

Band-Sifting Decomposition for Image Based Material Editing

Ivaylo Boyadzhiev and Kavita Bala

Cornell University

and

Sylvain Paris

Adobe

and

Edward Adelson

MIT

Photographers often “prep” their subjects to achieve various effects; for example, toning down overly shiny skin, covering blotches, etc. Making such adjustments digitally, after a shoot, is possible but difficult without good tools and good skills. Making such adjustments to video footage is harder still. We describe and study a set of 2D image operations, based on multi-scale image analysis, that are easy and straightforward, and that can consistently modify perceived material properties. These operators first build a subband decomposition of the image and then selectively modify the coefficients within the subbands. We call this selection process *band sifting*.

We show that different siftings of the coefficients can be used to modify the appearance of properties such as gloss, smoothness, pigmentation, or weathering. The band-sifting operators have particularly striking effects when applied to faces; they can provide “knobs” to make a face look wetter or drier, younger or older, and with heavy or light variation in pigmentation. Through user studies, we identify a set of operators that yield consistent subjective effects for a variety of materials and scenes. We demonstrate that these operators are also useful for processing video sequences.

Categories and Subject Descriptors: I.3.7 [Computer Graphics]: Imaging & Video - Image Processing—*Material Editing*

General Terms: Image-based Material Editing, Multiscale Decomposition

■

1. INTRODUCTION

Photographers care a great deal about the surface appearance of objects they photograph; indeed, much of the craft of traditional photography involves controlling material appearance using physical techniques. Portrait photographers control the appearance of skin wrinkles by adjusting the lighting, apply makeup to hide variation in skin color (e.g., blemishes or mottling) and powder to make skin appear less shiny. In product photography, dulling spray is used to reduce specular highlights, while in food photography, where specularities may be desirable, a glycerine spray may be used to make the food look fresher or juicier.

Such adjustments can be performed digitally, after the photo was taken, rather than physically during the photo session, which greatly simplifies the process and enables more control on the result. However, altering the appearance of material properties such as wetness, gloss, wrinkles, or mottled coloration remains a tedious task that requires advanced skills that are beyond the reach of casual users. Further, in the case of video, laborious manual retouching is simply impractical; not only are there multiple frames, but it is difficult to get the effects to align and adjust across the sequence without introducing temporal artifacts.

An alternative route to manipulating material appearance is to build a fully renderable 3D description of the scene, and to change the physical parameters as needed. For example, multi-image capture with multiple cameras and light sources can create a highly detailed model of an object. However, this acquisition pipeline is not helpful to a casual or professional photographer working with a single image from an ordinary camera.

Our goal is to work with an ordinary photograph, and to allow the photographer to alter the appearance of a 3D surface without using a 3D representation. By using 2D image operations, we gain both speed and simplicity. We want our image-based technique to accept many kinds of source images as input, and to avoid the errors that can arise when attempting a full 3D scene analysis. At the same time, we must recognize that 2D operations can be limited and work best when there is a straightforward mapping between 3D surface properties and 2D image properties.

This paper explores a space of image operators that can be used to modify a variety of visual surface properties. The operators decompose an image by first applying a series of splitting operations based on frequency, amplitude, sign, and then *sifting* through these decompositions and recombining them to compose a new image. This sifting operation can create a range of visual effects that affect perceived material properties, for example, by changing perceived shininess/gloss, aging/weathering, and glow. Figure 1 shows an example of one such *sifting* procedure for two example images.

By selectively modifying coefficients in different subsets, one can achieve a variety of distinct image operators that we call *sifting operators*. While some of these are known and well studied; for instance, increasing high-frequency coefficients enhances detail in an image, the combination of the several criteria has not been explored so far. This paper seeks to fill in this gap and focuses on aspects related to material perception in particular. That is, our goal is to characterize which band-sifting operators generate physically plausible change that lead to perceptually consistent effects. For example, in Figure 1 we show that the same band-sifting operator makes both the human skin (row 1) and the orange surface (row 2) look more shiny and wet.

We explored the space of band-sifting operators and found that depending on the selected coefficients, operators modify properties such as the material shininess or its degree of weathering. On faces, the effects were particularly interesting with, for instance, variations in the appearance of oiliness, glow, wrinkles and pigmentation of the skin. Figure 2 shows examples of these various effects. As expected, we also observed that applying a modification too strongly yields unnatural looking results. This motivated two user studies. First, for each band-sifting operator, we characterized how strongly it can be applied before producing an unnatural look. And second, we studied

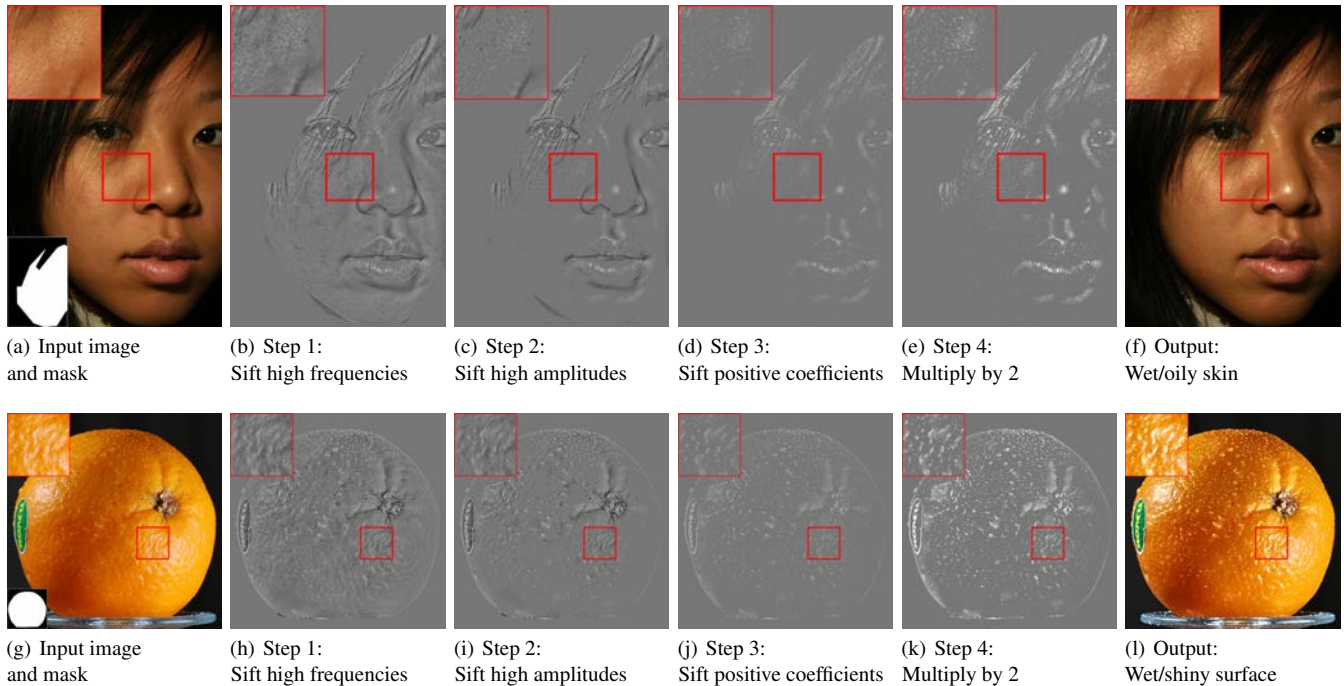


Fig. 1. Starting from a single input image and a mask (a), we selectively manipulate the subband coefficients of the luminosity channel by *sifting* them through a cascade of decisions based on the scale, amplitude and sign of the coefficients. Here we show one of these decision paths. First, we *sift* the high-spatial frequencies from the low-spatial frequencies (b). Then we *sift* the high amplitudes from the low amplitudes (c), and finally we *sift* the positive from the negative coefficients (d). Multiplying the *sifted* coefficients (e), adding them back, and reconstructing the image gives the skin a more oily or wet look (f). In the second row, we show that a similar perceptual effect is achieved on a non-face object, where the orange is given a more shiny or wet look. We found that *sifting* subband coefficients allows us to produce a variety of physically plausible effects that lead to perceptually consistent modifications across a variety of scenes.

how human observers describe the effect of each operator. This allowed us to isolate a subset of band-sifting operators that produce consistent effects across images. Finally, we also demonstrate the use of the band-sifting operators on videos. They are stable enough to achieve temporally coherent results without additional processing, they are fast enough to run at interactive rate, and they naturally “follow the scene content” without the need to estimate the optical flow explicitly. As an example, in a video of someone talking, we can add some glow onto them or make them look sweaty by simply applying our operators frame by frame, which is both simple and efficient.

Contributions. In this paper we introduce a new approach to image-based material editing based on multi-scale image decomposition and sifting.

- We propose and study a space of band-sifting operators that act along several criteria at the same time, based on scale, amplitude, and sign.
- We study the perceptual effects of our band-sifting operators; we validate their perceptual consistency through user studies; and we demonstrate their usefulness for both image and video post-process material editing.

1.1 Related Work

Image Decomposition. Splitting an image into components is a standard strategy to manipulate some properties independently of others. For instance, one can convert RGB colors into YIQ or CIE-Lab to edit luminance and chrominance independently. The coring operation used for denoising drives the low-amplitude coefficients

of a multi-scale decomposition towards zero without changing the high-amplitude coefficients [Donoho 1995; Simoncelli and Adelson 1996]. The classic Retinex algorithm by Land and McCann [1971] use a similar amplitude threshold in the gradient domain to separate the illumination from the reflectance of a scene. Mallic et al. [2006] describe a technique for separating specular and diffuse reflection components in images and videos. Durand and Dorsey [2002] separate large-scale variations from the small-scale ones for the purpose of HDR tone mapping, Bae et al. [2006] rely on a similar split for style transfer, and Farbman et al. [2008] for a variety of photo edits such as detail enhancement and local sharpening. Heeger and Bergen [1995] and Simoncelli and Portilla [2000] also use a multi-scale decomposition for their texture synthesis techniques. Motoyoshi et al. [2007] showed that manipulating the skewness of the coefficient distribution of the high-frequency bands affects the perceived gloss of materials.

Our work is related to this body of work since it splits images into components that are later modified separately. However, our purpose here is to look, more systematically, at the range of material-related manipulations that can be attained by doing modifications within the subband domain. Furthermore, material editing is different from other editing tasks because the visual information is distributed across space and subbands, and cannot be easily untangled.

Photo Editing. Many image operators exist to manipulate the level of texture in photographs, e.g., [Tomasi and Manduchi 1998; Farbman et al. 2008; Fattal 2009; He et al. 2010; Paris et al. 2011; Gastal and Oliveira 2011; 2012; Xu et al. 2011; Xu et al. 2012; Karacan et al. 2013]. These works focus on the signal processing




Fig. 2. Our band-sifting operators are particularly useful for manipulating material properties in human faces. (a) Original image, with detail inset at upper left, and mask inset at lower left. (b) We sift and then boost the high amplitude, positive coefficients in the high-spatial frequencies which gives the skin a more shiny or wet look. (c) We manipulate the positive low-spatial frequencies coefficients which gives the skin a soft glow. (d) We produce an aging effect by emphasizing blemishes and pores that are not noticeable in the input image. We achieve this by sifting and then boosting the low amplitude coefficients in the high-spatial frequencies. (e) We reverse the effect, i.e., reduce blemishes and pores, by decreasing the sifted coefficients from (d).

challenges, e.g., they improve the output or accelerate the computation. While these methods change the perceived properties of the materials, this aspect is not discussed in these papers. Fattal et al. [2007] implicitly use this effect to reveal details that would be hard to see otherwise. However, the perceptual aspects themselves are not studied. Our work is complementary to these articles and focuses on how image operators alter the observer’s material perception.

In parallel, the perceptual effect of some editing tools have been quantified. For instance, Mantiuk et al. studied contrast changes [2006] and whether image changes are visible [2011]. Trentacoste et al. studied the interaction between blur size and image resolution [2011] and showed that boosting the high frequencies of an image can be perceived as sharpening, halos, or countershading depending on the selected cutoff [2012]. In comparison, our work focuses on material perception.

Hybrid 2D/3D Material Editing. Khan et al. [2006] and Vergne et al. [2012] edit images to alter materials and their properties. The main difference with our work is their use of 3D data provided by users, as in [Vergne et al. 2012], or inferred 3D orientation, as in [Khan et al. 2006], from the images themselves — this allows Khan et al. to render new materials using standard 3D rendering, and Vergne et al. to warp images to convey shape and material properties. In comparison, our approach relies solely on the content of the input images and our operators are purely two-dimensional, thereby avoiding any sort of 3D reconstruction, that can be brittle on scenes with complex materials like those in which we are interested.

Weathering 3D Models. Several techniques exist to modify 3D models and simulate aging and weathering, e.g., [Golovinskiy et al. 2006; Mertens et al. 2006; Gloudu et al. 2012]. One of the effects we demonstrate makes people look older and objects more worn out, with the major difference being that we work purely in 2D. We also study several other effects beside aging, e.g., shininess and wetness.

Material Perception. A few techniques recognize materials depicted in photos, for instance, to differentiate plastic from wood, e.g., [Liu et al. 2010; Sharan et al. 2013; Bell et al. 2015]. Fleming et al. [2013] conduct user studies to explore the interactions between

material classification and judgments of material qualities, such as glossiness, roughness and hardness in the visual and semantic domains. In comparison, we are interested in altering the properties of a given material like its shininess or roughness. Researchers have also studied the interplay between physical sources, such as 3D geometry, surface reflectance and the light field in the perception of surface properties, such as gloss [Kim et al. 2011; Marlow et al. 2012; Marlow and Anderson 2013]. In our work we are interested in changing perceived surface properties, such as gloss, based on entirely image-based operations. More related to our approach, a few studies have shown a correlation between image statistics and the perception of properties such as translucency [Fleming and Bühlhoff 2005] and lightness [Motoyoshi et al. 2007; Sharan et al. 2008], and have proposed image filters that manipulate these statistics to alter the specific property that they study. Our work is inspired by these techniques and we build upon some of their findings. However, whereas these papers focus on a single effect, we explore a larger spectrum of effects, and systematically characterize how they affect material perception.

2. BAND-SIFTING OPERATORS

In this section, we describe the space of band-sifting operators that is at the core of our work. We strike a balance between two objectives. We define a space that is both expressive enough to include an interesting variety of effects and concise enough to allow for an exhaustive study.

2.1 Motivation for Band-sifting Operators

As we discussed in Section 1.1, many existing techniques can be interpreted as splitting an image according to a specific criterion like scale and amplitude, manipulating one of the generated components, usually with a simple operation like a multiplication, and recombining the result to form the final image. Our work extends this approach by decomposing images using several criteria at the same time. Figure 3 shows how we build our operators. We first split the original image into high and low-frequency subbands. We then separate the subbands into their high and low-amplitude parts. And we finally split the coefficients according to their sign, positive or negative. By placing multiplicative “control knobs” at specific

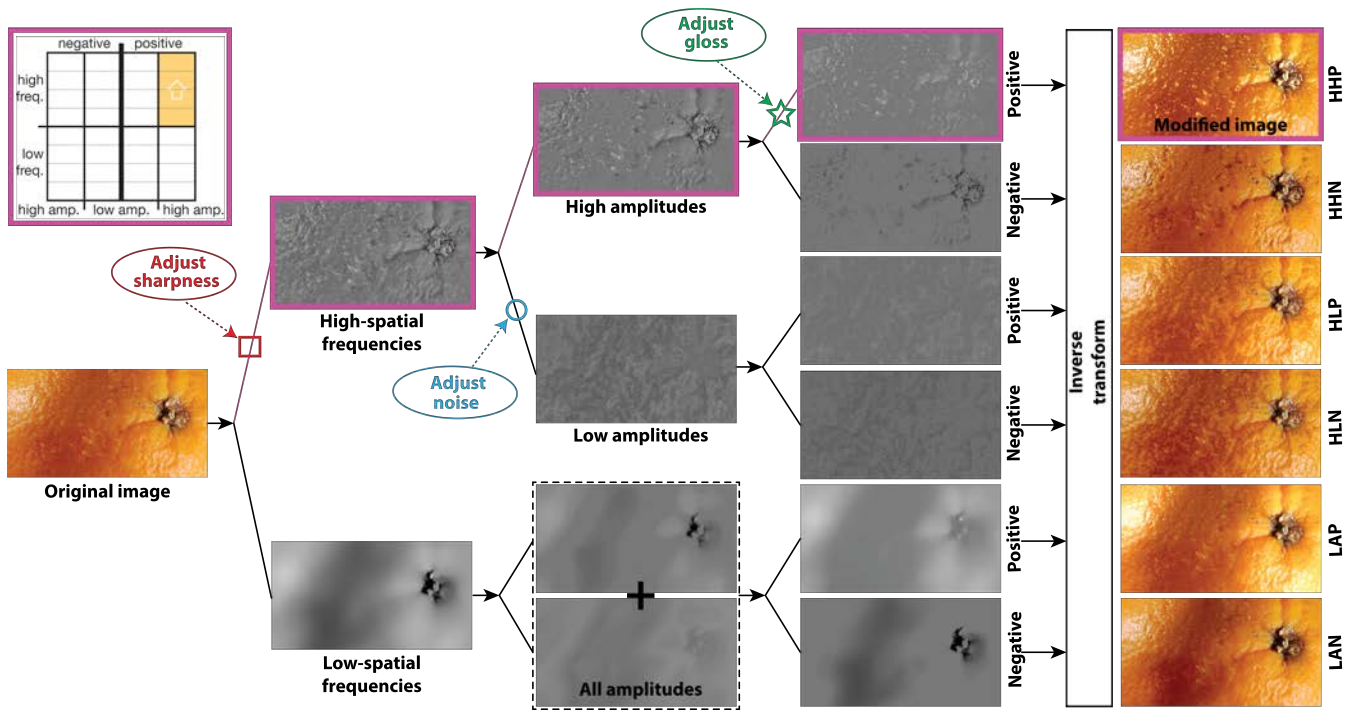


Fig. 3. Conceptual diagram of our band-sifting space. Given an image, we split it into high and low frequency subbands. These are then split into high and low amplitude parts. These are further split into positive and negative parts. For visualization purposes we show only two frequency splits, but in practice we create $\log_2(\min(\text{width}, \text{height}))$ frequency subbands and work on each one of them. Further, in order to make the size of the space more tractable, we “compress” the set of possible choices by looking at two categories of frequencies. We consider the high-to-mid frequencies as one category, which we refer to as “high spatial frequencies”, and we look at the mid-to-low frequencies as another category, which we call “low spatial frequencies”. Further, as shown in the diagram, we do not split the low-spatial frequencies category based on the amplitude of the coefficients, since numerically, *sifting* based on this criterion does not give much differentiation. However, the sign of the coefficients is still a useful *sifting* criterion along the low-spatial frequencies paths, e.g., it differentiates between broad-gloss and broad-shadow effects. With colored text and arrows we show how various operators can be mapped into paths in our space. With purple borders we show the path of *sifted* coefficients that was used to generate the orange result in Figure 1. In the upper left border we show an alternative, more compact diagram, of the same path. On the far right, next to each path, we show acronyms which we use to describe paths in our space. For example we use HHP for paths that manipulate (H)igh-spatial frequency, (H)igh amplitude, (P)ositive coefficients.

points in this flow diagram, one can modify sharpness (shown with a square), noise (shown with a circle), or gloss (shown with a star), similarly to previous work (§ 1.1). This illustration also suggests that there are many other ways to use our image decomposition, which raises several questions. What can one do by putting control knobs in other places? Are there more useful operators waiting to be found? What about yoked control knobs working on more than one component at a time? And what subband transforms are best?

Of course, there are any number of ways to increase the efficacy and complexity by adding in other techniques from image processing, computer vision and machine learning. However, our purpose here is to understand what is possible while staying within this scheme. Even with this restriction, there is plenty of territory to explore, and useful operators can serve as a starting point for later improvements.

Our goal here is to ask what can be done by simple manipulations of multiscale transforms. By staying close to the original image data, we maintain locality and avoid any propagation of artifacts. We also avoid the fragility that can occur, for example, when imposing specific physical models or elaborate priors. We accept arbitrary images as input, and in our experience, the image modifications look “natural” as long as they are not pushed too far.

2.2 Three Sifting Stages

We now describe the stages that we use to decompose images. We start by constructing a multiscale image decomposition and sift the subband coefficients based on three criteria: scale, amplitude, and sign.

Scale. Our design space follows a common trend and acts upon a multiscale image decomposition [Burt and Adelson 1983; Chen et al. 2007; Fattal et al. 2007; Farbmán et al. 2008; Fattal 2009; Hanika et al. 2011; Paris et al. 2011]. Such decomposition provides us with a set of subbands that can be thought of as an over-complete wavelet representation in the sense that each coefficient represents details at a given location and scale. This latter aspect is our first sifting stage: we allow our “sieve” to act selectively upon the large-scale or small-scale coefficients, or on all of them (i.e., both large and small scale). Intuitively, this separation differentiates between small elements, such as skin pores on a face, and bigger ones like large-scale shading and shadow variations.

Amplitude. Our second sifting criterion is the amplitude of the coefficients. We separately manipulate coefficients with a high or low amplitude, or both (high and low). This sieve separates low-




Fig. 4. In (a), we show a full decomposition into 8 subbands with each subband partitioned into 4 parts based on amplitude and sign. This gives 32 cells. Each cell can have one of three knob settings: boost, reduce, or leave the same. This gives 3^{32} potential configurations, which is hopelessly large. So we cluster the bands into two categories (high spatial frequencies and low spatial frequencies) giving 8 cells, as shown in (b). For each band-sifting operator, we show the knob setting of each cell with an arrow and a color, as shown in (c), where the low-amplitude positive high-frequency coefficients have been boosted, and the high-amplitude low-frequency negative coefficients have been reduced. The set of 3^8 configurations is still large; see text for further methods to reduce dimensionality.

contrast from high-contrast features. It is related to wavelet coring [Donoho 1995; Simoncelli and Adelson 1996] with the major difference that we keep and process the low-amplitude coefficients instead of discarding them.

Sign. The third sifting criterion differentiates coefficients based on their sign: positive or negative. Recent studies [Sharan et al. 2008] have shown that the skewness of the subband coefficient distributions, i.e., the asymmetry of the coefficient histograms, is correlated with the perception of lightness and gloss. From a numerical perspective, there is not a single well-defined way to alter skewness, that is, the same skewness value can be achieved with many different transforms. In our work, we modify coefficients based on their sign, which gives us a direct control on the distribution symmetry. This approach has an intuitive interpretation, the positive subband coefficients describe bright features like highlights, and the negative ones capture features like crevices, holes, and shadows in wrinkles.

2.3 Refining the Scale Sifting Criterion

The scale criterion raises two nontrivial issues: how many subbands to use and how to compute them. The rest of this section discusses these two issues.

Constructing the Scale Subbands. We started our study using a standard Laplacian pyramid that has the advantage of great simplicity. However, early in our investigations it became apparent that the Laplacian pyramid introduced artifacts at edges, which is a common problem when using linear filters on natural images. We therefore investigated pyramids based on edge-aware filters. We tried three such filters: the Bilateral Filter [Tomasi and Manduchi 1998], the Weighted Least-Squares filter [Farbman et al. 2008], and the Guided Filter [He et al. 2010]. All gave a significant reduction in edge artifacts (see the supplemental material for comparison). We chose to use the Guided Filter (used for all results in this paper), but other filters would presumably give similar results.

Number of Subbands. For our study, we used images at the resolution of typical monitors, e.g., the longer side set to 512 pixels. Using a factor of 2 in resolution between each subband, this yields 8 subbands. Then, the sign and amplitude sifting generates 4 components for each subband, and each of these components can either be boosted, reduced, or left unchanged. This would leave us

with 3^{32} operators to explore, which is impractical. We take a few steps to make this number more tractable. The first one is to group the subbands into two sets: the high and low-frequency subbands, which leaves us with 3^8 possible combinations. Figure 4 illustrates the decomposition we use. However, this number is still too large for the purposes of our exhaustive perceptual studies. In the next section, we further discuss how to reduce the space to a more manageable size, while ensuring that we still have a variety of distinct nontrivial effects to study.

2.4 Early Pruning

Even with the subband grouping described in the previous section, the space of possible band-sifting operators remain challenging to explore. In this section, we explain how we structured the space to make its exploration tractable.

Independent Criteria. First, we apply the sifting criteria independently of each other. For instance, for the sign, we choose between positive, negative, or both, and apply this choice to all the subbands. This gives us 3 sifting criteria (scale, amplitude and sign), with 3 options for each of them: {high (H), low (L), all (A)}, {high (H), low (L), all (A)}, and {positive (P), negative (N), all (A)} respectively. Once we have selected which coefficients to modify we can either boost (B) them or reduce (R) them. This defines $3 \times 3 \times 3 \times 2 = 54$ combinations in total. Figure 5 illustrates these 3 criteria.

Removing Redundancy. We explored and evaluated the space of operators in a set of pilot studies. We found that it could be reduced to a more useful set due to some redundancy in the effects achieved. Therefore, we applied the following pruning based on our observations.

- High frequencies tend to mask low frequencies and there is no visually significant difference between paths that sift only the high frequencies and those that sift all of them. Therefore, we do not include the latter in our study (pruned space: $2 \times 3 \times 3 \times 2 = 36$).
- All-amplitudes and high-amplitudes paths also produce visually similar results because they differ only due the low-amplitudes coefficients that are small by construction. We do not include the all-amplitudes paths in our study (pruned space: $2 \times 3 \times 2 \times 2 = 24$).
- Low-spatial frequency coefficients come from repetitively smoothing the input image and most of them are very small. Sifting these coefficients based on their amplitudes leave only very few significant values, and the corresponding modifications have almost no effect as can be seen in Figure 3. We avoid the high vs low-amplitudes paths and only include the all-amplitudes ones for the low-spatial frequency paths in the study. So for high spatial frequencies there are $(P|N|A) \times (H|L) \times (B|R) = 1 \times 3 \times 2 \times 2 = 12$ operators. And for low spatial frequencies there are $(P|N|A) \times (A) \times (B|R) = 1 \times 3 \times 1 \times 2 = 6$ operators. For a total of 18 possible operators.
- Also, reducing the low-amplitude coefficients does not have a noticeable effect since their value is already low. We do not include these paths in our study. This eliminates 3 possible operators, giving us a final total of 15 operators.

Another way of thinking of the space is: 9 operators where boost is applied, and 6 operators where reduce is applied.

2.5 Physical Observations

Appearance properties, such as luminance variations on surfaces, come from many physical sources, and a key reason why our approach works well is that in many cases these sources correspond to

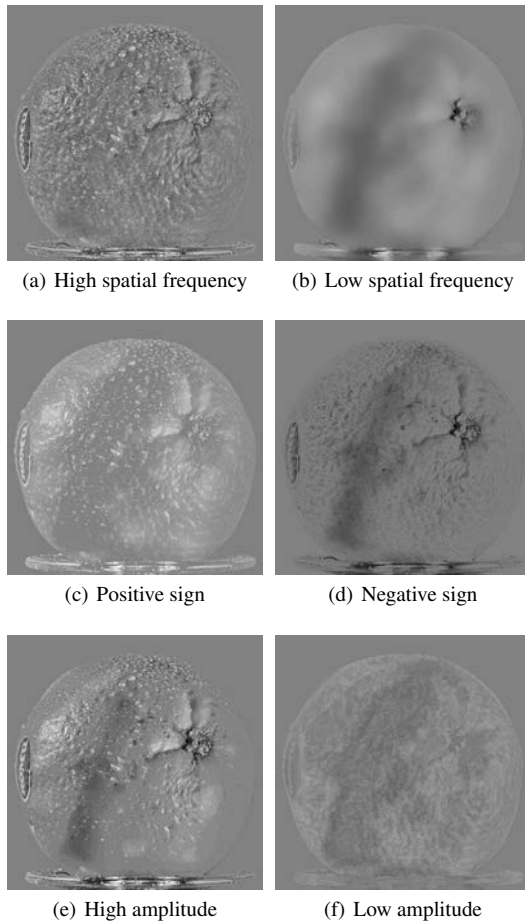


Fig. 5. Visualization of the subband coefficients *sifted* by each independent criterion. We visualize each possible selection by setting to 0 all the non-selected coefficients. We show the individual subbands for this input image in the supplemental material. High spatial frequencies (a) capture features like small-scale bumps and wrinkles, and low spatial frequencies (b) mostly represent large-scale shading variations. Positive (c) and negative (d) coefficients show highlights and shadows respectively. Finally, high-amplitude coefficients (e) represent specular highlights and deep shadows, while low-amplitude values (f) capture more subtle reflectance variations.

different bins of our subband decomposition. We discuss a few such examples below.

- Specularities typical of wet and glossy surfaces are bright and small, and mostly fall in the high spatial frequencies, with high amplitude and positive sign.
- Pits and grooves, including the wrinkles and pores of the skin, tend to be dark due to self-shadowing, and their magnitude is often medium or large. Because of this, they appear in the high spatial frequencies, high amplitude, negative coefficients.
- Variations in albedo, caused by dirt, stains, age, wear, or other degradations, tend to be low in amplitude compared to dark pits and bright highlights, and often shows up in the low-amplitude negative coefficients.

Such characteristics of physical objects are common and provides the basis for our approach. This also points at a limitation of our operators. If an object does not exhibit such properties, our operators

are not effective. For instance, we cannot make a perfectly smooth object look rough. To do so, one would need to hallucinate surface details that do not exist in the original image. We believe that this is an interesting direction for future work.

3. IMPLEMENTATION

In this section, we describe the actual implementation of the band-sifting operators that we described in the previous section. We provide detailed pseudo-code in Algorithm 1.

Our multi-scale decomposition is akin to that of Farbmán et al. [2008]. We repetitively process the input image with an edge-aware filter, doubling its spatial extent each time. This produces a series of images of increasing smoothness. Taking the difference between two such successive images gives us frequency bands, a.k.a. subbands, that contain details of a given size. Since we preserve edges in this construction, we do not downsample the subbands to prevent aliasing, i.e., each subband has the same resolution as the input. We use the Guided Filter with its default regularization parameter ($\sigma_r = 0.1^2$) [He et al. 2010] for edge-aware filtering.

In our prototype, the multi-scale decomposition and the subband-sifting procedure on the GPU were implemented using C++ and OpenCL. To accelerate the Guided Filter on the GPU, we implemented an efficient summed-area table algorithm [Harris et al. 2007] that we use as a building block for all the box filters, mean and standard deviation computations required by the Guided Filter approach. This allows us to achieve interactive frame rates (5–6 fps) for 1-megapixel videos, which is sufficient for preview purposes before running the full-resolution computation off-line.

For our study, we fix the long edge of the input image to 512 pixels, and compute 8 subbands. We split them into 4 low-frequency subbands and 4 high-frequency ones. For the amplitudes, we use the standard deviation of each subband as the threshold between the high and low categories. To avoid the artifacts that a hard threshold would introduce, we use a soft transition spanning $\pm 20\%$ around the standard deviation. For instance, if an operator multiplies by 2 the high-amplitude coefficients in a subband where the threshold is σ_t , the multiplication factor is 1 below $0.8\sigma_t$ and 2 above $1.2\sigma_t$, and smoothly varies in between. Finally, the increasing or decreasing of the selected coefficients is performed with a simple multiplication factor greater or lower than 1.

4. USER STUDIES AND RESULTS

In this section, we describe the user studies that we performed to characterize which operators produce effects that are natural and how they affect material appearance. We then present more results on still images and videos.

To understand the visual impact of our band-sifting operators, we conducted two user studies to validate their perceptual effects. The goal of the first study was to find which operators are natural, i.e., for a given image, what is the range of multiplication factors within which an operator produces a discernible and natural-looking change? This study tells us how much we can boost or reduce an effect before it starts to look unnatural on a certain image.

The second study asks users to describe the visual change that operators produce. This task shares some similarities with a recent line of work in computer vision, where the goal is to describe images through high-level *attributes* [Patterson et al. 2014]. In our work, we are interested in assigning attributes related to the perceptual changes produced by the band-sifting operators. We use 16 categories of words describing various material-specific properties. We designed the set of words in a pilot study between the authors by

looking at the perceptual effects produced by the operators on tens of examples. Figure 8 lists these words. Some of those categories of words describe low-level features, such as “wrinkled, pitted, bumpy” and other categories describe high-level properties, such as “young, new, fresh”. Then, in our study with casual users, participants were shown the original image and the modified image, and were asked to pick all categories of words that apply.

4.1 Study 1: Natural vs. Unnatural

The goal of this study is to find whether there is a reasonable range of multiplication factors where the operators produce natural looking results. We test a few multiplication factors and run a study to find the threshold between natural and unnatural.

Given a pair of an input image I and a band-sifting operator F , we seek to sample a few versions of the operator, $F(I, m_1), F(I, m_2), \dots, F(I, m_s)$, acting on the original image with different multiplication factors m_1, m_2, \dots, m_s . Our early experiments showed that using the same m factors across operators perform poorly; the same value can produce a strong effect with one operator and a weak one with another. Instead, we define the factor as $m_0 = 1$ and $m_{i+1} = \arg \min_m \|F(I, m_i) - F(I, m)\| > 1$ using the CIE-Lab L_2 norm. We use binary search to efficiently find m_{i+1} . This procedure generates samples regularly spaced in the CIE Lab color space akin to Ngan et al. [2006], which approximates a perceptually uniform sampling. We observed that for the effects that boost the coefficients, 5 iterations of the above procedure were usually enough to produce too strong results. For the effects that reduce the coefficients, 2 iterations were usually enough to make them close to 0.

With the above sampling procedure, we produced 5 images of different strength for each of the 9 operators that increase the coefficients plus 2 images for the 6 operators that decrease the coefficients, for a total of $5 \times 9 + 2 \times 6 = 57$ variations per image. We used 21 images, 11 faces of various genders and races and 10 non-face objects with uniform materials, e.g., metal, leather, ceramic, and fruits. Users were shown a single modified image at a time and were asked whether it looked natural to them. We provide a snapshot of this task as well as the full set of images in the supplemental material.

Every user was shown 15 sets of 21 images. Each set was made of each of the 21 test images modified using a randomly picked setting. Users saw the same scene 15 times. Occurrences of any scene were separated by 20 other images to limit the effect of users’ getting trained by the previous viewing of that scene. Users were asked to base their decision only on the current image, and they had no reference original image. This study had a total of 47 participants and on average we got 7.5 votes per setting, since we assigned them uniformly across participants.

Figure 6 summarizes the results of the study and confirms our initial observation that our band-sifting operators can produce non-trivial natural looking variations even for familiar objects, such as human faces. We also show separate statistics for face and non-face objects, which reveals some interesting trends. For example, in essentially all cases where we boost the signal, images of non-face objects withstand larger modifications than images of face objects. Another interesting observation is about reducing the high-frequency high-amplitude coefficients on faces. If we manipulate the positive or negative coefficients separately, we produce larger modifications compared to reducing them both at the same time. This happens because modifying the positive and negative coefficients at the same time leads to dampening of all high-frequency features, such as dark skin pores and bright skin gloss, which quickly produces unrealistic




Fig. 7. Our first user study characterized how much we can reduce (a) or boost (c) a set of coefficients while maintaining a natural look. Increasing the coefficients past this point eventually generates unnatural images (d). The insets explain which coefficients are affected, see Table I.

smoothing of the skin, compared to modifying skin pores independent of skin gloss. In the supplemental material, we show an example image that demonstrates this effect.

Figure 7 shows examples of natural and unnatural adjustments using our band-sifting operator that manipulates the high-frequency high-amplitude negative coefficients. We also mark the data points in Figure 6a that correspond to this operator. For example, as we pass beyond the realistic threshold of the boost operator, e.g., 3 steps, it starts to produce unrealistic looking images as we show in Figure 7d.

Statistical Significance. To confirm the statistical significance of our results, we assumed that each user has a naturalness threshold. We compared two hypotheses: purely random thresholds, i.e., uniformly distributed over the tested range, and Gaussian distributed thresholds centered on the value that we reported in Figure 6a (we used a unit variance for simplicity). We compared the probabilities of obtaining the users’ answers under these two assumptions. As shown in Figure 6b, for 76% of the face image and 97% of the non-face images, the results of our study are more than $300\times$ more probable under the Gaussian hypothesis than under the uniform one, which confirms the hypothesis of a consistent threshold across images and users.

4.2 Study 2: Name the Effects

The goal of the second study is to determine the perceptual effects associated with the band-sifting operators and to evaluate their consistency across different users and different scenes. Users were shown pairs of images where image A was the original input image I , and image B was a modified version, $F(I, m)$. We seek a parameter value m that produces a visible and natural effect, which we achieved with the quasi-median of the votes in the first study, i.e., the multiplication factor with an equal number of natural and unnatural votes above and below. We showed users the 16 groups of keywords, and for each group, asked them to choose between 3 options for the direction of the perceptual change: “Less”, “More” or “N/A”. 20 users participated in this study and half of them had not taken part in Study 1. On average, we got 60 responses per operator (30 for faces and 30 for non-face objects) for a total of 1100 responses. Figure 8 summarizes the responses for one of our operators for both face and non-face objects. We show similar plots for all the 15 operators in the supplemental material.

Statistical Significance. We tested our results against the null hypothesis that the choice between the 3 options is uniformly ran-




Fig. 6. The plot on the left reports the mean and standard deviations of the votes on the natural-vs-unnatural study (a). We use the notation introduced in Section 2.4 on the horizontal axis. We also indicate the median vote that we use later in the second user study. These votes confirm that the threshold between natural and unnatural settings is statistically significant since the uniform-distribution hypothesis is much less probable than the Gaussian-distribution hypothesis in most cases (b). With yellow and purple rectangles we indicate the boost and reduce statistics for one of the operators, for which we show a concrete qualitative example in Figure 7. See the text for details.




Fig. 8. Our second user study determined the perceptual effects associated with our band-sifting operators. In these plots we show results for one of our more consistent effects: boost high amplitude, positive-valued, high-spatial frequencies. The red bars show the percentage of votes for the “Less” option, the blue bars show the percentage of votes for the “More” option, and the difference to 100%, which we do not show on the plot, is the percentage of votes for the “N/A” option. The green lines indicate the probability of chance, i.e., where the results have been generated by picking between the three options, {“Less”, “More” or “N/A”}, uniformly at random. The blue lines indicate the 95% confidence interval, i.e., results above this threshold, $\approx 50\%$, are statistically significant with high probability. The red lines, at $\approx 47\%$, indicate the 90% confidence interval. The majority of participants agreed that this band-sifting operator tends to make human faces more wet, oily or sweaty, whereas, for objects, the band-sifting operator tends to make them look more shiny, glossy or metallic. See the text for details about the test of significance.

dom. This hypothesis corresponds to a standard multinomial distribution with a 33% mean. For 30 votes, the standard deviation is 8.6%, and using a 95% confidence interval, we can rule out the null hypothesis for results below 16% and above 50%. For a 90% confidence, the interval is [19%, 47%]. We give the detailed derivation of these numbers in the supplemental material. We show the 50% and 47% thresholds in Figure 8 and use them to report the results in Table I.

Consistent Effects. We found 7 operators that produce consistent and perceptually discriminative effects: boost/reduce shininess, boost/reduce roughness, boost weathering patterns and boost/reduce glow. Table I summarizes this finding. The number of word sets above the significance threshold ranges from 1 to 4. In general, operators have a more consistent effect on faces. We hypothesize that the diversity of scenes and materials present in the non-face images makes it “harder” for an operator to be consistent. In comparison, the only material in face images is skin and although human observers recognize subtle differences, these are not as dramatic as those between bronze and potatoes for instance.

From a photo editing perspective, these band-sifting operators cover several common tasks on objects such as reducing or increasing weathering, smoothness, and shininess. For faces, they provide a simple and effective means for attenuating blemishes and wrinkles, controlling the dryness of the skin, and adding a photographic glow typically observed in studio portraits.

4.3 Image Results

We now demonstrate the effects of band sifting on a range of static scenes and qualitatively discuss the results.

In Figure 9, we show that our purely image-based band-sifting operators can produce results visually similar to what can be achieved with a 3D model rendered with a physically inspired BRDF model. We used a photorealistic scanned 3D model of a face, where we control perceptual parameters related to shininess by changing physical parameters of the underlying rendering model [Ward 1992]. We rendered the face with two different values of the α parameter that controls the spread of the specular lobe. Smaller values of α increase the sharpness of the reflected image and make the object look shinier. The question is whether we can get a similar effect using just 2D image manipulations. We show that our band-sifting operator that boosts the positive high-amplitude high spatial frequencies produces a perceptually similar change in shininess.

In Figure 10(b) we show a failure case of one of our band-sifting operators. When the visual cues are not present in the original input image or they are not well isolated by our *sifting* criteria, our band-sifting operators fail to convey consistent perceptual effects (b).

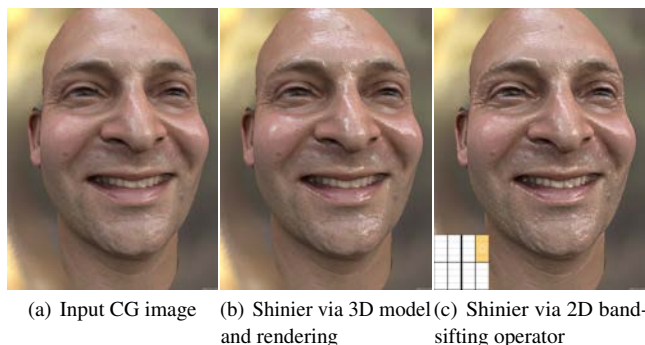


Fig. 9. Case study on a photorealistic scanned CG model of a face, courtesy of [von der Pahlen et al. 2014]. In (a) and (b) we rendered the face under natural lighting conditions [Debevec and Malik 1997], using the isotropic Ward BRDF model with two different values of the parameter α , that controls the spread of the specular lobe. In (c) we show that our band-sifting “wet/oily/shiny” operator, when applied to image (a), can produce a perceptually similar change in shininess.

Even though one operator may not work well on an image, we found that usually others might still be useful. For example, we can reduce skin defects (c) and blemishes (d) or add a smooth skin glow (e). Furthermore, the independent band-sifting operators that we studied can be combined to achieve even more advanced material editing effects. In (f) we combine the previous three operators to achieve the combined perceptual effect, a younger looking face with a nice skin glow, which is often seen in professional magazines.

Figure 11 illustrates the diversity of effects that can be achieved with band-sifting operators. For brevity’s sake, we use the notation previously introduced where the amplitude is selected in {H, L, A}, the frequency in {H, L, A}, and the sign in {P, N, A}.

Gargoyle. Boosting the HHP coefficients enhances the gloss, and also brings out whitish “distress” marks, which gives an overall shinier look. Boosting the AHN coefficients produces a patina with dark mottling.

Grapes. We show a combination of two of operators: boosting the ALP coefficients while reducing the AHN coefficients gives the grapes a luminous glow. We also achieve a weathering effect by boosting the LHN coefficients to bring out the patterning on the grape skins.

Onion. Reducing the HHA coefficients removes texture details, while retaining the smooth shiny appearance of the onion; whereas boosting the same coefficients reveals the mottled coloration of the onion skin.

Sweet potatoes. Boosting the LHN coefficients reveals dark blotchy patches while boosting the HHN ones reveals sharp dark spots.

Orange. Boosting the HHP coefficients emphasizes the highlights and makes the orange look shinier. Alternatively, we can emphasize pores and dark spots by boosting the HHN coefficients.

4.4 Video Results

We now demonstrate the effects of the band-sifting operators on video sequences. Editing videos consistently is particularly challenging and typically requires many hours of painstaking manual editing. In Figures 12 and 13, we show example input frames from the video sequences and the results on two different frames. We show results on three video categories where post-process material editing would be a desired tool: (1) closeup views of people giving interviews; (2) 360 spins for product photography; (3) static objects under dynamic

lighting. In each case, we applied our operators frame by frame. As can be seen in the supplemental video, the results are artifact-free and temporally consistent, which demonstrate the robustness and stability of our band-sifting operators.

Interview A. First, we boost the HHP coefficients to give the skin a more wet/oily look. Then we demonstrate our skin glow effect, which is a common appearance professional photographers aim to achieve through a combination of lighting and facial cosmetics. We achieve that effect through entirely image-based manipulations, by boosting the ALP coefficients. Finally, we emphasize blemishes and spots by boosting the LHA coefficients. We did not use a detailed mask around the face in this case to show that the operators could be directly applied and used in settings like this one.

Interview B. We demonstrate the perceptual consistency of our wet/oily and skin glow effects used in the previous example by applying them onto a different subject. Then, we reduce blemishes and spots to produce a cleaner looking face. To localize the effects

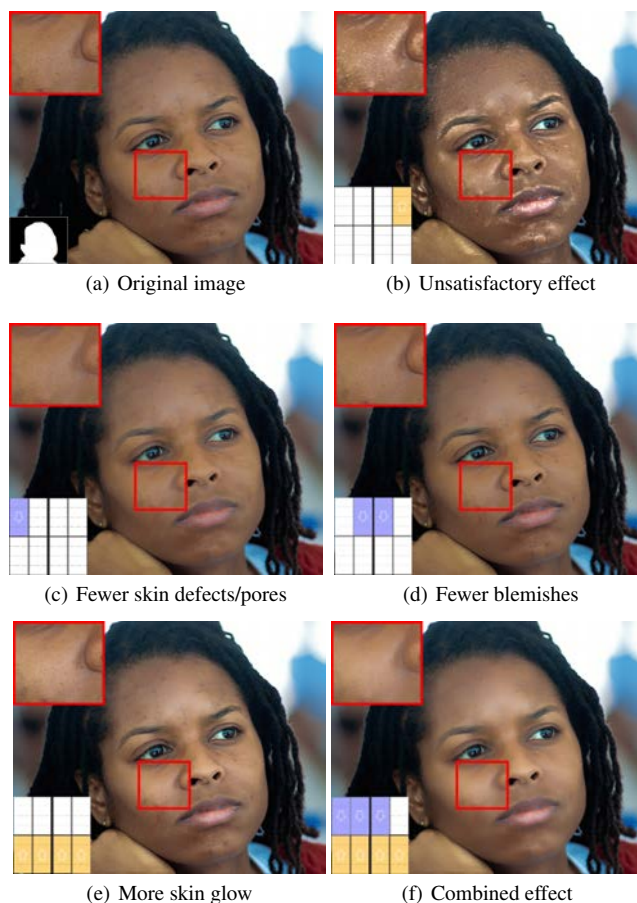


Fig. 10. Starting from an input image of a very dry and matte face (a), our “wet/oily/shiny” band-sifting operator fails to produce a plausible looking effect (b). Even though one operator may fail to produce a satisfactory effect on some image, others might still work. In (c) we reduce skin defects and pores. In (d) we reduce skin blemishes and pigmentation. In (e) we add a smooth skin glow. Furthermore, simple combinations of band-sifting operators can be used to achieve advanced material editing tasks. In (f) we show the combined effect, which achieves a younger look with a nice skin glow, an advanced effect often seen in professional magazines.

on the face only, we created and tracked a detailed facial mask, using the Roto Brush tool in Premiere Pro.

Leather shoes. Boosting the HHA coefficients emphasizes the highlights and some scratches, which gives an overall shinier looking leather. Then, we give the leather a smoother and more polished look by boosting the ALP coefficients. Finally, we achieve a weathering effect by boosting the LHA to bring out the patterning on the leather.

Grapes. Boosting the HHP coefficients emphasizes the highlights and makes the grapes look shinier and wetter. Then, we reduce the same coefficients to produce a more diffuse look. Finally, we bring out the weathering patterns, which also make the grapes look a bit more dirty.

Helmet. We demonstrate that the operators produce a consistent look under dynamic lighting conditions. Boosting the HHP coefficients emphasizes the highlights and some scratches, which makes the helmet appear shinier. Boosting the ALP brings out the broad gloss, which gives a smoother, less rough looking metal. Finally, boosting the LHA coefficients brings out the weathering patterns, which makes the metal look like it has more patina.

5. DISCUSSION AND LIMITATIONS

Our purpose in this paper is to devise 2D image operators, which can make visually plausible modifications of surface properties. Such operators are simple to implement and can be applied to arbitrary images. However, they are only useful for properties that are manifested simply within the distributions of subband coefficients. Prior research indicates that such descriptions are useful for certain tasks involving natural images (e.g., in denoising [Donoho 1995; Simoncelli and Adelson 1996], and in texture analysis [Heeger and Bergen 1995; Portilla and Simoncelli 2000]). Here we have tested the utility of similar representations in modifying material appearance.

We have identified several kinds of material appearance that are commonly associated with certain subband properties, and that can often be manipulated, when the signal is already present in the image. In our experience the observed perceptual changes are stable for materials where the albedo does not vary too rapidly, such as skin texture, leather texture, metals and the texture of some fruits. For example, specularities from fine-scale features show up in positive-valued, high-amplitude, high-spatial frequencies. Boosting them leads to an enhanced “glistening” appearance, which may be interpreted as oiliness or wetness on skin. Specularities from smooth, large-scale features show up in positive-valued, medium-spatial frequency coefficients. Boosting them leads to a smooth shine or sense of skin glow. Small spots, pits, and wrinkles, typically manifest themselves as small dark features that show up in the high-amplitude negative coefficients of high-spatial frequencies. Boosting them often emphasizes the visibility of fine-scale texture (both fine-scale geometry and fine-scale albedo). These features are often associated with the aging of human skin, or the weathering of natural surfaces. In addition, the low-amplitude coefficients of the high-spatial frequencies (both negative and positive) are often associated with splotchy or mottled pigmentation. These can also enhance the sense of age, weathering, and discoloration. In the past, those coefficients were typically artifacts of the imaging chain, i.e., sensor noise or image coding artifacts. However, modern digital cameras offer clean images, and the low amplitude signals tell us about the scene, not about the camera. The perceptual effects of our operators are less effective in the presence of noise or lossy compression artifacts. We show an example of this in the supplemental material.

We find it remarkable that these effects tend to look natural and realistic, rather than being the result of some artificial manipulation. The realism presumably results from the fact that the band-sifting

operators are not inventing any information that is not already there; they are just emphasizing or de-emphasizing visual patterns that are already part of the image. As we have discussed earlier, when the visual cues are not present in the original input image, or they are not well isolated by our *sifting* criteria, our band-sifting operators may fail to convey consistent perceptual effects. High frequency albedo can also lead to unsatisfactory effects for band-sifting operators that manipulate those frequencies. In Figure 14, starting from an input image that has high frequency texture (a), our band-sifting operator fails to produce a satisfactory “shiny” effect (b). Although one operator may fail on a certain image, other band-sifting operators might still be useful. In (c) we manipulate material properties related to smoothness, by boosting the low-spatial frequency coefficients, which gives the apple a more polished look.

6. CONCLUSION AND FUTURE WORK

We present band-sifting operators, and demonstrate their use in manipulating surface appearance. The band-sifting operators selectively alter coefficients within a subband decomposition, where the selection is based on spatial scale, sign, and amplitude. We explored a reasonable subspace of such operators and demonstrated their ability to modify a variety of surface properties in natural scenes. We use only 2D operations, but they can give the visual impression of acting on the materials in 3D scenes. We found some operators that were useful in controlling smoothness or gloss, which could alter the appearance of wetness, shininess, or degree of polish. Other operators altered the apparent pigmentation, roughness, or weathering of surfaces. We performed user studies and determined that there are certain operators that lead to consistent perceptual effects across various images and across multiple observers. Image class does matter: with images of faces, subjects reported that that the filters would change face-specific properties such as oiliness, blemishes, wrinkles, skin age, and skin health. Given the importance of perceptual surface qualities, we expect that these band-sifting operators can offer an important tool for photography. Our band-sifting operators can also be used with video sequences. The visual effects tend to be consistent across a sequence, making it possible, for example, to change the apparent shininess of an actor’s skin. In the future, further exploration of band-sifting, e.g., by a finer subdivision of our proposed space or by introducing new sifting criteria, could open the door for many more operators for image-based material editing. Identifying the conditions when an operator would succeed or fail to produce a desired effect, based on the content of an arbitrary given image, is an interesting avenue for future work.




Fig. 14. Working on an object that has high frequency albedo (a), our “shiny/glossy/metallic” band-sifting operator is less effective in conveying the more shiny look since it picks mainly on the albedo (b). Although one band-sifting operator may fail, others might still be useful. In (c) we manipulate material properties related to smoothness, which makes the apple look more polished.

Table I. Recap of the band-sifting operators with a consistent effect

	coefficient selection	associated properties	input	output
O1.		<p>Objects (strength: 4) more shiny/glossy/metallic (83%) more bright/glowing (64%)</p> <p>Faces (strength: 2) more wet/oily/sweaty (76%) more shiny/glossy/metallic (67%) less dry/dull/matte (48%)</p>		
O2.		<p>Objects (strength: 2) more dry/dull/matte (62%) less shiny/glossy/metallic (54%)</p> <p>Faces (strength: 2) more dry/dull/matte (62%) less shiny/glossy/metallic (55%)</p>		
O3.		<p>Objects (strength: 3.5) more sharp/crisp (48%)</p> <p>Faces (strength: 2) more hard/dark shadows (60%)</p>		
O4.		<p>Objects (strength: 2)</p> <p>Faces (strength: 2) more smooth/polished (56%) less wrinkled/pitted/bumpy (53%) less hard/dark shadows (53%) more young/new/fresh (47%)</p>		
O5.		<p>Objects (strength: 3.5) more sharp/crisp (57%) more old/used/worn (49%) more blemished/stained (47%)</p> <p>Faces (strength: 2) more blemished/stained (62%) more old/used/worn (55%) more wrinkled/pitted/bumpy (55%)</p>		
O6.		<p>Objects (strength: 5) more shiny/glossy/metallic (79%) more bright/glowing (56%)</p> <p>Faces (strength: 4) more bright/glowing (70%) more shiny/glossy/metallic (68%)</p>		
O7.		<p>Objects (strength: 2) more dry/dull/matte (58%) less shiny/glossy/metallic (52%)</p> <p>Faces (strength: 2) more dry/dull/matte (64%) less shiny/glossy/metallic (57%) less bright/glowing (57%) less sharp/crisp (54%)</p>		

Band-sifting operator #1 makes objects look shinier and more metallic, and faces wetter and oilier. #2 has the opposite effect. #3 reveals small defects in objects and makes face pores look deeper. #4's effect on objects is not consistent but it renders people younger. #5 reveals object details and make them look stained. On faces, it emphasizes blemishes and makes people look older. #6 renders objects more metallic, and gives a shiny glow to people. #7 has the opposite effect. The strength numbers in the rightmost column are the result of our first user study to estimate how much we can push an effect before it starts looking unnatural. For our second user study, we report the statistically significant perceptual votes. Votes larger than 50% are statistically significant at the 95% confidence interval, i.e., with a very high probability the votes are not a result of chance. In gray we show votes that are statistically significant at the 90% confidence interval.




Fig. 11. Showing a variety of effects produced with our band-sifting operators. In row (1) we make the gargoyle look more glossy (1,b) or more weathered, by emphasizing the patina (1,c). The grapes can be given a more glowing (2,b) or more dirty look (2,c) by emphasizing the patterning on the skin. The skin of the onion can be given a more clean (3,b) or a more mottled look (3,c). In row (4) we show that by treating the low and high amplitude coefficients separately we can get very different perceptual effects, weathering patterns (4,b) vs. surface roughness (4,c). Finally, we can make the orange look more shiny and wet (5,b) by manipulating the positive coefficients independently from the negative. In (5,c) we show that pores and dark spots are well captured in the negative coefficients which can be used to manipulate those material properties independently from highlights and gloss.



Fig. 12. Results on videos of faces, downloaded from the Internet. Our band-sifting operators can be used to efficiently post-process material appearance in videos without introducing temporal artifacts, see the supplemental video. For example, we can make the actor’s skin look more oily, column 2, or add more skin glow, column 3. In column 4 we control wrinkles and blemishes by manipulating coefficients in the corresponding combination of band-sifting paths. We boost the coefficients in rows 1 and 2 which gives the face a more aged look, whereas in rows 3 and 4 we reduce them to render a more clean and young looking face. The simplicity of our model allows all this to be done interactively by manipulating a few sliders, without having to model the effects pixel-by-pixel on every frame. We show the abbreviations of the used operators in brackets.

APPENDIX

In this appendix we provide a detailed pseudo-code implementation of our band-sifting operators. In practice all major steps can be implemented efficiently on a GPU using a fast summed-area table algorithm [Harris et al. 2007].

REFERENCES

BAE, S., PARIS, S., AND DURAND, F. 2006. Two-scale tone management for photographic look. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 25, 3, 637 – 645.

BELL, S., UPCHURCH, P., SNAVELY, N., AND BALA, K. 2015. Material recognition in the wild with the materials in context database. *Computer ACM Transactions on Graphics*, Vol. 32, No. 5, Article XXX, Publication date: 09 2015.

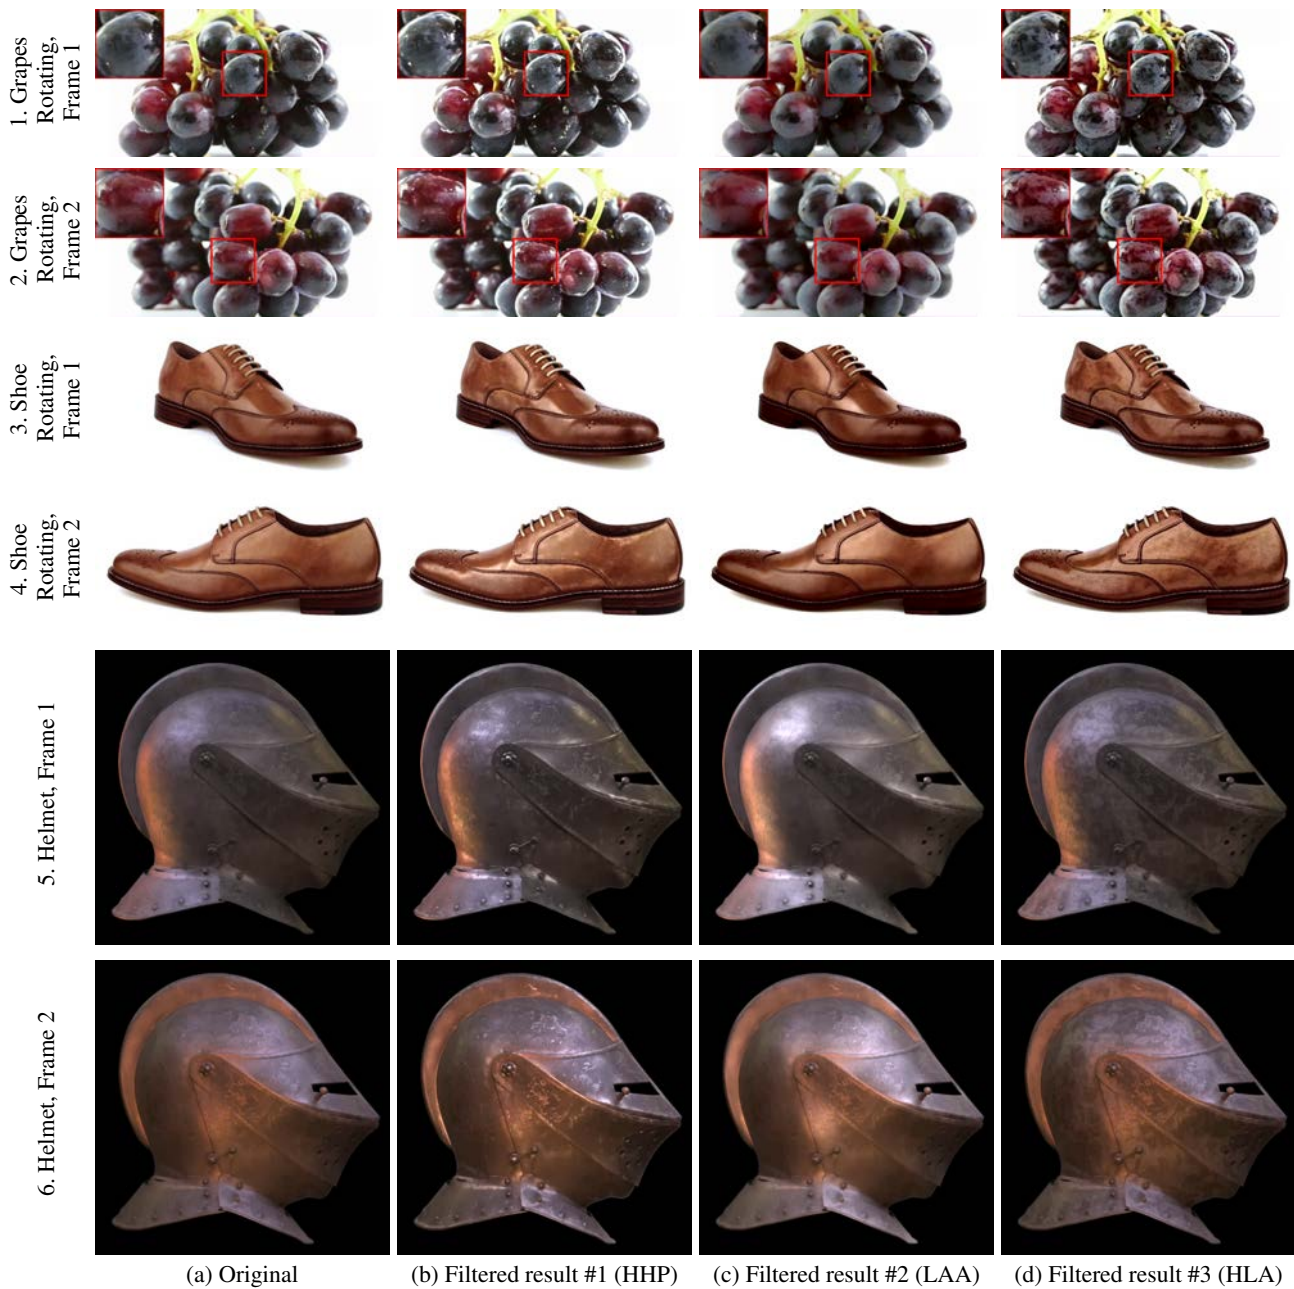


Fig. 13. We evaluate our band-sifting operators on videos of 360 spins for product photography, rows 1 through 4, and videos of static objects under dynamic lighting, rows 5 and 6. The helmet scene is courtesy of [Debevec et al. 2000]. In column (a) we show example input frames from the video sequences, and the next columns show some of the effects our band-sifting operators can achieve. In column (b) we demonstrate that our “shiny/glossy/metallic” effect, which boosts the high amplitude positive-valued, high-spatial frequencies, produces a perceptually consistent effect across different scenes. In column (c), rows 1 and 2 we reduce the corresponding coefficients, which gives the grapes a more diffuse look. In the rest of column (c), rows 3 through 6, we show our “bright/glowing” effect, which boosts the positive-valued, low-spatial frequency coefficients. This gives the leather shoes a more smooth and polished appearance, rows 3 and 4, and it makes the metal helmet look more smooth and less rough, rows 5 and 6. In column (d) we show our “old/used/worn” effect, which boosts the low amplitude, high-spatial frequencies. In agreement with our user study #2, this effect produces a persistent perceptual effect of aging or damage by bringing out weathering patterns such as spots and dust on fruits, rows 1 and 2, leather stains, rows 3 and 4, and patina on metals, rows 5 and 6.

Vision and Pattern Recognition (CVPR).

BURT, P. J. AND ADELSON, E. H. 1983. The Laplacian pyramid as a compact image code. *IEEE Transactions on Communication* 31, 4, 532–540.

CHEN, J., PARIS, S., AND DURAND, F. 2007. Real-time edge-aware image processing with the bilateral grid. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 26, 3.

DEBEVEC, P., HAWKINS, T., TCHOU, C., DUIKER, H.-P., SAROKIN, W.,

AND SAGAR, M. 2000. Acquiring the reflectance field of a human face. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '00. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 145–156.

DEBEVEC, P. E. AND MALIK, J. 1997. Recovering high dynamic range radiance maps from photographs. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '97. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 369–378.

DONOHO, D. 1995. De-noising by soft-thresholding. *IEEE Transactions on Information Theory*.

DURAND, F. AND DORSEY, J. 2002. Fast bilateral filtering for the display of high-dynamic-range images. *ACM Transactions on Graphics (Proc. SIGGRAPH) 21*, 3.

FARBMAN, Z., FATTAL, R., LISCHINSKI, D., AND SZELISKI, R. 2008. Edge-preserving decompositions for multi-scale tone and detail manipulation. *ACM Transactions on Graphics (Proc. SIGGRAPH) 27*, 3.

FATTAL, R. 2009. Edge-avoiding wavelets and their applications. *ACM Transactions on Graphics (Proc. SIGGRAPH) 28*, 3.

Algorithm 1 BANDSIFTINGOPERATOR

input: image I , multiplication factor λ , $sign \in \{\text{pos, neg, all}\}$, $freq \in \{\text{high, low, all}\}$, $amp \in \{\text{high, low, all}\}$
output: image O

- 1: $L \leftarrow \text{luminance}(I)$ // process only the L channel of Lab
- 2: $\{S_\ell\} \leftarrow \text{DECOMPOSE}(\log(L+\epsilon))$ // multiscale decomposition of log luminance (Alg. 2), use small ϵ to avoid 0 values
- 3: $n \leftarrow \#\{S_\ell\}$ // define n , the number of subbands
- 4: **for all levels** $\ell \in [1; n]$ **do**
- 5: $R \leftarrow S_\ell$ // keep a copy of the subband as reference
- 6: **for all coefficients** $c \in S_\ell$ **do**
- 7: // check if c selected by its sign and frequency (Alg. 3)
- 8: **if** SIGNANDFREQUENCYSELECTED(c , $sign$, $freq$) **then**
- 9: // check if c is selected by its amplitude
- 10: **if** $amp = \text{all}$ **then**
- 11: $c \leftarrow \lambda c$ // if all amplitudes, directly apply multiplier
- 12: **else**
- 13: // else smooth transition between high and low amplitudes
- 14: $\sigma = \text{STDDEV}(S_\ell)$
- 15: $\alpha \leftarrow \text{SMOOTHSTEP}(0.8\sigma, 1.2\sigma, |c|)$ // returns 0 if $|c| < 0.8\sigma$, 1 if $> 1.2\sigma$; transitions in between
- 16: // orient transition depending on amplitude selection
- 17: **if** $amp = \text{high}$ **then**
- 18: $c \leftarrow c \times (1 + \alpha(\lambda - 1))$
- 19: **else if** $amp = \text{low}$ **then**
- 20: $c \leftarrow c \times (1 + (1 - \alpha)(\lambda - 1))$
- 21: **end if**
- 22: **end if**
- 23: **end if**
- 24: **end for**
- 25: // smooth applied gain map
- 26: $Gain \leftarrow S_\ell/R$
- 27: $Gain \leftarrow Gain \otimes G(2^\ell)$
- 28: $S_\ell \leftarrow R \times Gain$
- 29: **end for**
- 30: $O = \exp(\sum_1^n S_\ell) - \epsilon$ // sum subbands to get output

Algorithm 2 DECOMPOSE

input: single-channel image C
output: multiscale stack $\{S_\ell\}$

- 1: $n \leftarrow \log_2(\min(C.width, C.height))$ // number of layers
- 2: $Tmp_1 \leftarrow \text{GUIDEDFILTER}(C, 0.1^2, 2)$ // first operand is input image, second is range sensitivity, third is spatial extent
- 3: **for** $\ell = 1 \dots n$ **do**
- 4: $Tmp_2 \leftarrow Tmp_1$
- 5: $Tmp_1 \leftarrow \text{GUIDEDFILTER}(Tmp_2, 0.1^2, 2^\ell)$ // double the spatial extent each time
- 6: $S_\ell = Tmp_2 - Tmp_1$ // a subband is the difference of two successively filtered versions of the image
- 7: **end for**

Algorithm 3 SIGNANDFREQUENCYSELECTED

input: $c \in \mathbb{R}$, $sign \in \{\text{pos, neg, all}\}$, $freq \in \{\text{high, low, all}\}$
output: $isSelected \in \{\text{true, false}\}$

- 1: $isSelected \leftarrow sign = \text{all}$
 or ($c < 0$ **and** $sign = \text{neg}$)
 or ($c > 0$ **and** $sign = \text{pos}$)
- 2: $isSelected \leftarrow freq = \text{all}$
 or ($\ell \leq n/2$ **and** $freq = \text{low}$)
 or ($\ell > n/2$ **and** $freq = \text{high}$)
- 3: $isSelected \leftarrow isSelected$ **and** $freqSelected$

FATTAL, R., AGRAWALA, M., AND RUSINKIEWICZ, S. 2007. Multiscale shape and detail enhancement from multi-light image collections. *ACM Transactions on Graphics (Proc. SIGGRAPH) 26*, 3.

FLEMING, R. W. AND BÜLTHOFF, H. H. 2005. Low-level image cues in the perception of translucent materials. *ACM Transactions on Applied Perception 2*, 3.

FLEMING, R. W., WIEBEL, C., AND GEGENFURTNER, K. 2013. Perceptual qualities and material classes. *Journal of Vision 13*, 8, 9.

GASTAL, E. S. L. AND OLIVEIRA, M. M. 2011. Domain transform for edge-aware image and video processing. *ACM Transactions on Graphics (Proc. SIGGRAPH) 30*, 3.

GASTAL, E. S. L. AND OLIVEIRA, M. M. 2012. Adaptive manifolds for real-time high-dimensional filtering. *ACM Transactions on Graphics (Proc. SIGGRAPH) 31*, 4.

GLONDU, L., MUGUERCIA, L., MARCHAL, M., BOSCH, C., RUSHMEIER, H., DUMONT, G., AND DRETTAKIS, G. 2012. Example-based fractured appearance. *Computer Graphics Forum (Proc. Eurographics Symposium on Rendering) 31*, 4.

GOLOVINSKIY, A., MATUSIK, W., PFISTER, H., RUSINKIEWICZ, S., AND FUNKHOUSER, T. 2006. A statistical model for synthesis of detailed facial geometry. *ACM Transactions on Graphics (Proc. SIGGRAPH) 25*, 3.

HANIKA, J., DAMMERTZ, H., AND LENSCH, H. P. A. 2011. Edge-optimized à-trous wavelets for local contrast enhancement with robust denoising. *Computer Graphics Forum 30*, 7.

HARRIS, M., SENGUPTA, S., AND OWENS, J. D. 2007. Parallel prefix sum (scan) with CUDA. In *GPU Gems 3*, H. Nguyen, Ed. Addison Wesley, Chapter 39, 851–876.

HE, K., SUN, J., AND TANG, X. 2010. Guided image filtering. In *Proceedings of European Conference on Computer Vision*.

HEGER, D. J. AND BERGEN, J. R. 1995. Pyramid-based texture analysis/synthesis. In *Proc. of ACM SIGGRAPH*.

KARACAN, L., ERDEM, E., AND ERDEM, A. 2013. Structure preserving image smoothing via region covariances. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia) 32*, 6.

- KHAN, E. A., REINHARD, E., FLEMING, R. W., AND BÜLTHOFF, H. H. 2006. Image-based material editing. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 25, 3.
- KIM, J., MARLOW, P., AND ANDERSON, B. L. 2011. The perception of gloss depends on highlight congruence with surface shading. *Journal of Vision* 11, 9, 4.
- LAND, E. H., JOHN, AND MCCANN, J. 1971. Lightness and retinex theory. *Journal of the Optical Society of America*, 1–11.
- LIU, C., SHARAN, L., ADELSON, E. H., AND ROSENHOLTZ, R. 2010. Exploring features in a bayesian framework for material recognition. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*.
- MALLICK, S. P., ZICKLER, T., BELHUMEUR, P. N., AND KRIEGMAN, D. J. 2006. Specularity removal in images and videos: A pde approach. In *In Proc. of ECCV*. 550–563.
- MANTIUK, R., KIM, K. J., REMPEL, A. G., AND HEIDRICH, W. 2011. HDR-VDP-2: A calibrated visual metric for visibility and quality predictions in all luminance conditions. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 30, 4.
- MANTIUK, R., MYSZKOWSKI, K., AND SEIDEL, H.-P. 2006. A perceptual framework for contrast processing of high dynamic range images. *ACM Transactions on Applied Perception*.
- MARLOW, P., KIM, J., AND ANDERSON, B. 2012. The perception and misperception of specular surface reflectance. *Current Biology* 22, 20, 1909 – 1913.
- MARLOW, P. J. AND ANDERSON, B. L. 2013. Generative constraints on image cues for perceived gloss. *Journal of Vision* 13, 14, 2.
- MERTENS, T., KAUTZ, J., CHEN, J., BEKAERT, P., AND DURAND, F. 2006. Texture transfer using geometry correlation. In *Proc. of Eurographics Symposium on Rendering*.
- MOTOYOSHI, I., NISHIDA, S., SHARAN, L., AND ADELSON, E. H. 2007. Image statistics and the perception of surface qualities. *Nature*.
- NGAN, A., DURAND, F., AND MATUSIK, W. 2006. Image-driven navigation of analytical BRDF models. In *Proceedings of the Eurographics Symposium on Rendering*.
- PARIS, S., HASINOFF, S. W., AND KAUTZ, J. 2011. Local Laplacian filters: Edge-aware image processing with a Laplacian pyramid. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 30, 4.
- PATTERSON, G., XU, C., SU, H., AND HAYS, J. 2014. The sun attribute database: Beyond categories for deeper scene understanding. *Int. J. Comput. Vision* 108, 1-2 (May), 59–81.
- PORTILLA, J. AND SIMONCELLI, E. P. 2000. A parametric texture model based on joint statistics of complex wavelet coefficients. *Int. Journal Computer Vision* 40, 1 (Oct.), 49–70.
- SHARAN, L., LI, Y., MOTOYOSHI, I., NISHIDA, S., AND ADELSON, E. H. 2008. Image statistics for surface reflectance perception. *Journal of the Optical Society of America A* 25, 4.
- SHARAN, L., LIU, C., ROSENHOLTZ, R., AND ADELSON, E. H. 2013. Recognizing materials using perceptually inspired features. *International Journal of Computer Vision*.
- SIMONCELLI, E. P. AND ADELSON, E. H. 1996. Noise removal via bayesian wavelet coring. In *Proc. of IEEE International Conference on Image Processing*.
- TOMASI, C. AND MANDUCHI, R. 1998. Bilateral filtering for gray and color images. In *Proc. of IEEE Int. Conf. on Computer Vision*.
- TRENTACOSTE, M., MANTIUK, R., AND HEIDRICH, W. 2011. Blur-aware image downsizing. *Computer Graphics Forum (Proc. Eurographics)*.
- TRENTACOSTE, M., MANTIUK, R., HEIDRICH, W., AND DUFROT, F. 2012. Unsharp masking, countershading and halos: Enhancements or artifacts? *Computer Graphics Forum (Proc. Eurographics)*.
- VERGNE, R., BARLA, P., FLEMING, R., AND GRANIER, X. 2012. Surface flows for image-based shading design. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 31, 3.
- VON DER PAHLEN, J., JIMENEZ, J., DANVOYE, E., DEBEVEC, P., FYFFE, G., AND ALEXANDER, O. 2014. Digital ira and beyond: Creating real-time photoreal digital actors. In *ACM SIGGRAPH 2014 Courses*. SIGGRAPH '14. ACM, New York, NY, USA, 1:1–1:384.
- WARD, G. J. 1992. Measuring and modeling anisotropic reflection. *SIGGRAPH Comput. Graph.* 26, 2 (July), 265–272.
- XU, L., LU, C., XU, Y., AND JIA, J. 2011. Image smoothing via L0 gradient minimization. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 30, 6.
- XU, L., YAN, Q., XIA, Y., AND JIA, J. 2012. Structure extraction from texture via relative total variation. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 31, 6.