $\theta_o = 0$), while even-order profiles start at the same value in different dimensions for a same order. This is simply due to the fact that all BRDF slices are radially-symmetric at normal incidence.

We now make a number of observations for each type of statistical property:

- The energy profiles α exhibit a constant behavior up to grazing angles, at which point they tend to increase progressively. Such an increase might be due to Fresnel effects or asperity scattering [108], and manifests visually as an increase in reflected intensity at grazing angles (even though it is attenuated by the cosine term in Equation 2.6).¹²
- For the mean $\mu_{1,0}$, one would expect profiles to tend toward lines, with a slope of -1 in the case of a mirror material where all the energy is reflected in the specular direction. Even though there is a global linear trend for small viewing elevations, the slopes frequently depart from -1 , and some profiles even depart from linearity toward grazing angles. Such departures might be due to the presence of different material layers for instance.

¹²Some materials exhibit energy profiles that drop down at extremely grazing angles, but it is likely that this behavior might be due to measurement issues and/or post-processing of data in the MERL database.




Figure 3.16: Moment profiles computed from our selected BRDFs are shown at increasing orders: (a) energy, (b) mean, (c) co-variance, (d) co-skewness and (e) co-kurtosis.

- The variances exhibit relatively flat profiles at low-grazing angles overall. However, $\Sigma_{2,0}$ tends to increase toward grazing angles, whereas $\Sigma_{0,2}$ rather decreases. This suggests that the shape of the BRDF slice tends to become roughly elliptical toward grazing angles, which is in accordance with micro-facet theory. Visually, this means that the material will blur differently the environment at different viewing angles. A few materials even show non-monotonic variance profiles, suggesting that blur strength will smoothly oscillate when going from normal to grazing viewing incidences.
- Both co-skewness and co-kurtosis profiles increase strongly toward grazing angles, even tremendously in the case of excess kurtosis. It is highly likely that these effects are due to the clipping of BRDF slices by hemispherical boundaries, whose effect becomes more and more pronounced at higher orders. It still suggests that the shape of BRDF slices becomes more complex than elliptical at grazing angles, and that this complexity varies among materials. Moreover, some materials yield kurtotic BRDF slices at normal incidence.

Our preliminary study has raised a number of questions both in terms of optics (i.e., how to explain statistical properties of materials in terms of light transport at the micro-scale) and perception (i.e., what is the effect of material properties on visual appearance). One particular material property has been the subject of recent investigations in the graphics literature: the soft falloff of glossy highlights. It may be due to a particular distribution of microfacets [109], to multi-layered materials [110], or to diffraction effects [111]. Visually, it may be related to haze in

Hunter’s classification. Intuitively, hazy reflections are those in which a relatively sharp (distinct) reflection is superimposed with, or surrounded by, a blurry bloom or fringe, similar to the effect of viewing a light source through haze or mist (see Figure 3.13a).

To date, the visual perception of such hazy gloss materials has barely been investigated. Among the six appearance aspects of gloss suggested by Hunter [104], only two independent dimensions were found in perceptual studies on a single class of materials such as real paint samples [112] or rendered materials [105, 113]. A natural hypothesis is that haze is related to the kurtosis of a BRDF slice at normal incidence, since it characterizes the falloff of the material response. We set out to test this hypothesis using a perceptual rating experiment [VBF17].¹³

Rating experiment The computer-generated stimuli were random blob shapes rotating back and forth, made of a silvery metallic material and illuminated by the Galileo light probe. The specular material was made of two Ward components [114], yielding a BRDF of the form:

$$(3.8) \quad \rho_{\mathbf{n}} = F(\mathbf{h} \cdot \boldsymbol{\omega}_{i|o}) \cdot (\rho_{s,n} + \rho_{s,w}), \quad \text{with} \quad \rho_{s,n|w} = \frac{k_{n|w}}{4\pi\alpha_{n|w}^2 \sqrt{\cos\theta_i \cos\theta_o}} \exp\left(-\frac{\tan^2\theta_h}{\alpha_{n|w}^2}\right),$$

where $\rho_{s,n}$ and $\rho_{s,w}$ are narrow and wide specular components respectively, and F is the Fresnel term (we set $\eta = 0.145$ and $\kappa = 3.19$, which is typical of silver). This sum of two specular components is the simplest approximation to the true BRDF of various types of hazy materials.




Figure 3.17: Left: space of two-component glossy materials used in the rating experiment. Example materials of the four corners of the space and 2D polar plots of their BRDF are displayed at their correct location in the space. Right: representative screenshots of the stimulus presentation for the rating task. Subjects had to rate materials according to six different qualities.

The sharpness of the narrow component was set to $\alpha_n = 0.0396$ and the sharpness of the wide component was chosen from the range $\alpha_w \in \{0.0833, 0.1125, 0.1417, 0.1708, 0.2\}$. In other words, the difference in sharpness $\Delta\alpha = |\alpha_n - \alpha_w|/2$ was chosen from the range $\Delta\alpha \in \{0.0218, 0.0364, 0.0510, 0.0656, 0.0802\}$. The reflectances of the two components summed to the total $k_{\text{total}} = 0.95$, typical for silver, and the ratio of the reflectances of the two components was chosen from the range $\frac{k_{n|w}}{k_{\text{total}}} \in \{1/6, 2/6, 3/6, 4/6, 5/6\}$. The space of stimuli is visualized in Figure 3.17-left. In each trial a single stimulus was presented and each of the 14 participants was

¹³This project was initiated with Carlos Zubiaga at the end of his PhD, and pursued after the end of the PRISM ITN in collaboration with Peter Vangorp and Roland Fleming.




Figure 3.18: Left: mean ratings across participants (a) for each of the six properties. Grey values indicate mean rating, consistently normalized such that black represents the lowest and white the highest average ratings across all stimuli and properties. Fits of a third-order polynomial surface with superimposed contours (b) facilitate visualization of the main trends in each plot. Right: directions of the six rating scales represented by unit vectors in the principal component space, projected onto the plane formed by the first two principal components which together account for 93.6% of the variance. Nearly orthogonal rating scale directions, such as blurry and hazy, are almost completely uncorrelated, which indicates distinct perceptual dimensions.

instructed to *rate the presented material on the following six different continuous scales related to gloss appearance: glossy vs. matte, sharp vs. blurry, not hazy vs. hazy, polished vs. unpolished, low vs. high friction, not coated/varnished vs. coated/varnished*. Each participant provided subjective ratings of the six different surface qualities by adjusting the slider positions as shown in the right of Figure 3.17, performing this task for 25 materials \times 2 repetitions.

Figure 3.18-left shows the mean results of the rating task across participants. There are several notable aspects of the results. We find that for this range of stimuli, matteness, blurriness and appearing unpolished are associated with materials with a prominent wide component more-or-less irrespective of the difference in sharpness, which is illustrated by the fact that the highest rating values appear all across the bottom edge of the stimulus space. In contrast, haze and appearing coated or varnished are also associated with a prominent wide component, but only when there is a large difference in sharpness, which is illustrated by the fact that the highest rating values appear only on the right side of the bottom edge. All the materials were essentially perceived as having low friction. The difference between blurry-matte-unpolished appearances and hazy-coated/varnished appearances is made clearer through a principal component analysis on the mean responses. Figure 3.18-right plots the factor loadings of each property in the space spanned by the first two principal components, which together account for 93.6% of the variance in the mean data. The plot shows how blur, matteness and appearing unpolished are correlated with one another, while haziness and coatedness are approximately orthogonal to this cluster. This suggests that haze truly is a distinct perceptual dimension of reflectance, independent of blur, and associated with the presence of both sharp and blurry reflections simultaneously.

Candidate interpretations Hunter [104] proposed a number of haze measurement techniques, which are widely used in industrial applications even though they are not intended as





Figure 3.19: Candidate interpretations for the mean haziness ratings (A); each plot has been normalized independently as shown by their respective color bars. Hunter’s haze measurements (B) at $\theta_{\text{off}} = 5^\circ$ do not appear related to perceived haze. Only when pushed to $\theta_{\text{off}} = 20^\circ$ (C) do they begin to show a similar distribution. Excess kurtosis (D) is necessary to elicit hazy gloss, but is not related to perceived haze either. The energy k_h of the halo component (E) from our decomposition is distributed similarly to haze ratings. The sharpness α_c of the central component (F) is distributed along a different direction and is similar to (B).

models of human visual processing. The one retained by the ASTM [115] is a ratio of reflectances H/S , with S measured in the specular direction and H measured in an off-specular direction, for an incident light elevation of $\theta_i = 30^\circ$. The choice of the θ_{off} angle between specular and off-specular directions vary depending on the materials considered; standard values are 2° and 5° , and give essentially the same results when plotted in the stimulus space (Figure 3.19b); they are not aligned with the haze ratings of Figure 3.19a. When θ_{off} is set to 20° (Figure 3.19c), results become correlated to ratings, but this manual adjustment is not general enough.

Our initial hypothesis was that kurtosis of the material response might be related to perceived haziness. The excess kurtosis of our two-component BRDF—a mixture of Gaussians—is positive and can be computed analytically. However, as shown in Figure 3.19d, we find that kurtosis is not aligned with participants’ haziness ratings, hence contradicting our hypothesis. In other words, even though a material must exhibit kurtosis to elicit hazy gloss, kurtosis does not seem to characterize our subjective experience of haziness, at least for this range of materials.

We propose that the visual system may represent haze by decomposing the material response into two distinct components: a central BRDF component that matches the reflectance peak and spread around the specular direction, and a surrounding halo component that corresponds to the (positive) residual reflectance (if there is any). The inset figure illustrates the procedure we use to perform this decomposition: we scale the $\rho_{s,n}$ component to match the peak reflectance (dashed curve) and make it as wide as possible (in blue) to “fit inside” the compound specular term (red outline); the halo component (in orange) is then simply obtained by subtraction.



We write k_c and α_c for the reflectance and sharpness of the central component, and use $k_h = 1 - \frac{k_c}{k_n + k_w}$ to characterize the halo energy. The latter is equal to 0 when $\alpha_c = \alpha_n = \alpha_w$ as the central component encompasses the whole compound specular term. We plot the halo energy in the rectangular stimulus space in Figure 3.19e. It shows a clear alignment with mean haziness ratings. We also plot the sharpness α_c of the central component in Figure 3.19f.

Outlook on perceptual decomposition Figure 3.20 shows the effect of manipulating the intensity of the halo term while leaving the central peak of our decomposition unchanged. The material appears less hazy or more hazy when the halo intensity is decreased or increased respectively, suggesting that the decomposition could be useful for material editing. Moreover,




Figure 3.20: The halo intensity of the original material (dashed curve) may be increased (right) or decreased (left), yielding a corresponding increase or decrease in perceived haziness.

since the central peak and halo are additive, they may be computed separately during rendering, and the resulting peak and halo images added together in post-process. As a result, the editing of Figure 3.20 may be equivalently obtained by manipulating the intensity of the halo image before adding it to the peak image. This raises the question of which image cues are involved in the perception of hazy gloss, and how they could be manipulated to alter the perception of haziness.

3.2.2 Anisotropic reflections


Up to this point, we have only considered the appearance due to homogeneous and isotropic BRDFs. *Homogeneity* implies that surface irregularities responsible for material appearance are small enough not to be seen even in extreme close-up views and that their distribution is statistically the same everywhere on the object surface. This restriction is seldom satisfied by real-world materials: they often exhibit several types of spatial variations, in either reflectance or geometric details. *Isotropy* means that the statistical distribution of these microscopic surface details has no preferred tangential direction; as a result, the highlight produced by an isotropic light source is solely shaped by surface curvatures. However, highlights may also come in extended, arched or looped shapes, which are due to microscopic properties of the material presenting directional elongations across the surface. Examples abound in nature (e.g., petals, leaves or hairs) and in man-made objects (e.g., brushed metals, silk or satin).

We now turn to the study of materials that exhibit both spatial variations and anisotropic structure. One difficulty of such materials is that they require the full four angular dimensions of a BRDF, plus two additional dimensions to encode spatial variations. In particular, the BRDF now needs to be oriented in the surface tangent frame and these orientations will most likely vary across the surface. This is the case for instance for brushed metal surfaces as shown in Figure 3.13c, where orientations corresponds to brushing directions. Moreover, even though elongated micro-structures such as metal scratches are small, they may still be resolved in reasonably close views of the material as seen in Figure 3.13b. Reproducing the appearance of individual scratches and of their distribution at multiple scales is a thus challenging problem.


In the following, we will call anisotropic reflections or highlights the visual patterns due to anisotropic BRDFs. The complex shape of anisotropic highlights is determined by BRDF orientations across the surface; but it also depends on object shape, lighting and viewing configurations.

It is thus extremely difficult to predict the appearance of anisotropic highlights on arbitrary-shaped objects. Our solution lies in an inverse design approach [RGB⁺14] whereby an artist deforms highlights as desired; then BRDF orientations matching this goal are automatically inferred, exhibiting robustness to changes in light or view direction, and even materials.¹⁴

Let's first assume a distribution of mirror microfacets strongly elongated in a tangential direction given by a vector field \mathbf{u} across the surface. Since most microfacet normals are orthogonal to \mathbf{u} , then locally rotating \mathbf{u} to be perpendicular to \mathbf{h} everywhere maximizes reflected radiance. This simple manipulation is shown in the inset figure where the object is lit by a single directional key light source. However, even a slight rotation of the light direction turns the highlight back into a curve. In other words, highlight regions correspond to degenerate configurations; only highlight curves are generic.



Differential analysis Lu et al. [116] have conducted a differential analysis to examine the behavior of a single highlight curve on a quadratic surface. As shown in Figure 3.21, they consider a densely packed distribution of specular 'fibers' that may exhibit curvature and/or splay. The highlight curve and its tangent direction vary significantly with surface curvature, fibers trajectories and halfway vector elevation. They also study highlight curves for simple closed geometric primitives and degenerate cases, but do not for arbitrary-shaped objects.



(a) Fiber distrib. (b) Straight fibers (c) Adding curvature (d) Adding splay (e) Rotating \mathbf{h}

Figure 3.21: Assuming a distribution of *densely packed* specular fibers (a) in the tangent plane, we observe the behavior of the corresponding specular highlight curve (b-e). Straight fibers (b) applied to a curved 3D surface produce a highlight curve (in red) with its tangent direction (in cyan) nearly orthogonal to fibers. When fibers bend (c) or diverge (d), the highlight curve rotates significantly and its tangent nearly aligns with fibers. Decreasing the elevation of the half-way vector \mathbf{h} rotates even further the specular curve in the tangent plane.

We have extended their method to an arbitrary surface described by a manifold M , with BRDF orientations controlled by a direction field $\mathbf{u} : M \rightarrow TM$ where TM denotes the tangent bundle of M . Four differential properties are relevant to our analysis: the curvature $\kappa(\mathbf{u})$ and torsion $\tau(\mathbf{u})$ of the surface in the direction \mathbf{u} ; as well as the signed curl magnitude $|\nabla \times \mathbf{u}| := (\nabla \times \mathbf{u}) \cdot \mathbf{n}$ and divergence $\nabla \cdot \mathbf{u}$. As detailed in the original paper, the highlight tangent field $\ell : TM \rightarrow TM$ is determined up to a sign ambiguity since a highlight curve has no intrinsic parametrization:

$$(3.9) \quad \ell(\mathbf{u}) = \pm \mathbf{U} \begin{bmatrix} 0 & -\nabla \cdot \mathbf{u} & -\tau(\mathbf{u}) \\ 0 & |\nabla \times \mathbf{u}| & \kappa(\mathbf{u}) \\ 0 & 0 & 0 \end{bmatrix} \mathbf{U}^T \mathbf{h},$$

¹⁴This project has been developed with Boris Raymond during his PhD.

where $\mathbf{U} = [\mathbf{u}\mathbf{v}\mathbf{n}]$ is the local-to-world transform (see inset figure). The last row of the central matrix is null since ℓ is a tangent vector. Note however that it is *not* a unit vector; in particular, $\ell = 0$ occurs in degenerate cases. The first column is null as well, which amounts to project \mathbf{h} onto the plane orthogonal to \mathbf{u} . This is equivalent (up to a multiplicative factor) to the highlight tangent that would be obtained for a nearby lighting direction. As shown in Figure 3.22(a-c), ℓ indeed remains accurate for small changes of lighting or viewing direction: it thus characterizes *potential* highlight curves around the current one. A critical configuration occurs when $\mathbf{h} = \mathbf{n}$, which corresponds to the location of isotropic highlights. In this case, ℓ only becomes a function of directional surface curvatures and torsions; hence varying \mathbf{u} has no effect at critical points as shown in Figure 3.22(d,e).





Figure 3.22: Left: the highlight tangent field ℓ of a white key light source is visualized using LIC [38], and compared to the highlight curves of 3 light sources: (a) the red light source is close to the key light, showing an accurate match between ℓ and the red highlight curve; (b,c) when the light source gets farther away from the key light, the accuracy of ℓ begins to decrease. Right: For a given key light, the specular curve always goes through critical points (in red) where $\mathbf{h} \cdot \mathbf{n} \approx 1$, irrespective of the choice of \mathbf{u} (shown using LIC in inset images at top right).

Inverse design If we re-express the highlight tangent field ℓ as a function of the angle φ between \mathbf{u} and \mathbf{t} , then it may be re-written as $\ell(\varphi) = \pm [\mathbf{b} - \mathbf{t}] [\nabla\varphi \bar{\mathbf{v}}^T \mathbf{W}\bar{\mathbf{u}}] [\mathbf{t} \mathbf{b} \mathbf{n}]^T \mathbf{h}$, where $\bar{\mathbf{u}} = [\cos\varphi \sin\varphi]^T$ and $\bar{\mathbf{v}} = [-\sin\varphi \cos\varphi]^T$ are the 2D versions of \mathbf{u} and \mathbf{v} , \mathbf{W} is the 2×2 Weingarten map, and $\nabla\varphi$ is the 2D gradient vector of covariant derivatives. This formulation is better adapted to the inverse design problem, where the goal is to find the scalar field φ^* so that the resulting highlight tangent field comes as close as possible to a user-prescribed field ℓ^* . Formally, the optimization problem is stated as:

$$(3.10) \quad \varphi^* = \operatorname{argmin}_{\varphi} E(\varphi) \quad \text{with} \quad E(\varphi) = \int_M \|\mathbf{L}(\varphi) - \mathbf{L}^*\|^2 dA,$$

where we represent non-oriented highlight tangents by normalized tensors of the form $\mathbf{L} = \frac{\ell\ell^T}{\|\ell\|^2}$. Since the highlight tangent direction is orthogonal to the gradient produced by the highlight, \mathbf{L} is similar to the structure tensors we investigated in Chapter 2, the only difference being that we consider a specific configuration (i.e., anisotropic BRDF, single light source).

The target highlight tangent field ℓ^* is obtained by warping the initial ℓ as shown in Figure 3.23a. Formally, the target highlight tangent at a point \mathbf{p} is $\ell_{\mathbf{p}} = \mathbf{J}_W(\mathbf{q}) \ell_{\mathbf{q}}$, where W is a warping function over a finite support, $\mathbf{q} = W^{-1}(\mathbf{p})$ and \mathbf{J}_W is the Jacobian of W . The effect of the warping operator is seen when comparing Figures 3.23b and c.

As detailed in the paper, the energy $E(\varphi)$ in Equation 3.10 may be discretized over the triangles of a surface mesh. The minimization turns into a non-linear least-squares optimization problem expressed in matrix form that we solve using the Levenberg-Marquardt method. Such




Figure 3.23: Left: the target highlight tangent field ℓ^* is obtained by (a) warping the initial field ℓ inside a finite support: ℓ is shown in (b), ℓ^* in (c). Right: Starting with an initial direction field and a target highlight tangent field (shown with red curves), we apply a few iterations of our solver which here converges in 10 steps.

a method requires the Jacobian of the discretized version of $E(\varphi)$, a very sparse matrix having the same structure as the Laplacian matrix, which we use to speed up computation. As with all non-convex energies, the method requires a proper initial solution; we naturally pick the angles φ associated with the initial tangent direction field. This is shown in the right of Figure 3.23, which depicts a few iterations of our solver. It usually converges in less than 10 iterations.




Figure 3.24: Rendering an object using an anisotropic BRDF [79] produces elongated highlights (a). Our approach lets users directly modify highlight shape (b) by inferring appropriate BRDF orientations. The edited highlights retain their deformed shape for nearby lighting (c) or viewing (d) configurations. Highlight shape is still preserved when a rougher material is used (e).

Micro-sculpt results As seen in Figure 3.24, our warping strategy permits to design complex anisotropic material appearances that are retained in different lighting and viewing conditions, as well as when changing the anisotropic BRDF (only retaining its orientation). In comparison, directly manipulating BRDF orientations is impractical. A limitation of our approach is that critical points cannot be edited since they depend on the light-view configuration, not on BRDF orientations. Highlight editing is thus significantly restricted when dealing with complex surfaces where several critical points occur. Materials with slanted micro-facet distributions like velvet [117] are not well accounted by our current model since they are made of non-tangential fibers. A limitation that we share with Lu et al. [116] is that lighting and viewing must be distant.

Anisotropic materials are not only characterized by their tangential orientations, but also by their complex reflectance response. They exhibit several angular peaks and spatial variations

that produce visually salient spatial variations. This is the case of materials exhibiting many tiny scratches, as found mostly on metals, but also plastics or finished woods. The scratches may be due to polished or deteriorated objects, and are usually distributed all over the surface. They produce a number of interesting visual effects that are visible even from afar: glint lines around highlights, patterned highlight silhouettes, and directionally-smearred reflections (see Figure 3.13b,c). In the following I will present a multi-scale SV-BRDF model [RGB16] tailored to the reproduction of such scratched materials.¹⁵

Scratch primitives As recently demonstrated by Yan et al. [118], homogeneous BRDF models are not adapted to reproduce the appearance of surfaces exhibiting high-frequency irregularities. In the case of scratches, classic microfacet-based BRDF models relying on smooth NDFs are not adapted either: each scratch should be considered as an individual primitive that may be resolved in close-up views. At the same time, for distant viewing conditions, individual scratches will not be visible anymore, while appearance will be dominated by collections of near-parallel scratches yielding anisotropic reflections.

The key insight of our work is to exploit the *regular micro-structure* of scratched materials. Individual scratches are mostly one-dimensional, and many of them are distributed parallel to each other. In the following, we call a scratch *layer* a set of similar parallel scratches; a scratch *distribution* becomes a combination of one or more layers. This is not to be confused with the concept of layered materials (e.g., [119]): here we are only concerned with interface on top.

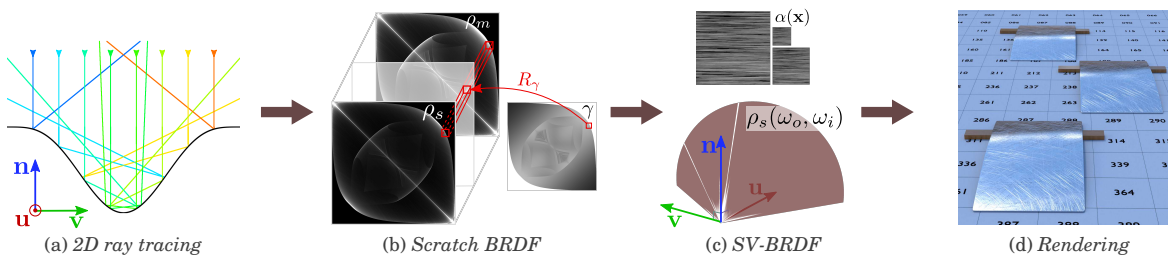


Figure 3.25: (a) The BRDF for a single mirror scratch is simulated using 2D ray tracing (each ray is given a consistent color across reflections for illustration purpose). (b) The full 3D scratch BRDF with Fresnel effects is approximated using an adequate interpolation between a pair of 2D BRDF slices, controlled by an exponent map γ . (c) At render time, a SV-BRDF is constructed by linear combinations of scratch BRDFs (bottom) under a pixel or ray footprint based on the relative area of scratches (top). (d) Our SV-BRDF model is efficiently evaluated at multiple scales, exhibiting both high spatial and angular frequencies.

Since scratches of a same layer are carved inside the surface, their BRDF may be computed *independently* of other scratches, then combined based on their spatial distribution to yield a SV-BRDF. The process is illustrated in Figure 3.25 and detailed in the following pages. The BRDF for a single scratch is simulated using a 2D ray-tracing algorithm, including inter-reflections and realistic Fresnel effects. The result is stored in the three channels of a 2D texture thanks to a parametrization inspired by the mirror cone of Kajiyama [120]. The SV-BRDF is obtained as a linear combination of scratch BRDFs driven by the area of scratches under a pixel footprint.

¹⁵This project has been developed in the context of the RichShape project with Boris Raymond during his PhD.

SV-BRDF model We start by building a BRDF model for a single *mirror* scratch extruded from an arbitrary profile. An important property of such an ideal scratch is that incoming directions contributing to reflection lie on a half-cone around the scratch direction \mathbf{u} . This property was first observed by Kajiya [120] for specular fibers. In our case, this leads to the following BRDF ρ_m , where m stands for “mirror”:

$$(3.11) \quad \rho_m(\boldsymbol{\omega}_o, \boldsymbol{\omega}_i) = \delta(\theta_i - \theta_o) \rho_m(\phi_o, \phi_i) \quad \text{with} \quad \begin{cases} \theta_{i|o} &= \sin^{-1}(\boldsymbol{\omega}_{i|o} \cdot \mathbf{u}) \\ \phi_{i|o} &= \text{atan2}(\boldsymbol{\omega}_{i|o} \cdot \mathbf{v}, \boldsymbol{\omega}_{i|o} \cdot \mathbf{n}), \end{cases}$$

where $(\theta_{i|o}, \phi_{i|o})$ are the coordinates of $\boldsymbol{\omega}_{i|o}$ in Kajiya’s angular parametrization, δ is the Kronecker function, and $\rho_m(\phi_o, \phi_i)$ is a 2D scratch BRDF. In other words, we have reduced the dimensionality of the material function from 4D to 2D in the ideal mirror case.

As illustrated in Figure 3.25a, the 2D scratch BRDF is obtained by a 2D light transport simulation. The BRDFs obtained for analytic and measured scratch profiles are shown in the left of Figure 3.26. The axis of the 2D BRDF space correspond to ϕ_o and ϕ_i and may be interchanged due to reciprocity. The symmetry across the anti-diagonal is due to the symmetry of the profile itself, and thus only occurs in the analytic case. Colors correspond to inter-reflection orders.

Since no real-world material comes anywhere close to a perfect mirror, Fresnel effects must be re-introduced for physical realism. This is not as direct as in Equation 3.6, where a single Fresnel term is sufficient since only single scattering is considered. However, in the limit of grazing angles (i.e., $\theta_{i|o} = \frac{\pi}{2}$), the BRDF becomes identical to ρ_m since Fresnel reflectance is equal to 1. As detailed in the original paper, we found empirically that the scratch BRDF ρ_s with Fresnel effects may be approximated by a linear combination of a pair of 2D BRDFs:

$$(3.12) \quad \rho_s(\boldsymbol{\omega}_o, \boldsymbol{\omega}_i) = \text{lerp}(\rho_0(\boldsymbol{\omega}_o, \boldsymbol{\omega}_i), \rho_m(\boldsymbol{\omega}_o, \boldsymbol{\omega}_i), R_\gamma(\theta_{i|o})) \quad \text{with} \quad R_\gamma(\theta) = \left(\frac{2|\theta|}{\pi}\right)^\gamma,$$

where ρ_0 denotes the 2D BRDF at normal incidence (i.e., $\theta_{i|o} = 0$) that we also compute using a 2D light transport simulation, this time using Fresnel reflectance at each bounce; and R_γ is used to interpolate between the two extreme incidence angles. The exponent γ must be estimated for each (ϕ_o, ϕ_i) pair. As a result, the full scratch BRDF consists of a triplet of 2D functions: $\rho_0(\phi_o, \phi_i)$, $\rho_m(\phi_o, \phi_i)$ and $\gamma(\phi_o, \phi_i)$, which we store in the 3 channels of a 2D texture (see Figure 3.25b).

The BRDF of Equation 3.12 may already be used to model uniform materials made of densely packed parallel scratches (with no space in-between), which will exhibit directionally-smearred reflections. However, variations in the spacing, width, height and orientation of scratches are necessary to reproduce glint lines and patterned highlights. In our approach, scratches in a same layer are localized by an indicator function $\mathbf{1}_s : \mathbb{R}^2 \rightarrow \{0, 1\}$. Different solutions could be taken to design $\mathbf{1}_s$; ours consists in stretching, rotating and binarizing a 2D cellular noise, as illustrated in the right of Figure 3.26. The *relative area* α_s of scratches under a given pixel or ray footprint P is obtained by integration: $\alpha_s(\mathbf{x}) = \int_{P(\mathbf{x})} \mathbf{1}_s(\mathbf{y}) d\mathbf{y}$. The different scratch layers are then linearly combined to yield $\bar{\rho}_\mathbf{n}(\mathbf{x}, \boldsymbol{\omega}_o, \boldsymbol{\omega}_i) = \sum_s \alpha_s(\mathbf{x}) \rho_s(\boldsymbol{\omega}_o, \boldsymbol{\omega}_i)$ where the s subscript is used to indicate the layer number. This simple linear combination is not guaranteed to conserve energy though, since at locations where scratches overlap, the combined relative area $\bar{\alpha}(\mathbf{x}) = \sum \alpha_s(\mathbf{x})$ might be greater than 1. In order to correct this issue, we define the final SV-BRDF by:

$$(3.13) \quad \rho_\mathbf{n}(\mathbf{x}, \boldsymbol{\omega}_o, \boldsymbol{\omega}_i) = \begin{cases} \bar{\rho}(\mathbf{x}, \boldsymbol{\omega}_o, \boldsymbol{\omega}_i) / \bar{\alpha}(\mathbf{x}) & \text{if } \bar{\alpha}(\mathbf{x}) > 1 \\ \text{lerp}(\rho_b(\mathbf{x}, \boldsymbol{\omega}_o, \boldsymbol{\omega}_i), \bar{\rho}(\mathbf{x}, \boldsymbol{\omega}_o, \boldsymbol{\omega}_i), \bar{\alpha}(\mathbf{x})) & \text{otherwise,} \end{cases}$$

which guarantees energy conservation in places where many scratches accumulate, while letting an arbitrary base material ρ_b appear where the distribution of scratches is less dense.




Figure 3.26: Left: a pair of analytic and measured scratch profiles (shown on top) and their corresponding mirror scratch BRDFs (1st, 2nd and higher-order bounces in red, green and blue respectively). Even though the two profiles have a similar global shape, their BRDFs differ significantly: the peak on the anti-diagonal in the BRDF (i.e., mirror reflection) is absent for the measured profile due to its lack of flatness. Right: A 2D cellular noise is stretched and rotated, then binarized to obtain the indicator function used for scratch localization.

Equation 3.13 avoids the computation of the complex BRDF occurring at scratch intersections, replacing it with a simpler average of scratch reflectances. As shown in Figure 3.25d, our model is sufficient to reproduce convincing glint lines and patterned highlights at multiple scales.

Material appearance Our SV-BRDF model is naturally adapted to importance sampling but requires special considerations to work with light sampling, as detailed in the original paper. It is sufficiently efficient to provide interactive feedback when implemented in a progressive renderer running on the GPU. This is required for the design of scratch profiles and distributions. Once this is done, the material may be rendered in a more advanced global illumination engine. Figure 3.27 shows several such renderings on a same, slightly curved metal plate, varying the profile and distribution of scratches. We also compare our approach to a NDF-based approach similar to that of Yan et al. [118], which does not consider inter-reflections inside scratches and approximates the masking-shadowing term. When the scratches are deep enough, the NDF-based approach yields a dimmer appearance and reflections that look less complex.

Our method does not only apply to metals, but also to other types of scratched materials, as demonstrated in Figure 3.28. Inset images reveal the behavior of our SV-BRDF model in close-up views, where different reflection patterns emerge compared to zoomed out views.

Advantages and drawbacks Our SV-BRDF model is tailored to scratched materials, which brings advantages compared to more general approaches [118, 119, 121, 122]. It produces complex highlights due to inter-reflections inside scratches and accurate masking-shadowing thanks to light transport simulation. The simplicity of Equations 3.12 and 3.13 not only enable material design at interactive rates but also eases multi-scale rendering, since it decouples spatial and angular dimensions. In spite of its simplicity, the model adequately reproduces the glint lines, patterned highlights and directionally-smearred reflections characteristic of scratched materials.

The main drawback of our approach is that the micro-surface is restricted to a particular type of structure: we consider idealized and repeated scratches, devoid of imperfections along their extrusion direction. Imperfections might be added by incorporating punctual glitter effects (e.g. [119]) inside scratches. Another way to increase visual richness would be to let all scratch parameters (profile, width, height, etc) vary per layer. This would require to devise an analytic

scratch BRDF model that could be instanced on the fly, as opposed to running a simulation per scratch profile in preprocess. The choice of distribution — statistical or primitive-based — for describing the microscopic structure at an object surface depends on the density and layout of irregularities. I believe the two approaches should be combined to achieve highly realistic results.




Figure 3.27: The choice of scratch profile has a strong impact on material appearance, in either brushed (top row) or scratched (bottom row) patterns. In the first two columns, we use nearly flat profiles, which exhibit multiple highlights due to off-specular peaks in the scratch BRDF. The extreme zooms show the limitation of our method: individual scratches span a few pixels. Steeper scratch profiles yield more contrasted reflections due to multiple bounces being attenuated by Fresnel effects. The material appears dimmer with NDFs since many scratches reflect little of the incoming radiance when inter-reflections are discarded.




Figure 3.28: Materials rendered in global illumination; side images show separate contributions of base (top) and scratch (bottom) materials. (a) Gold (4 layers) is obtained by a regular distribution of sparse tiny scratches combined with a base mirror-like material. Layers are oriented around a main direction on the surface, which produces an elongated highlight blended with a mirror reflection. Such highlights are also visible in macro-scale inter-reflections, taking an orange tint. (b) Painted metal (16 layers) makes use of a sparse distribution of scratches in multiple orientations, and a paint BRDF for the base material. A sparse set of multi-directional layers reproduces the random deteriorations expected with such materials. (c) Plastic (32 layers) uses Fresnel formula for dielectrics and adds a colored diffuse term to the dense multi-directional scratch layers. (d) Finished wood (16 layers) is similar to Plastic, except that it uses a sparser set of scratches and a textured diffuse term.

3.3 Discussion

Apparent shape Even though the light warping and radiance scaling techniques of Section 3.1.1 have been presented as alterations of the incoming radiance, they might equally well be conceived as local modifications of the BRDF slices that we have studied in Section 3.2.1. Indeed, since radiance scaling works by modifying light intensities, modifications of the (view-dependent) material energy could lead to visually equivalent results. Similarly, light warping might be reproduced by modifying the (view-dependent) material mean and variance. This suggests a more principled approach to local surface shape enhancement: investigate how the different properties of BRDF slices may be altered to better convey local shape properties. Whether these material variations are physically-realistic (or better, physically-achievable by fabrication) is an independent, albeit intriguing question for future work.

The additional degrees of freedom provided by anisotropic materials (i.e., orientation, complex reflectance response) could also be employed to convey specific shape features in visually surprising ways. This direction of research will demand a better understanding of how surface curvatures become visually entangled with tangent field curvatures when rendering with anisotropic BRDFs. Indeed, it is highly likely that the perception of shape is easily altered by the use of curving tangent fields such as those we have manipulated in Section 3.2.2, as shown by the brushed metal artwork of Steve Prachyl (see Figure 3.29): here the metal surface is *flat*, hence apparent shape is solely due to tangent field curvatures. Future work on apparent shape should thus study the potential perceptual confusions between differential properties of the surface and of its tangent field, before trying to make use of these ambiguities either to create shape illusions or to better convey existing shape.




Figure 3.29: This artwork (“Metal Madonna”, by Steve Prachyl) has been produced by brushing a flat piece of metal in carefully chosen directions. As highlighted on the right, it produces a strong illusion of 3D shape: not only different parts appear to pass in front of others (outlined in red), but the front part appears to be curving and bending toward the view (purple mesh). The illusion is even more vivid when the main source of light is moved around.

Multi-scale analysis There are actually a few different strategies for computing surface or tangent field curvatures. The view-dependent surface curvatures of Section 3.1.1 are computed in screen-space using only visible neighbor points. In contrast, neighbor points were computed in object space in the method of Section 3.1.2, but required a base/relief separation. In later work, we have explored yet another approach, this time based on the fitting of algebraic spheres to neighbor points [37]. Our first variant [MBG⁺13] uses ray tracing to collect neighbor points; it has the advantage of working on fully dynamic shapes without requiring any pre-process, and has

been incorporated inside the Modo software. A second variant [MGB⁺12] works directly in object space, but collects neighbors at increasing 3D euclidean distances and measures the parametric variations of the fitted sphere to identify stable (i.e., pertinent) scales. It would be interesting to explore the potential of this fitting approach as a replacement of the method of Section 3.1.2: it could not only obviate the need of a base/relief separation, but also potentially suggest pertinent scales, a likely useful tool for shape exploration. The tangent field curvatures of Section 3.2.2 might also benefit from a proper multi-scale treatment; however, this will require to re-introduce the smoothed-out tangential variations inside the material function, as discussed further below.

A more involved research problem would be to identify object parts in a bottom-up fashion. This should help devising novel illustrative shading techniques that convey objects with complex structures (e.g., plants, molecules, organs, etc) in a more legible manner. It might also shed light on how specific shape patterns help convey a specific material appearance. A famous example is the veil effect mastered by 19th century sculptors, shown in Figure 3.30a: the transparency of the veil is artfully conveyed out of a homogeneous chunk of material. Similar perceptual correlation effects exist with anisotropic materials. This is easily observed by comparing a sphere and a folded surface rendered out of a same anisotropic silk material: the former looks hardly as 'silky' as the latter. The mechanical properties of materials may also be conveyed with the right choice of shape patterns, as shown in Figure 3.30b and c, which respectively convey the impression of a thin piece of cloth wrapped around a head and of a thick coat worn by a virtual character.




Figure 3.30: How specific shape patterns may affect perceived optical and mechanical material properties: (a) although made of the same material, the statue convincingly conveys the impression of a transparent veil; (b) carving folds out of wood gives a sense of delicacy and fragility that contrasts with what we know of the material; (c) through skillfull 3D sculpting, the artist is able to convey the thickness of the jacket even before working out its material appearance.

Material appearance The computation of view-centered surface properties at the right scale is not only necessary for enhancing shape through shading, but it is also required to correctly convey the material appearance of distant objects. Indeed, when surface details become smaller than the pixel size, two approaches may be taken: either sub-sample the pixel sufficiently to capture all variations, or pre-filter radiance inside the pixel area by incorporating surface details in the material function. The latter approach has the potential to be much more efficient and to

provide more insights on the image patterns produced by materials at various distances, as we have shown in Section 3.2.2. It is still an active area of research though, since not all material functions are directly amenable to pre-filtering. When a large number of surface details are found under a pixel, then statistical material models seem best adapted to the task. However, this is only true for distant objects or very fine surface details; in most situations, both smooth reflectance and sharp details remain visible over different scales. In other words, the material not only acts as a filter, but also as a texture, and modeling its multi-scale behavior is an exciting direction of future work.

The modeling of materials from surface (micro-)structures is only one side of material appearance; the other resides in the visual effects that are produced, as we have started to explore in Section 3.2.1. An essential task of this line of research is to identify the different material components that lead to visually different effects in an image. Besides the classic diffuse/specular separation, there are a number of other distinct visual components that remain to be identified, such as haze, luster or asperity scattering effects. This should again involve perceptual considerations, so that an ideal material model could not only be described in terms of its underlying physical parameters, but also in terms of perceptually-relevant visual effects. The former description not only enables multi-scale rendering, but also opens to the design and eventually to the fabrication of objects with desired material appearances. The latter description is more likely to appeal to artists for the manipulation of material appearance, and will permit to better structure the final image as discussed in the final chapter. Understanding the visual patterns produced by materials and their variations at an object surface is an important topic of future investigation [82], since this pervades our everyday life when trying to answer such questions as: Is this table wet or greasy? Would this object feel soft or hard to the touch? Could it be deformed, torn or teared apart? Has it been used recently or has it been sitting here for a while? Etc. Material designers already exploit our tendency to assign tactile properties by relying only on the visual appearance of materials, as shown in Figure 3.31.




Figure 3.31: Even though this material (Star ©Foglizzo Leather) looks like a vintage patina, it is actually soft and made of leather as revealed by the side view.

Lighting design We have mostly discussed two of the three main components that determine appearance in object space — namely material and shape — and left out lighting as if it were independent of the two others¹⁶. This is of course not so simple, as shown by the examples in

¹⁶The Light Warping technique only modifies lighting directions locally, and hence is not considered as a global lighting design tool.

Figure 3.32: lighting may be made to convey specific aspects of an object shape, and even alter its material appearance [123]. There has been some work on this topic in computer graphics, for instance with the goal of automating the orientation of the lighting environment to better convey shape or materials [124]. In the domain of human vision, experiments have shown the tremendous impact that lighting may have on the perception of shape or material (e.g., [125, 126]). Nevertheless, a proper understanding of how and why lighting may mingle with our perception of object shape and material is still lacking and represents a fascinating direction of future work.



Figure 3.32: The photograph on the left (©Urs Recher / broncolor) makes use of multiple soft box light sources and a black reflective backdrop to properly convey different shape features of the Buddha statue. Lighting may also strongly affect material appearance as shown on the right [123], where the photographer switches from a small sharp light source, to a large soft one.

CONCLUSIONS AND PERSPECTIVES

Our journey from the image plane to the object surface has made us encounter a variety of approaches to the manipulation of appearance. The projects I have been specifically involved in span the four stages of appearance, as shown in Figure 4.1. Yet many more research directions remain to be investigated, as pointed out in the discussions of Sections 2.3 and 3.3.




Figure 4.1: The techniques introduced in previous chapters span the four stages of appearance: from simple color images (in red) or images augmented with auxiliary buffers (in orange), to radiance scattered by a surface (in cyan) or inside its microscopic structure (in blue).

Looking back at Figure 4.1, there might seem that there is not a single recipe for manipulating appearance: approaches differ most notably depending on whether they work in the image plane or at the object surface, which is why I have devoted them different chapters. Nevertheless, over the course of these projects, I have come to realize that the same elements of appearance recur consistently. In other words, even if the recipes differ, the ingredients broadly remain the same.

Edges & contours represent natural image boundaries. They are at the heart of techniques based on diffusion curves where they act as drawing primitives. Methods manipulating image

flows may also be augmented with image boundaries that clip pixel neighborhoods. Automatically identifying them in bitmap images remains a difficult research problem. Edges & contours also implicitly emerge in surface flows and MatCap-based techniques, since any discontinuities in depth or normal buffers will show up as image discontinuities. In radiance tweaking techniques, they represent boundaries that should not be crossed when computing view-dependent curvatures. At the material design stage, the form taken by the material function at grazing angles potentially has a strong impact on the appearance of surface regions close to object boundaries.

Compression & stretching patterns are explicitly exaggerated in flow-guided warping, and even created in the surface flows technique. The nature of the distortion should differ when textures or shading are the source of such patterns, and disambiguating them in bitmap images is a challenging problem for future work. Such patterns should also reveal useful to extend the MatCap-based appearance manipulation techniques beyond simple spheres. In rendering, compression & stretching patterns are due to the interplay of shape, material and lighting. They may be exaggerated to convey shape, as in the light warping technique. When anisotropic materials are employed, more pronounced distortions are produced. Manipulating these distortions as in the micro-sculpt technique remains an intriguing direction of research.

Filters & textures constitute two different, yet complementary aspects of material appearance. The filtering due to the material may be visually confounded with lighting, and both might affect the way we perceive shape as shown in the flow-guided warping technique. Filter properties might be altered using blurring or sharpening operators, but these should be applied prior to distortion by shape, as we have shown in the surface flows and MatCap-based methods. Material appearance is actually more complex than a simple amount of blur: hazy gloss and directionally-smearred reflections (due to densely scratched materials) constitute compelling examples. Spatial variations of the material may be used to depict shape such as in radiance tweaking techniques. But they may also convey properties of the material itself, as we have shown with the glint lines and patterned highlights exhibited by scratched and brushed surfaces.

Layers & scales represent two different ways to structure the final image. A proper separation in different appearance layers is necessary, as we have explicitly done in the surface flows technique. Obtaining a similar layering from an input bitmap image represents a difficult (yet important) challenge when no additional information is provided. We have experimented with simple diffuse/specular separations both in expressive shading techniques and when studying measured materials; however, a more principled decomposition remains to be found. For some material qualities such as haziness, the decomposition should rely on visual perception. The notion of scale occurs in image flows techniques through the size of the pixel neighborhood that may be varied in the image. Finding the right scale(s) at the right location is an important research problem. The same issue occurs in object space when studying surface curvatures or larger shape features, and becomes crucial in the context of multi-scale rendering, where different image patterns will show up at different scales.

Tying together these elements in a perceptually-relevant theory of appearance that stretches from Physical Optics to Visual Perception is a difficult endeavor. In the next section, I will propose a blueprint for such a theory and identify remaining issues for future work.

4.1 Blueprint for a theory of appearance


We are now going to traverse the diagram of Figure 4.1 in the opposite direction, from “below” the object surface where physical or geometric optics apply, to “above” the image plane that is directly subject to visual perception.

4.1.1 Below the object surface

By “below the object surface” I mean everything that is not visible to the naked eye, which obviously depends on the distance to objects. In the following, we may thus irrespectively refer to the surface as that of an apple seen at arm length, or that of a canopy seen from an aircraft.


The optics of appearance A material acts as a black box that encapsulates all of light transport occurring at scales not directly observable given object distance and image resolution. It has another important property that distinguishes it from light transport in general: it should be *relocatable* [95]. Intuitively, this means that a material corresponds to the “stuff” an object is made of, and for a homogeneous material we expect that roughly the same stuff will be found at different surface locations, so that the optical response will be the same. When defined statistically, the material corresponds to the optical response expectation of *multiple realizations* of a same statistical process at the micro-scale. It is usually defined for a locally flat surface and a pair of collimated ingoing and outgoing beams (for BSDFs), and might possibly consider different entry and exit points (for BSSDFs) when lateral transport becomes significant.

Having outlined the general concept of a material, we may now look inside the black box at the different optics phenomena that should be accounted for. From a physical optics point of view, the vast majority of optical phenomena of importance to a theory of appearance stem from Maxwell equations, which describe the propagation of electro-magnetic waves in various media defined by their refractive indices. This is, however, too fine-grained both for practical and theoretical reasons, as any material would require extremely time-consuming simulations, which would in return provide little insights on material appearance. Fortunately, for a number of special configurations of interest, it has been possible to derive new formula from Maxwell equations: a famous example is Fresnel equations, which describe (among other things) the reflectivity and transmissivity of a perfectly flat interface between two media of different indices. There are many other configurations of interest, such as distributions of flat mirror-like micro-surfaces as frequently used in graphics, the layering of different media with potentially rough interfaces, distributions of particles on top of the surface or inside a particular medium, distributions of surface elements small enough to diffract light waves, etc. This requires an extensive survey of the optics literature [101], which is out of the scope of this draft account. We rather discuss a number of topics that should be part of a theory of appearance.



Structuration The way the 4D space of BSDFs or the 6D space of BSSDFs is “filled in” by a given material model will most likely reflect some properties of the micro-surface. A known example is that of isotropic BRDFs which only span three of the four general dimensions due to the rotational invariance of the underlying micro-structure. The space of BSDFs and BSSDFs may be structured further in a variety of ways: by making use of symmetries in the physical structure; by separating single, from double, from multiple scattering; etc.

However, it is likely that the structuration of a material that is best adapted to the modeling of appearance will be slightly more involved. The general idea will be to find families of light paths at the micro-scale that correspond to material components that can be characterized *par-simoniously* in a B[S]SDF. A good example is that of gloss models using the halfway vector parametrization: complex material responses may be obtained with a one-dimensional material function parametrized by the angle made by the halfway vector and the normal. The disentangling of the space of materials into parsimonious components might even require some consideration of human vision at this early stage, as we have shown in the case of hazy gloss.



The structuration of a material function might not only be a matter of decomposition, but also of integration, as is best exemplified with colors. Most models in optics characterize scattering spectrally: results depend on incoming light wavelengths, as well as on the spectral properties of the medium (e.g., the refractive index generally varies with wavelengths). In contrast, material appearance rather requires a tristimulus color model, so that color effects of the material may be directly characterized in perceptually-significant terms. The required spectral integration might become challenging for materials that exhibit peaked spectra.


Thanks to the linearity of light transport, the outcome of this structuration process will be a separation of the final image in material-related components. The classic diffuse and specular components are simple instances, but owing to the complexity of real-world materials, we should expect more complex decompositions.

Heterogeneity When considering heterogeneous materials, one should answer the question of how material components vary from one surface location to the other. In the general case of surface-varying BSDFs or BSSDFs, each material component may be described by a set of parameters (e.g., albedo, roughness), which themselves vary across the surface. No matter how these variations are produced (e.g., through painting, fluid flow, aging or weathering,) having the same set of material parameters at two different surface points should yield the same material response: hence relocatability remains valid.

In some specific instances, the spatial variations may be themselves considered as part of the material, such as in the case of randomly bumpy, glittery or scratched materials. When the material is observed from afar, these spatial variations are no longer resolved; hence they should be incorporated inside the material response, yielding a multi-scale model of material appearance. This will only work at a restricted range of scales, since the local curvature of the surface (or of the tangent field for anisotropic materials) may become too pronounced to satisfy the local planar assumption of BSDFs or BSSDFs.




When the object is covered with distributions of tiny scatterers that stick out of the surface, such as with furry or fluffy materials, then this same planar assumption becomes untenable: scatterers tend to spread apart in convex regions, and bundle together in concave regions. It then becomes necessary to consider a non-planar surface in the definition of the material function.



Heterogeneity of the material function thus gives rise to *textural* aspects of the material, whose visual importance might depend on the *scale* at which they appear.

Translucency Up until now we have made an a priori distinction between BSDFs and BSSDFs, assuming that one type of model or the other would be considered depending on the underlying physical structure. In practice, whenever lateral transport in the material can be neglected, a BSDF model is sufficient. In reality, the relevance of lateral transport to the final translucent appearance depends on *scale*: materials such as wax appear opaque at a sufficiently large distance; in other words, one could transition from a BSSDF to a BSDF in a multi-scale fashion.

However, with increasingly large lateral transport, the conventional assumption of a locally flat surface once again becomes untenable for BSSDF models. Take the example of a human head lit from behind: light will get absorbed for most of the head, except in *thin regions* such as the ear; in other words, light transport then depends on the object shape and is not relocatable anymore. One may still consider a conventional BSSDF model in this case (i.e., assuming the surface is flat), but it is likely that the resulting appearance will not be faithful to the ground truth. One challenge with translucency is thus to come up with new types of materials that are compatible with both relocatability and large lateral transport.



Future issues


- How should the connections between material components and light paths be established? Which aspects of a light transport simulation at the micro-scale are expected to structure the material function?
- What are the best models for reproducing the vibrant colors or complex transparent effects observed in nature as well as man-made objects? What approximations to Wave Optics could be tolerated in order to obtain parsimonious models?
- To which extent material components could be directly manipulated by artists? What would be the impact on material appearance? What could be visually (i.e., perceptually) tolerated in terms of material appearance manipulation?
- How may the physical micro-structure of a material be inferred from its material appearance? Could it be fabricated? How could unconventional micro-structures (e.g., distributions of nano-particles) help achieve surprising material effects?
- What would be the best approach to the modeling of dynamic materials, such as when a surface gets wet or dirty? How may the definition of a B[S]SDF be extended to account for non-parallel layers at a meso-scale?
- How should materials that depend on shape (e.g., translucent, granular, furry) be defined while maintaining relocatability? At which point the flat surface and collimated directions assumptions of a B[S]SDF will not be valid anymore?

4.1.2 Between the object and the image

We now consider a full 3D scene where objects are made of different materials, surrounded by one or more light sources, and a camera has been put in place. Only a subset of the scene is visible from the current viewpoint, and we consider that the level of details found in visible surface regions is locally adapted to the current scale. This means that even for a distant object, the

surface intersected by a *beam* of minimal section (e.g., pixel size) emanating from the camera is assumed to be gently curving; all the remaining high-frequency details of the original surface should thus be incorporated in the material definition. In the following, we will interchangeably use the terms view direction or view ray to refer to the central direction of a beam emanated from the viewpoint, and refer to its *footprint* to denote the intersection of the beam with a surface.

View-dependency If we intersect the view ray with the closest surface in the scene and express its direction in the local reference frame, we then obtain a *material slice* that consists in a 2D function for a BSDF slice and a 4D function for a BSSDF slice.¹ Intuitively, a material slice characterizes which parts of the incoming radiance will contribute the most to the radiance reflected or transmitted toward the viewpoint, which will greatly differ between a diffuse material where most of the environment contributes little energy, and a specular material where a small portion of the environment contributes the most. In spite of its reduced dimensionality, a material slice might still have a complex profile, which is likely to have a significant impact on appearance: it acts as an *angular filter* on the local environment lighting, and its behavior at grazing viewing angles might influence appearance close to *occluding contours*.



The first step in characterizing a material slice is to rely on the structuration of materials discussed in Section 4.1.1. When the material can be separated in additive components (such as diffuse + specular), then this translates to an additive decomposition of material slices. When the components are multiplicative (such as the Fresnel term in microfacet theory), then characterizing the effect on a slice will be more involved as it will likely depend on the view direction. The general dependence of a material slice on view direction is also a consequence of the symmetries in the material structure: for instance, isotropic BSDFs are insensitive to azimuthal variations of the view direction. Ideally, a material slice should be itself structured into a set of material lobes that vary in a number of parameters (e.g., intensity, color, direction, extent, and even location for BSSDF slices) and may be directly traced back to parameters of the corresponding material component, and thus by transitivity to parameters of the underlying physical structure (e.g., albedo, roughness, refractive index, etc).

Inside the footprint For a nearby surface point intersected by a neighboring view ray, both the material slice and the local lighting will likely vary, which should be taken into account. Our next step is then to look inside the footprint at infinitesimal variations of surface shape (i.e., normal and tangential curvatures), of material (i.e., spatially-varying parameters), and of incoming lighting (i.e., spatially-varying light field). The outcome of this study is the characterization of *radiance gradients*. Formally, the challenge consists in performing a differential analysis of the rendering equation that decomposes the gradient of emitted, reflected or transmitted radiance into a sum of gradients due to shape, lighting and material. An important prospect of this analysis is the explanation of *masking* and *facilitation* effects in shape perception that occur with specific configurations of material and lighting: if the radiance gradient does not contain any information about local shape features, then there is no hope for these features to be perceived; conversely, if the radiance gradient is exaggerated by material and lighting gradients that are




¹Since conventional material models assume collimated incoming and outgoing directions, the approximation of a locally orthographic view beam is also implicitly assumed.

congruent to shape features, then it is likely that these will appear as exceptionally legible in the final image.

The result of this analysis should not be taken literally as a perceptual decomposition of the image gradient for three reasons. First, it is likely that the human visual system will conflate some terms from the gradient decomposition; I would expect in particular that variations due to anisotropic reflections will likely be confused with classic reflections.² Second, the view projection and sensor properties might alter radiance variations as we will see at the end of this section. Third, what happens inside the footprint will project onto a very small region in the image and thus to a small portion of the retina; however, spatial pooling occurs even in the very first stages of vision (in the retina), which requires to study how gradients are connected in larger patterns. Nevertheless, such a differential analysis of the rendering equation will further decompose the imaging process, yielding *layers* related to variations of shape, material or lighting.


Outside the aperture Once we have decomposed radiance gradients at each and every point of a surface, we may connect them to form layered *gradient fields* living on the surface. Each layer thus formed corresponds to a different type of variations, which may be due to material textures, shadows and inter-reflections, environment lighting, or glossy reflections among other things. These fields are interdependent since they are governed by the same underlying surface shape (i.e., normal and tangential curvatures). As a result, they will be *compressed and stretched* in ways that are characteristic of differential shape properties across the object surface.



Visually identifying material properties is more delicate, as it requires to estimate *filtering effects* (e.g., directional blur, contrast, haze) in potentially broad surface neighborhoods around the aperture. The ideal situation is that of a surface that is slightly bent and smooth on the macro-scale: in this case, the material slice will vary little from one surface point to the other and may be considered to act as a convolution filter of the local environment lighting. When the surface has a more complex shape, the key challenge is to understand how to disentangle the distortions due to shape, from those of the material acting as a filter. In the case of rough surface features on the meso-scale, material filtering effects might become harder to discriminate simply because there is not enough pictorial space to observe filtering properties.

An additional difficulty is that shape features might occur at multiple *scales*, nested inside each other. Only a few scales are then relevant, and these may be considered as different *layers* on top of each other, with small-scale features “perturbing” larger ones. At the largest of scales, object parts and eventually the object as a whole may be considered in the way they globally reflect or transmit light from the environment. This should help understand how some lighting environments are better apt at conveying objects with given shapes and materials.

Projection & sensor The last stage carries out the previous analysis from the object surface to the image plane by projection. In doing so, it introduces *compressions* due to the orientation of the surface with respect to the view direction (i.e., the surface slant). The limiting behavior of this process is achieved at *occluding contours*, which are also due to projection. All the layered gradient fields get affected in the same way close to the contour on the occluding side,



²This would make sense from an evolutionary point of view because of the relatively rare occurrence of anisotropic materials in the natural environment.

creating correlations among layers in the image plane. Projection also hides away information through occlusion, either by a part of the same object or by a different object, placed in front. This has the effect of interrupting gradient fields when they pass behind an occluder, which is visually different from the compression of gradient fields on the occluder, and potentially acts as an important visual cue for parsing an image into different objects and parts.

Finally, the properties of the sensor itself will tend to add blur, glare or chromatic effects in the final image. One might think that these filter effects could be confused with those due to the material. However, sensor-induced effects are not tied to surfaces: they are not affected by distortions due to shape, and are not necessarily clipped at occluding contours.


Future issues

- How are the filter parameters of material slices systematically related to the physical parameters of B[S]SDFs in a way dependent on the view direction? What governs their behavior at grazing viewing angles?
- What is the general form of the decomposition of radiance gradients given arbitrary material functions and lighting environments? When are the gradient components confused perceptually and how might this confusion be used to control appearance?
- Once component gradients have been connected to form fields, what characterizes the resulting compression and stretching patterns? How do they differ or correlate between different layers, in particular close to occluding contours?
- To which extent is the filtering due to the material affected by compression and stretching due to shape? How should both material and shape be taken into account in the design of lighting environments as in Photography?
- What characterizes the meso-scale patterns and surface parts occurring on objects having complex shapes? How should they be analyzed to identify a few persistent scales? Where should light sources and reflectors be placed to convey them legibly?
- In which way should the multi-scale layered structure reaching the sensor be represented to grant further appearance manipulation in post-process?

4.1.3 Above the image plane

We have come back to the image plane, and are once again in the position to consider how an image is structured in ways that communicate shape, materials and lighting. However, the description of image patterns responsible for object appearance may now rely on the structuration of the image formation process that we have introduced in the last two sections.

Nested structure Objects in an image are bounded by their *occluding contours*, which are usually the locii of gradient fields discontinuities. The *layered* gradient fields running inside these boundaries are produced by variations in shading or texture, whose patterns of *compression and stretching* indicate the underlying surface shape. Since shape features may occur




at different finite *scales*, then gradient fields may also be decomposed at multiple scales, in a way similar to how some artists paint layers one after the other, progressively adding details.

In close-up views, an object material may appear as a *texture*, itself producing specific gradient fields (e.g., a distribution of visible scratches). In general, a material will also have the effect of a *filter* in the image, but its properties might not be directly observable as it will also be affected by compression and stretching patterns. In other words, these patterns form a distorted canvas onto which the material filter may be thought to operate. Since material functions are separated into different components, then the corresponding material filters are themselves *layered*: for instance, sharp highlights appear as put on top of other diffuse shading effects.

This cascaded set of transforms gives rise to a nested structure for each object depicted in the image, which are themselves connected by the global perspective cues in the image.

Disentangling patterns In order to exactly retrieve this nested structure from an image, one has to distinguish among patterns due to different physical causes. However, since only pictorial information is available, the disentanglement of image patterns will likely be a heuristic process. Different *layers* might be assumed to exhibit different orientations, color saturations, intensities, degrees of blur, etc; these criteria may then be used as priors to separate gradient fields of seemingly different nature. In this respect, it will be particularly interesting to study when humans make errors compared to ground-truth separations. This will happen in particular when different terms of the differential analysis of Section 4.1.2 become confounded. Another source of error is the coupling of material and lighting: for instance, with a low-frequency lighting, material components will appear blended together.




Once a layering of gradient fields has been retrieved, then inner image boundaries remains to be identified. A potential cue is then the set of field discontinuities which are typical of *occluding contours*. In particular, different layers tend to be compressed close to a contour on the occluder side, while there is no a priori reason that they exhibit this distortion on the occluded side.

Equipped with layered gradient fields clipped at occluding contours, the image may now be undistorted to study the *filtering properties* of the material. Such an inverse distortion might reveal challenging for complex shapes as it might not yield a bijective mapping. Fortunately, the goal of this analysis is not to produce an undistorted image, but rather to sufficiently discount distortions on local image neighborhoods, in order to recover filtering parameters. These estimated parameters might differ from the true filter parameters implemented by a material slice; but then again, the patterns of errors might be interesting in their own right to better understand the visual perception of materials.

Ambiguity & completion We have alluded to how patterns of errors in the identification of image structure might provide insights on human vision. However, in some situations, there is a priori not enough pictorial information for vision to do its work. This happens when the cues provided by an image might be compatible with more than a single interpretation. For instance, a smooth highlight may be due to a rough material reflecting a sharp light source, or a mirror material reflecting a soft light source. This ambiguous configuration appears to be effortlessly resolved by the human visual system. The trick might reside in the pooling of multiple observations from different locations inside a same object: if the same filter characteristics are identified in different places, then the material is assumed to be the cause of highlight smoothness

since this is the most generic interpretation at hand. Such a long-range pooling strategy might fail though, as some of the best illusions (e.g., using bi-stable stimuli) demonstrate effectively.

A much more common example of the partial information provided by pictorial cues is when an object is occluded by another object (or by itself). Once again, a variety of interpretations are possible; yet the human visual system usually ends up with so-called amodal completions that are close to the physical truth in most cases. Completion amounts to find a portion of the image structure that would appear if we removed the object in front. It is far from being a simple problem, since the objects might be touching and get compressed against each other, or be separated in depth and remain unaffected by each other, leading to altogether different image structures. A challenging direction of future work resides in the exploitation of the visible image structure for the completion of occluded regions.



Understanding the resilience to ambiguities exhibited by human vision is actually fundamental if we wish to identify image structure in paintings and drawings. In such cases, the image structure may depart much more profoundly from that of a photograph for instance: the properties of layered gradient fields, the visibility of occluding contours, and even the global deformation due to perspective effects might all be approximately conveyed. I believe this represents the ultimate test case for image structuration and appearance analysis.

Future issues

- How should the image creation process be organized to maintain a plausible nested image structure while drawing? How could it be used to manipulate appearance afterwards, yet without imposing constraints that sacrifice the style of the artist?
- Which heuristics permit to decompose an image into a nested structured representation? How are these heuristics related to the image formation process? To which extent do these relations explain why the heuristics might work at all?
- What are the patterns of errors made by the human visual system when it comes to image structuration? Is it possible to understand them in terms of the image formation process or do they rely on cognitive inferences? How might these errors be used to create controllable illusions?
- How are ambiguities resolved by means of a spatial pooling — or any other strategy — in Human Vision? How could these strategies be taken into account when the goal is to create legible images?
- To which extent are completion mechanisms able to recover plausible hidden parts? How do they rely on available image structure to that end? Would it be possible to modify an image structure so that amodal completion becomes easier?

4.2 Beyond appearance: Ecological Optics

As Jonathan Miller nicely expressed in his book on reflection [71]: “The natural world reveals itself in borrowed light”. Light, by the time it gets impinged on our retina, gets structured in formidably complex ways. Even though the ramifications of this process (down to Physical Optics) remain to be explored and cartographed, the basic concept had already been laid out by Gibson in 1961 in what he called Ecological Optics [127]. It is defined as the “Analysis of the ways in which radiant light interacts with surfaces of the environment, and, in turn, the manner in which the resulting reflected light is structured by those surfaces. The array of reflected light serves as potential information for perceiving, and as such creates the possibility for direct perception of environmental layout”. Whether or not perception *directly* works from this structure (as opposed to being mediated by *inferences*) remains a debate that belongs to the domain of Psychology, and thus clearly falls out of the scope of this dissertation.




Figure 4.2: Left: the optic array of light reaching the eye has a complex structure and stops at the boundaries of the ocular window. Right: during observer motion, parts of the scene disappears (deletions) while others appear (accretions) around occluding contours.

As is apparent on a first reading of Gibson’s work, there’s more than appearance that meets the eye. Even though he insists on the importance of occluding contours, his emphasis is on the way they structure images *in motion* through the production of accretions and deletions, which denote the appearance and disappearance of image patterns behind an occluder depending on the direction of motion (see Figure 4.2). Contours help build a layout that informs on proximity relationships of objects in a scene, whether they touch or sit upon each others, etc. Gibson even suggests that disparity in stereoscopic vision merely provides the cues that would be brought about by an instantaneous movement. I find the idea of explicitly controlling or altering motion to manipulate the way a scene layout is perceived to be a fascinating direction of research.

Motion is not only limited to object surfaces, but also concerns thin objects or even volumetric effects. Thin objects usually come in great numbers, such as fur, hair or the dense foliage of a tree seen at an intermediate distance. Thanks to the group motion of objects (e.g., leaves on a same branch), our visual system usually finds no problem at perceiving such myriads of objects in depth. To some extent, perceiving the layout of complex volumetric (even dynamic non-rigid) effects such as smoke or clouds is also made possible by the (rigid) motion of the viewpoint. The connection with stereopsis is striking here as well, as is made evident by looking at a tree foliage

with both eyes open, then closing one eye: it then becomes difficult to tell which branches are in front of others! Said differently, the layout then seems to vanish! Even though I have set aside the study of tiny objects or volumetric effects in this document since they do not exhibit an easily identifiable visible surface, I believe there are potentially many ways to manipulate them in images once we consider motion cues. In the case of static images, layout may still be conveyed by depicting *potential motion* of the observer. This already occurs in scientific illustrations where combinations of halos and contour lines are used to emphasize image regions susceptible to undergo accretion or deletion with a slight motion of the viewpoint.

The scene layout perceived at each and every moment allows us to orient ourselves in the environment, and the way objects appear to us permit to recognize their shape and material without having to come closer and touch them. What happens next when we start to interact with our environment? What are the image cues that allow us to act on objects, to manipulate them? These activities quickly require a perception of objects *dynamics*, which are related to the mechanical properties of objects: their viscosity, fragility, elasticity, weight, etc. Understanding how these are conveyed through motion, layout and appearance cues represents a major challenge that has defied researchers in the Human Vision and Artificial Intelligence communities.

Finally, the environment is not simply filled by inanimate objects, but it is also populated by living creatures. Of course, the images they produce might be studied with the same tools since creatures exhibit specific shapes and (optical and mechanical) materials, and they may be located with respect to other objects or creatures in the environment. What then makes them different from other objects from a purely visual point of view? The short answer is that they move differently; yet it is not clear how we perceive *animacy* in objects, and how these percepts rely on more than motion cues. Nevertheless, with enough artistic skills, motion cues are sufficient to give the illusion of life as demonstrated by more than a century of hand-made 2D animation!

In conclusion, it appears that we should once again take inspiration from artistic practices in the study of visual cues to layout, motion, dynamics and animacy, which constitute different elements of image structure lying at the heart of Ecological Optics. I, for my part, believe that a true understanding of these all-important perceptual cues will require a *systematic* account, and I do not see any better tool than Computer Graphics for this purpose.

AUTHOR'S BIBLIOGRAPHY

- [ABG⁺12] Lucas Ammann, Pascal Barla, Xavier Granier, Gael Guennebaud, and Patrick Reuter. Surface Relief Analysis for Illustrative Shading. *Computer Graphics Forum*, 31(4):1481–1490, June 2012.
- [BBG12] Simon Boyé, Pascal Barla, and Gael Guennebaud. A Vectorial Solver for Free-form Vector Gradient. *ACM Transaction on Graphics*, page p., November 2012.
- [CGBG13] Jiazhou Chen, Gael Guennebaud, Pascal Barla, and Xavier Granier. Non-oriented MLS Gradient Fields. *Computer Graphics Forum*, page p., December 2013.
- [MBG⁺13] Nicolas Mellado, Pascal Barla, Gael Guennebaud, Patrick Reuter, and Gregory Duquesne. Screen-Space Curvature for Production-Quality Rendering and Compositing. In *SIGGRAPH Talk Program*, Anaheim, United States, August 2013. ACM New York.
- [MGB⁺12] Nicolas Mellado, Gaël Guennebaud, Pascal Barla, Patrick Reuter, and Christophe Schlick. Growing Least Squares for the Analysis of Manifolds in Scale-Space. *Computer Graphics Forum*, 31(5):1691–1701, July 2012.
- [OBBT07] Alexandrina Orzan, Adrien Bousseau, Pascal Barla, and Joëlle Thollot. Structure-preserving manipulation of photographs. In *International Symposium on Non-Photorealistic Animation and Rendering (NPAR)*, pages 103–110, San Diego, United States, August 2007. ACM.
- [OBW⁺08] Alexandrina Orzan, Adrien Bousseau, Holger Winnemöller, Pascal Barla, Joëlle Thollot, and David Salesin. Diffusion Curves: A Vector Representation for Smooth-Shaded Images. *ACM Transactions on Graphics*, 27(3):92:1–8, August 2008.
- [RGB⁺14] Boris Raymond, Gaël Guennebaud, Pascal Barla, Romain Pacanowski, and Xavier Granier. Optimizing BRDF Orientations for the Manipulation of Anisotropic Highlights. *Computer Graphics Forum*, 33(2):9, April 2014.
- [RGB16] Boris Raymond, Gael Guennebaud, and Pascal Barla. Multi-Scale Rendering of Scratched Materials using a Structured SV-BRDF Model. *ACM Transactions on Graphics*, July 2016.

- [VB15] Romain Vergne and Pascal Barla.
Designing Gratin, A GPU-Tailored Node-Based System.
Journal of Computer Graphics Techniques, 4(4):54–71, 2015.
- [VBBF16] Romain Vergne, Pascal Barla, Georges-Pierre Bonneau, and Roland Fleming.
Flow-Guided Warping for Image-Based Shape Manipulation.
ACM Transactions on Graphics, 34(4):Article No. 93, July 2016.
- [VBF17] Peter Vangorp, Pascal Barla, and Roland W Fleming.
The Perception of Hazy Gloss.
Journal of Vision, 17(19), June 2017.
- [VBFG12] Romain Vergne, Pascal Barla, Roland Fleming, and Xavier Granier.
Surface Flows for Image-based Shading Design.
ACM Transactions on Graphics, 31(3):94:1–94:9, August 2012.
- [VPB⁺09] Romain Vergne, Romain Pacanowski, Pascal Barla, Xavier Granier, and Christophe Schlick.
Light Warping for Enhanced Surface Depiction.
ACM Transaction on Graphics (Proceedings of SIGGRAPH 2009), 28(3), 2009.
- [VPB⁺11] Romain Vergne, Romain Pacanowski, Pascal Barla, Xavier Granier, and Christophe Schlick.
Improving Shape Depiction under Arbitrary Rendering.
IEEE Transactions on Visualization and Computer Graphics, 17(8):1071 – 1081, June 2011.
- [ZBB⁺15] Carlos J. Zubiaga, Laurent Belcour, Carles Bosch, Adolfo Muñoz, and Pascal Barla.
Statistical Analysis of Bidirectional Reflectance Distribution Functions.
In *Measuring, Modeling, and Reproducing Material Appearance 2015*, San Francisco, United States, February 2015.
- [ZGVB16] Carlos J. Zubiaga, Gael Guennebaud, Romain Vergne, and Pascal Barla.
Local Shape Editing at the Compositing Stage.
In *27th Eurographics Symposium on Rendering (ESRG)*, EGSR '16 Proceedings of the Eurographics Symposium on Rendering: Experimental Ideas & Implementations, pages 23–32, Dublin, Ireland, June 2016. The Eurographics Association.
- [ZMB⁺15] Carlos J. Zubiaga, Adolfo Muñoz, Laurent Belcour, Carles Bosch, and Pascal Barla.
MatCap Decomposition for Dynamic Appearance Manipulation.
In *Eurographics Symposium on Rendering 2015*, Darmstadt, Germany, June 2015.

BIBLIOGRAPHY

- [1] R. W. Fleming, A. Torralba, and E. H. Adelson, “Specular reflections and the perception of shape,” *J. Vis.*, vol. 4, pp. 798–820, 9 2004.
- [2] I. E. Sutherland, *Sketchpad: A man-machine graphical communication system (Outstanding dissertations in the computer sciences)*. New York, NY, USA: Garland Publishing, Inc., 1980.
- [3] A. R. Smith, “Digital Paint Systems: An Anecdotal and Historical Overview,” *IEEE Annals of the History of Computing*, vol. 23, pp. 4–30, apr–jun 2001.
- [4] J. H. Elder, “Are edges incomplete?,” *International Journal of Computer Vision*, vol. 34, no. 2-3, pp. 97–122, 1999.
- [5] J. H. Elder and R. M. Goldberg, “Image editing in the contour domain,” *IEEE Transactions on Pattern Analysis Machine Intelligence*, vol. 23, no. 3, pp. 291–296, 2001.
- [6] J. Canny, “A computational approach to edge detection,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, no. 6, pp. 679–698, 1986.
- [7] T. Lindeberg, “Edge detection and ridge detection with automatic scale selection,” in *CVPR ’96: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, p. 465, 1996.
- [8] H. Hnaidi, E. Guerin, S. Akkouche, A. Peytavie, and E. Galin, “Feature based terrain generation using diffusion equation,” *Computer Graphics Forum*, 2010.
- [9] K. Takayama, O. Sorkine, A. Nealen, and T. Igarashi, “Volumetric modeling with diffusion surfaces,” *ACM Trans. Graph.*, vol. 29, pp. 180:1–180:8, Dec. 2010.
- [10] H. Winnemöller, A. Orzan, L. Boissieux, and J. Thollot, “Texture design and draping in 2d images,” *Computer Graphics Forum*, vol. 28, no. 4, pp. 1091–1099, 2009.
- [11] S. Jeschke, D. Cline, and P. Wonka, “Rendering surface details with diffusion curves,” *ACM Trans. Graph.*, vol. 28, pp. 117:1–117:8, Dec. 2009.
- [12] S. Jeschke, D. Cline, and P. Wonka, “Estimating color and texture parameters for vector graphics,” *Computer Graphics Forum*, vol. 30, no. 2, pp. 523–532, 2011.
- [13] G. Xie, X. Sun, X. Tong, and D. Nowrouzezahrai, “Hierarchical diffusion curves for accurate automatic image vectorization,” *ACM Trans. Graph.*, vol. 33, pp. 230:1–230:11, Nov. 2014.
- [14] S. Zhao, F. Durand, and C. Zheng, “Inverse diffusion curves using shape optimization,” *CoRR*, vol. abs/1610.02769, 2016.

BIBLIOGRAPHY

- [15] H. Bezerra, E. Eisemann, D. DeCarlo, and J. Thollot, “Diffusion constraints for vector graphics,” in *Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering*, NPAR ’10, (New York, NY, USA), pp. 35–42, ACM, 2010.
- [16] M. Finch, J. Snyder, and H. Hoppe, “Freeform vector graphics with controlled thin-plate splines,” *ACM Trans. Graph.*, vol. 30, pp. 166:1–166:10, Dec. 2011.
- [17] S. Jeschke, “Generalized diffusion curves: An improved vector representation for smooth-shaded images,” *Comput. Graph. Forum*, vol. 35, pp. 71–79, May 2016.
- [18] S. Jeschke, D. Cline, and P. Wonka, “A gpu laplacian solver for diffusion curves and poisson image editing,” *ACM Trans. Graph.*, vol. 28, pp. 116:1–116:8, Dec. 2009.
- [19] J. C. Bowers, J. Leahey, and R. Wang, “A ray tracing approach to diffusion curves,” *Computer Graphics Forum*, vol. 30, no. 4, pp. 1345–1352, 2011.
- [20] R. Prévost, W. Jarosz, and O. Sorkine-Hornung, “A vectorial framework for ray traced diffusion curves,” *Computer Graphics Forum*, vol. 34, no. 1, pp. 253–264, 2015.
- [21] W. M. Pang, J. Qin, M. Cohen, P. A. Heng, and K. S. Choi, “Fast rendering of diffusion curves with triangles,” *IEEE Computer Graphics and Applications*, vol. 32, pp. 68–78, July 2012.
- [22] X. Sun, G. Xie, Y. Dong, S. Lin, W. Xu, W. Wang, X. Tong, and B. Guo, “Diffusion curve textures for resolution independent texture mapping,” *ACM Trans. Graph.*, vol. 31, pp. 74:1–74:9, July 2012.
- [23] P. Ilbery, L. Kendall, C. Concolato, and M. McCosker, “Biharmonic diffusion curve images from boundary elements,” *ACM Trans. Graph.*, vol. 32, pp. 219:1–219:12, Nov. 2013.
- [24] T. Sun, P. Thamjaroenporn, and C. Zheng, “Fast multipole representation of diffusion curves and points,” *ACM Trans. Graph.*, vol. 33, pp. 53:1–53:12, July 2014.
- [25] O. Ben-Shahar and S. Zucker, “The perceptual organization of texture flow: A contextual inference approach,” *IEEE Transaction on Pattern Analysis and Machine Intelligence*, vol. 25, no. 4, pp. 401–417, 2003.
- [26] O. Ben-Shahar, Y. Vasilyev, Y. Adato, and T. Zickler, “Shape from specular flow,” *IEEE Transactions on Pattern Analysis & Machine Intelligence*, vol. 32, pp. 2054–2070, 2010.
- [27] R. W. Fleming, D. Holtmann-Rice, and H. H. Bühlhoff, “Estimation of 3d shape from image orientations,” *Proceedings of the National Academy of Sciences*, vol. 108, no. 51, pp. 20438–20443, 2011.
- [28] H. Kang, S. Lee, and C. Chui, “Coherent line drawing,” in *NPAR’07: Proc. symposium on Non-photorealistic animation and rendering*, pp. 43–50, ACM/Eurographics, 2007.
- [29] H. Kang, S. Lee, and C. K. Chui, “Flow-based image abstraction,” *IEEE Trans. Visualization and Comput. Graph.*, vol. 15, no. 1, pp. 62–76, 2009.
- [30] A. Hertzmann, “Painterly rendering with curved brush strokes of multiple sizes,” in *Proc. SIGGRAPH ’98*, pp. 453–460, ACM, 1998.

-
- [31] J. Hays and I. Essa, “Image and video based painterly animation,” in *NPAR’04: Proc. symposium on Non-photorealistic animation and rendering*, pp. 113–120, ACM/Eurographics, 2004.
- [32] G. Papari and N. Petkov, “Continuous glass patterns for painterly rendering,” *IEEE Trans. Image Processing*, vol. 18, no. 3, pp. 652–664, 2009.
- [33] J. Bigün and G. H. Granlund, “Optimal orientation detection of linear symmetry,” in *Proc. IEEE International Conference on Computer Vision*, pp. 433–438, 1987.
- [34] E. Zhang, J. Hays, and G. Turk, “Interactive tensor field design and visualization on surfaces,” *IEEE Trans. Visualization and Computer Graphics*, vol. 13, no. 1, pp. 94–107, 2007.
- [35] J. Kyprianidis and H. Kang, “Image and video abstraction by coherence-enhancing filtering,” *Computer Graphics Forum*, vol. 30, no. 2, pp. 593–602, 2011.
- [36] J. Kyprianidis, H. Kang, and J. Döllner, “Image and video abstraction by anisotropic kuwahara filtering,” *Computer Graphics Forum*, vol. 28, no. 7, pp. 1955–1963, 2009.
- [37] G. Guennebaud and M. Gross, “Algebraic point set surfaces,” *ACM Trans. Graph.*, vol. 26, July 2007.
- [38] B. Cabral and L. Leedom, “Imaging vector fields using line integral convolution,” in *Proc. 20th annual conference on Computer graphics and interactive techniques*, pp. 263–270, ACM, 1993.
- [39] H. Barrow, “Recovering intrinsic scene characteristics from images,” *Computer Vision Systems*, pp. 3–26, 1978.
- [40] A. Bousseau, S. Paris, and F. Durand, “User assisted intrinsic images,” *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2009)*, vol. 28, no. 5, 2009.
- [41] R. Carroll, R. Ramamoorthi, and M. Agrawala, “Illumination decomposition for material recoloring with consistent interreflections,” *ACM Trans. Graph.*, vol. 30, pp. 43:1–43:10, July 2011.
- [42] S. K. Yeung, C.-K. Tang, M. S. Brown, and S. B. Kang, “Matting and compositing of transparent and refractive objects,” *ACM Trans. Graph.*, vol. 30, no. 1, pp. 2:1–2:13, 2011.
- [43] E. A. Khan, E. Reinhard, R. W. Fleming, and H. H. Bühlhoff, “Image-based material editing,” *ACM Trans. Graph.*, vol. 25, pp. 654–663, 2006.
- [44] D. Gutierrez, F. J. Seron, J. Lopez-Moreno, M. P. Sanchez, J. Fandos, and E. Reinhard, “Depicting procedural caustics in single images,” *ACM Trans. Graph.*, vol. 27, pp. 120:1–120:9, Dec. 2008.
- [45] T. Chen, Z. Zhu, A. Shamir, S.-M. Hu, and D. Cohen-Or, “3swepp: Extracting editable objects from a single photo,” *ACM Trans. Graph.*, vol. 32, pp. 195:1–195:10, Nov. 2013.
- [46] N. Kholgade, T. Simon, A. Efros, and Y. Sheikh, “3d object manipulation in a single photograph using stock 3d models,” *ACM Trans. Graph.*, vol. 33, pp. 127:1–127:12, July 2014.

BIBLIOGRAPHY

- [47] M. W. A. Wijntjes, K. Doerschner, G. Kucukoglu, and S. C. Pont, “Relative flattening between velvet and matte 3d shapes: Evidence for similar shape-from-shading computations,” *Journal of Vision*, vol. 12, no. 1, p. 2, 2012.
- [48] F. Caniard and R. W. Fleming, “Distortion in 3d shape estimation with changes in illumination,” in *Proceedings of the 4th Symposium on Applied Perception in Graphics and Visualization*, APGV ’07, (New York, NY, USA), pp. 99–105, ACM, 2007.
- [49] S. W. Mooney and B. L. Anderson, “Specular image structure modulates the perception of three-dimensional shape,” *Current Biology*, vol. 24, no. 22, pp. 2737 – 2742, 2014.
- [50] J. E. Kyprianidis and J. Döllner, “Image abstraction by structure adaptive filtering,” in *Proc. EG UK Theory and Practice of Computer Graphics*, pp. 51–58, 2008.
- [51] J. T. Kajiya, “The rendering equation,” *SIGGRAPH Comput. Graph.*, vol. 20, pp. 143–150, Aug. 1986.
- [52] R. Ramamoorthi, D. Mahajan, and P. Belhumeur, “A first-order analysis of lighting, shading, and shadows,” *ACM Trans. Graph.*, vol. 26, no. 1, 2007.
- [53] D. Shepard, “A two-dimensional interpolation function for irregularly-spaced data,” in *Proc. 1968 23rd ACM national conference*, (New York, NY, USA), pp. 517–524, ACM, 1968.
- [54] S. F. Johnston, “Lumo: Illumination for cel animation,” in *Proceedings of the 2Nd International Symposium on Non-photorealistic Animation and Rendering*, NPAR ’02, (New York, NY, USA), pp. 45–ff, ACM, 2002.
- [55] I. Motoyoshi, “Highlight-shading relationship as a cue for the perception of translucent and transparent materials.,” *J Vis*, vol. 10, no. 9, p. 6, 2010.
- [56] M. Okabe, Y. Matsushita, L. Shen, and T. Igarashi, “Illumination Brush: Interactive Design of All-Frequency Lighting,” in *Pacific Conference on Computer Graphics and Applications*, pp. 171–180, 2007.
- [57] F. Pellacini, F. Battaglia, R. K. Morley, and A. Finkelstein, “Lighting with paint,” *ACM Trans. Graph.*, vol. 26, no. 2, 2007.
- [58] F. Pellacini, “envylight: an interface for editing natural illumination,” *ACM Trans. Graph.*, vol. 29, pp. 34:1–34:8, 2010.
- [59] T. Ritschel, M. Okabe, T. Thormählen, and H.-P. Seidel, “Interactive reflection editing,” *ACM Trans. Graph.*, vol. 28, no. 5, pp. 129:1–129:7, 2009.
- [60] T. Ritschel, T. Thormählen, C. Dachsbacher, J. Kautz, and H.-P. Seidel, “Interactive on-surface signal deformation,” *ACM Trans. Graph.*, vol. 29, no. 4, 2010.
- [61] P.-P. J. Sloan, W. Martin, A. Gooch, and B. Gooch, “The lit sphere: A model for capturing npr shading from art,” in *Proc. of Graphics Interface 2001*, pp. 143–150, 2001.
- [62] R. W. Fleming, R. O. Dror, and E. H. Adelson, “Real-world illumination and the perception of surface reflectance properties,” *Journal of vision*, vol. 3, pp. 347–68, Jan. 2003.

-
- [63] K. Doerschner, H. Boyaci, and L. Maloney, “Estimating the glossiness transfer function induced by illumination change and testing its transitivity,” *J. of Vision*, vol. 10, pp. 1–9, 2010.
- [64] P. J. Marlow, J. Kim, and B. L. Anderson, “The perception and misperception of specular surface reflectance,” *Current biology*, vol. 22, pp. 1909–13, Oct. 2012.
- [65] W. Heidrich and H.-P. Seidel, “View-independent environment maps,” in *Proceedings of the Workshop on Graphics Hardware*, (New York, NY, USA), pp. 39–ff., ACM, 1998.
- [66] S. R. Sternberg, “Grayscale morphology,” *Comput. Vision Graph. Image Process.*, vol. 35, pp. 333–355, Sept. 1986.
- [67] R. Ramamoorthi and P. Hanrahan, “A signal-processing framework for inverse rendering,” *ACM Siggraph*, 2001.
- [68] F. Durand, N. Holzschuch, C. Soler, E. Chan, and F. X. Sillion, “A Frequency Analysis of Light Transport,” *ACM Trans. on Graphics*, vol. 24, pp. 1115 – 1126, Aug. 2005.
- [69] P. Pérez, M. Gangnet, and A. Blake, “Poisson image editing,” *ACM Trans. Graph.*, vol. 22, pp. 313–318, July 2003.
- [70] R. Ramamoorthi and P. Hanrahan, “An efficient representation for irradiance environment maps,” in *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH ’01, (New York, NY, USA), pp. 497–500, ACM, 2001.
- [71] J. Miller, V. Mendes, and N. G. G. Britain), *On Reflection*. National Gallery Publications, 1998.
- [72] M. Gleicher, G. Cipriano, and G. N. P. Jr., “Multi-scale surface descriptors,” *IEEE Transactions on Visualization & Computer Graphics*, vol. 15, pp. 1201–1208, 2009.
- [73] G. Miller, “Efficient algorithms for local and global accessibility shading,” in *SIGGRAPH ’94: Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pp. 319–326, ACM, 1994.
- [74] G. Kindlmann, R. Whitaker, T. Tasdizen, and T. Möller, “Curvature-Based Transfer Functions for Direct Volume Rendering: Methods and Applications,” in *Proc. IEEE Visualization 2003*, pp. 513–520, October 2003.
- [75] H. H. B. Michael J. Tarr, Daniel Kersten, “Why the visual recognition system might encode the effects of illumination,” *Vision Research*, vol. 38, no. 15-16, pp. 2259–2275, 1998.
- [76] Y.-X. Ho, M. S. Landy, and L. T. Maloney, “How direction of illumination affects visually perceived surface roughness,” *J. Vis.*, vol. 6, pp. 634–648, 5 2006.
- [77] Y. Ostrovsky, P. Cavanagh, and P. Sinha, “Perceiving Illumination Inconsistencies in Scenes,” in *MIT AIM*, 2001.
- [78] R. W. Fleming and M. Singh, “Visual perception of 3d shape,” in *ACM SIGGRAPH 2009 Courses*, SIGGRAPH ’09, (New York, NY, USA), pp. 24:1–24:94, ACM, 2009.

BIBLIOGRAPHY

- [79] M. Ashikhmin, S. Premoze, and P. Shirley, “A microfacet-based BRDF generator,” in *Proc. ACM SIGGRAPH '00*, pp. 65–74, ACM, 2000.
- [80] D. A. v. Koenderink J.J., *The visual neurosciences*, ch. Shape and shading, pp. 1090–1105. MIT Press, Cambridge, 2003.
- [81] J. P. O’Shea, M. S. Banks, and M. Agrawala, “The assumed light direction for perceiving shape from shading,” in *APGV '08: Proc. symposium on Applied perception in graphics and visualization*, pp. 135–142, ACM, 2008.
- [82] E. H. Adelson, “On seeing stuff: the perception of materials by humans and machines,” in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* (B. E. Rogowitz and T. N. Pappas, eds.), vol. 4299 of *Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference*, pp. 1–12, June 2001.
- [83] P. Vangorp, J. Laurijssen, and P. Dutré, “The influence of shape on the perception of material reflectance,” *ACM Trans. Graph. (Proc. SIGGRAPH 2007)*, vol. 26, no. 3, p. 77, 2007.
- [84] K. J. Pont S.C., *Computer Analysis of Images and Patterns*, ch. Illuminance flow, pp. 90–97. Springer, Berlin, 2003.
- [85] P. Wood and P. McDonnell, *Scientific Illustration: A Guide to Biological, Zoological, and Medical Rendering Techniques, Design, Printing, and Display*. A VNR book, Wiley, 1994.
- [86] F. Cazals and M. Pouget, “Estimating differential quantities using polynomial fitting of osculating jets,” *Computer Aided Geometric Design*, vol. 22, no. 2, 2005.
- [87] G. Cipriano, G. Philips Jr, and M. Gleicher, “Multi-Scale Surface Descriptors,” *IEEE Trans. Visualization and Comput. Graph.*, vol. 15, no. 6, pp. 1201–1208, 2009.
- [88] Y. Ohtake, A. Belyaev, and H.-P. Seidel, “Ridge-valley lines on meshes via implicit surface fitting,” in *ACM SIGGRAPH 2004 Papers*, SIGGRAPH '04, (New York, NY, USA), pp. 609–612, ACM, 2004.
- [89] M. Kolomenkin, I. Shimshoni, and A. Tal, “Demarcating Curves for Shape Illustration,” *ACM Trans. Graph. (Proc. SIGGRAPH Asia)*, vol. 27, no. 5, 2008.
- [90] P. Cignoni, R. Scopigno, and M. Tarini, “A simple normal enhancement technique for interactive non-photorealistic renderings,” *Computers & Graphics*, vol. 29, pp. 125–133, 2005.
- [91] X. Zhang, W. Chen, J. Fang, R. Wang, and Q. Peng, “Perceptually-motivated Shape Exaggeration,” *The Visual Comp. (Proc. CGI)*, vol. 26, no. 6-8, pp. 985–995, 2010.
- [92] S. Rusinkiewicz, M. Burns, and D. DeCarlo, “Exaggerated shading for depicting shape and detail,” *ACM Trans. Graph. (Proc. SIGGRAPH)*, vol. 25, p. 1199, July 2006.
- [93] R. Zatzarinni, A. Tal, and A. Shamir, “Relief Analysis and Extraction,” *ACM Trans. Graph. (Proc. SIGGRAPH Asia)*, vol. 28, no. 5, pp. 1–10, 2009.

- [94] G. Harary and A. Tal, “3d euler spirals for 3d curve completion,” in *Proceedings of the Twenty-sixth Annual Symposium on Computational Geometry*, SoCG ’10, (New York, NY, USA), pp. 393–402, ACM, 2010.
- [95] F. Nicodemus, *Geometrical Considerations and Nomenclature for Reflectance*. NBS Monograph, U.S. Government Printing Office, 1977.
- [96] H. W. Jensen, S. R. Marschner, M. Levoy, and P. Hanrahan, “A practical model for subsurface light transport,” in *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH ’01, (New York, NY, USA), pp. 511–518, ACM, 2001.
- [97] “Anyone can cook — inside ratatouille’s kitchen.” <https://graphics.pixar.com/library/AnyoneCanCook/paper.pdf>, 2007.
- [98] B. T. Phong, “Illumination for computer generated pictures,” *Commun. ACM*, vol. 18, pp. 311–317, June 1975.
- [99] K. E. Torrance and E. M. Sparrow, “Theory for off-specular reflection from roughened surfaces,” *J. Opt. Soc. Am.*, vol. 57, pp. 1105–1112, Sep 1967.
- [100] E. Heitz, “Understanding the masking-shadowing function in microfacet-based brdfs,” *Journal of Computer Graphics Techniques (JCGT)*, vol. 3, pp. 48–107, June 2014.
- [101] E. Hecht, *Optics*. Addison-Wesley, 4th edition ed., 2001.
- [102] P. Beckmann and A. Spizzichino, *The scattering of electromagnetic waves from rough surfaces*. International series of monographs on electromagnetic waves, Pergamon Press; [distributed in the Western Hemisphere by Macmillan, New York], 1963.
- [103] D. Guarnera, G. Guarnera, A. Ghosh, C. Denk, and M. Glencross, “Brdf representation and acquisition,” *Comput. Graph. Forum*, vol. 35, pp. 625–650, May 2016.
- [104] R. S. Hunter, *The measurement of appearance*. John Wiley & Sons, 1987.
- [105] J. Wills, S. Agarwal, D. Kriegman, and S. Belongie, “Toward a perceptual space for gloss,” *ACM Trans. Graph.*, vol. 28, pp. 103:1–103:15, Sept. 2009.
- [106] A. Ngan, F. Durand, and W. Matusik, “Experimental analysis of brdf models,” in *Proceedings of the Eurographics Symposium on Rendering*, pp. 117–226, Eurographics Association, 2005.
- [107] W. Matusik, H. Pfister, M. Brand, and L. McMillan, “A data-driven reflectance model,” *ACM Transactions on Graphics*, vol. 22, no. 3, pp. 759–769, 2003.
- [108] J. Koenderink and S. Pont, “The secret of velvety skin,” *Mach. Vision Appl.*, vol. 14, pp. 260–268, Sept. 2003.

- [109] B. Walter, S. R. Marschner, H. Li, and K. E. Torrance, “Microfacet models for refraction through rough surfaces,” in *Proceedings of the 18th Eurographics Conference on Rendering Techniques*, EGSR’07, (Aire-la-Ville, Switzerland, Switzerland), pp. 195–206, Eurographics Association, 2007.
- [110] A. Weidlich and A. Wilkie, “Thinking in layers: Modeling with layered materials,” in *SIGGRAPH Asia 2011 Courses*, SA ’11, (New York, NY, USA), pp. 20:1–20:43, ACM, 2011.
- [111] N. Holzschuch and R. Pacanowski, “A Two-Scale Microfacet Reflectance Model Combining Reflection and Diffraction,” *ACM Transactions on Graphics*, vol. 36, p. 12, Aug. 2017. Article 66.
- [112] F. X. D. O’Donnell and F. W. Billmeyer, Jr., “Psychometric scaling of gloss,” in *Review and Evaluation of Appearance: Methods and Techniques. ASTM STP 914* (J. J. Rennilson and W. N. Hale, eds.), (West Conshohocken, PA), pp. 14–32, American Society for Testing and Materials, 1986.
- [113] F. Pellacini, J. A. Ferwerda, and D. P. Greenberg, “Toward a psychophysically-based light reflection model for image synthesis,” in *Proceedings of ACM SIGGRAPH 2000* (K. Akeley, ed.), (New York, NY), pp. 55–64, ACM, 2000.
- [114] G. J. Ward, “Measuring and modeling anisotropic reflection,” *SIGGRAPH Comput. Graph.*, vol. 26, pp. 265–272, July 1992.
- [115] ASTM, “Standard test methods for measurement of gloss of high-gloss surfaces by goniphotometry,” Standard E430-97, American Society for Testing and Materials, West Conshohocken, PA, 1997.
- [116] R. Lu, J. J. Koenderink, and A. M. L. Kappers, “Specularities on Surfaces with Tangential Hairs or Grooves,” *Computer Vision and Image Understanding*, vol. 1, no. 3, pp. 320–335, 2000.
- [117] R. Lu and J. J. Koenderink, “Optical properties (bidirectional reflection distribution functions) of velvet,” *Applied Optics*, vol. 37, pp. 5974–5984, 1998.
- [118] L.-Q. Yan, M. Hašan, W. Jakob, J. Lawrence, S. Marschner, and R. Ramamoorthi, “Rendering glints on high-resolution normal-mapped specular surfaces,” *ACM Trans. Graph.*, vol. 33, pp. 116:1–116:9, July 2014.
- [119] W. Jakob, M. Hašan, L.-Q. Yan, J. Lawrence, R. Ramamoorthi, and S. Marschner, “Discrete stochastic microfacet models,” *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2014)*, vol. 33, no. 4, 2014.
- [120] J. T. Kajiya and T. L. Kay, “Rendering fur with three dimensional textures,” *SIGGRAPH Comput. Graph.*, vol. 23, pp. 271–280, July 1989.
- [121] L.-Q. Yan, M. Hašan, S. Marschner, and R. Ramamoorthi, “Position-normal distributions for efficient rendering of specular microstructure,” *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2016)*, vol. 35, no. 4, 2016.

- [122] Z. Dong, B. Walter, S. Marschner, and D. P. Greenberg, "Predicting appearance from measured microgeometry of metal surfaces," *ACM Trans. Graph.*, vol. 35, no. 1, pp. 9:1–9:13, 2015.
- [123] F. Hunter, S. Biver, and P. Fuqua, *Light–science & Magic: An Introduction to Photographic Lighting*. Focal Press, Focal Press, 2007.
- [124] A. Bousseau, E. Chapoulie, R. Ramamoorthi, and M. Agrawala, "Optimizing environment maps for material depiction," *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)*, vol. 30, 07 2011.
- [125] K. Doerschner, H. Boyaci, and L. T. Maloney, "Estimating the glossiness transfer function induced by illumination change and testing its transitivity," *Journal of Vision*, vol. 10, no. 4, p. 8, 2010.
- [126] P. J. Marlow and B. L. Anderson, "Generative constraints on image cues for perceived gloss," *Journal of Vision*, vol. 13, no. 14, p. 2, 2013.
- [127] J. Gibson, "Ecological optics," *Vision Research*, vol. 1, no. 3, pp. 253 – 262, 1961.

