

Non-Parametric Filtering for Geometric Detail Extraction and Material Representation

Zicheng Liao¹ Jason Rock¹ Yang Wang² David Forsyth¹
¹University of Illinois at Urbana-Champaign ²University of Manitoba
{liao17, jjrock2, daf}@illinois.edu ywang@cs.umanitoba.ca




Figure 1. We decompose an image (a) into three components: (b) albedo, (c) coarse-scale shading, and (d) shading detail. The albedo and coarse-scale shading represent surface color and directional lighting effect. The shading detail image captures fine-scale surface geometry, or material property.

Abstract

*Geometric detail is a universal phenomenon in real world objects. It is an important component in object modeling, but not accounted for in current intrinsic image works. In this work, we explore using a non-parametric method to separate geometric detail from intrinsic image components. We further decompose an image as albedo * (coarse-scale shading + shading detail). Our decomposition offers quantitative improvement in albedo recovery and material classification. Our method also enables interesting image editing activities, including bump removal, geometric detail smoothing/enhancement and material transfer.*

1. Introduction

It is usual to decompose the brightness $I(x, y)$ of an image pixel into its *albedo* (the percentage of that light reflected to the camera, $A(x, y)$), and its *irradiance* (the amount of light collected by the surface viewed by the pixel, $S(x, y)$, sometimes called shading). Assuming a calibrated sensor we have:

$$I(x, y) = A(x, y)S(x, y).$$

These two components are sometimes known as *intrinsic images*, after [2].

However, this decomposition fails to model *geometric detail* properly. Geometric detail refers to fast changing patterns of surface geometry. These are important real world phenomena, strongly associated with particular *materials*. Examples include wrinkles on a crusted paper, grooves on skin, and cracks on a trunk. Geometric detail occurs nearly universally on real world objects, and is an important component of object modelling. Geometric detail does not fit naturally into an albedo-irradiance decomposition, because it can generate large, localized gradients (like albedo) but is essentially a shading effect. Because geometric detail has a drastically different local spatial pattern than albedo or shading, we use a non-parametric patch-based model learned from albedo to remove geometric detail error in albedo; similarly for separating geometric detail in shading. Experimental results show the non-parametric method is very effective. We call it *patch-based filtering*.

Based on the patch-based filtering, we decompose irradiance into *coarse-scale shading* and *shading detail*. Coarse-scale shading is a relatively smooth signal which represents the average irradiance over a patch of surface. Shading detail is a fine scale detail, emphasizing effects caused by local surface geometry like bumps and grooves. Figure 1 shows our decomposition; notice that the shading detail represents geometric effects like the creases and fibers in the paper,

whereas the shading represents the irradiance at a longer spatial scale. Dropping pixel coordinates, we have

$$I = A(S_c + S_d) \quad (1)$$

where A is albedo, S_c is coarse-scale shading, and S_d is shading detail. We achieve the decomposition using a novel strategy that combines conventional intrinsic image methods with a reestimation step using our patch-based filtering.

Contributions: (1) We propose an effective method of separating geometric detail from image. (2) Our decomposition is novel. We demonstrate a number of applications from them. By applying the patch-based filter to images, we effectively remove geometric details such as bumps (section 3.1). Our decomposition results in improved albedo estimation because it does not force the albedo signal to account for all fast local effects (section 3.2). The shading detail image permits attractive image editing applications, where the apparent material of an object in an image can be smoothed, enhanced or transferred (section 3.3). Finally, the shading detail signal is a powerful cue to material identity, as we demonstrate in classification experiments (section 3.4).

1.1. Related work

Intrinsic image decomposition: Land’s influential Retinex model assumes effective albedo displays sharp, localized changes (which result in large image gradients), and that shading has small gradients [19]; important variants include [16, 3, 6]. Sharp shading changes occur at shadow boundaries or normal discontinuities, but using chromaticity [13] or multiple images under different lighting conditions [30] yields improved estimates. Discriminative methods to classify edges into albedo or shading help [27]; chromaticity cues can contribute [10], as can user input [5]. Tapen *et al.* regress local intrinsic image patches against the image, exploiting the constraint that patches join up [26]. When more than one image is available, recent methods can recover quite complex surface properties [12].

Patch-based image de-noising methods use a dictionary of “typical” image patches to recover a noise-free image from given data (for example, [7, 17, 9]; a comprehensive comparison is in [32]). **Discriminative dictionary methods** build a dictionary that is explicitly discriminative, typically focussed on texture discrimination [22, 23]. Our patch-based filtering resembles the patch-based de-noising work but our motivation is to learn a discriminative dictionary to separate signals with structured spatial pattern instead of noise.

Image-based editing: White and Forsyth [31] show that, for screenprinted surfaces, one can separate shading and albedo sufficiently accurately to composite the shading together with a new albedo map. Khan *et al.* [18] show how to simulate changes in the apparent material given

an approximate normal field. In our work, the separation of shading into coarse-scale shading and shading detail provides well-behaved editing primitives: The coarse-scale shading captures gross shading; and the detail shading carries only fine-scale material details. We exploit the image analogy technique [15] to synthesize and transfer shading details from one object to another. Transferring detail from one image to another produces visually realistic results, even though local shading inconsistencies are inevitable; but as [24] shows, human observers are unable to spot inconsistency in fine-scale shading, as long as gross shading is correct.

Material recognition is an established topic, with few strong results. There is a clear distinction between material and texture (eg [21], Figs 2 and 3), although texture methods apply generally. The CURET database [8] consists of a set of views of 61 material surfaces under differing illumination conditions. Varma and Zisserman show very strong results on this and other datasets, using nearest neighbor methods applied to a patch based representation [29]. Liu *et al.* offer a demanding dataset of material images, and show that the methods of Varma and Zisserman are not particularly strong at discriminating between these images [21]. However, Liu *et al.*’s dataset aggressively mixes spatial scales (50 pure material, 50 object scale images per category), and covers a relatively small range of materials. This suggests a need for dataset more demanding than CURET, but less broad in scope than Liu *et al.*’s, and we offer one in section 3.4.

1.2. Outline

Our patch-based filter is the base of this work. It allows us to remove geometric detail from albedo, shading or image, and extract the shading detail image. The operations on shading and albedo lead to our new image decomposition.

Our strategy to decompose an image is as follows. First, we use the state-of-the-art decomposition method of [1], or a variant of it (section 3.4), to obtain preliminary estimates of albedo and irradiance. This approach uses explicit representations of a prior on albedo, together with a prior on irradiance (as represented by illumination and geometry) to produce an initial albedo field \hat{A} and an initial irradiance field \hat{S} . However, as noted above, the preliminary albedo field tends to contain small signals caused by local geometric bumps. We correct \hat{A} using our patch-based filter (section 2); the result is the albedo field A . Furthermore, we compute the shading field S_c from the irradiance field using a similar filter learned from shading (again, section 2). Finally, we compute the shading detail field by $S_d = \hat{S} - S_c$, (Fig. 2). Note there is missing shading details in the albedo correction step. Recovering it for S_d requires division by A (as Equation 1 suggests), which causes artifact around albedo edges. So we extract an approximated S_d only from the shading field instead.

2. Patch-based filtering

Assume we have obtained a preliminary albedo field \hat{A} from existing methods. We wish to polish it. Typical errors in albedo fields from modern methods look like the shading on local geometric detail; this appears in albedo maps because shading priors strongly discourage large gradients. We can cope with these errors, caused by the relatively weak parametric representation of the prior, by a non-parametric model. We expect that the exact albedo fields are properly described using a patch dictionary. Typical albedo patches will have large constant domains, while geometric detail has a wildly different spatial pattern (see an illustration in Fig. 3). So the error of geometric detail in albedo estimates will be poorly encoded by an albedo dictionary.

We apply a standard method ([20]), described for reference here. We represent \hat{A} as a collection of patches $X \in R^{k \times m}$ (each column is a sample patch). Write $P(\hat{A})$ for the mapping that takes \hat{A} to X . We assume we have a dictionary $B \in R^{k \times n}$, learned using the procedure below. We now choose S to minimize

$$\min_S \|P(\hat{A}) - BS\|_F^2 + \lambda \|S\|_1 \quad (2)$$

where $S \in R^{n \times m}$ are the patch coefficients. We then reconstruct by identifying A such that $P(A) = BS$. In practice, this is straightforward. Each patch is placed over its location in the image, covering multiple pixels. Each pixel may be covered by many patches, and so there are many predicted values. At each pixel, A takes the median of all candidate values from the over-sampled patches. This helps to avoid oversmoothing.

Learning: We use a standard procedure with L1 regularization. Given sample patches $X \in R^{k \times m}$ from a set of training examples, we want to learn a dictionary $B \in R^{k \times n}$ that minimizes:

$$\begin{aligned} \min_B \quad & \|X - BS\|_F^2 + \lambda \|S\|_1 \\ \text{subject to} \quad & \forall j = 1, \dots, n, \sum_{i=1}^k B_{i,j}^2 \leq c \end{aligned} \quad (3)$$

where $S \in R^{n \times m}$ are the coefficient matrix. We use the fast sparse coding solver from Lee *et al.* [20]. It iteratively solves for the coefficient matrix S uses the feature-sign search algorithm, and optimizes the basis B using the Lagrangian dual.

We apply the same method to learn the albedo and shading dictionaries using the (clean) ground truth albedo and shading images of the MIT intrinsic image dataset [14], respectively. Figure 3 also shows exemplar entries of a learned geometric detail dictionary, for which we use a set of typical bump images collected online.




Figure 2. Decomposing initial shading \hat{S} (Left) into coarse-scale shading S_c (Middle) and shading detail image S_d (Right).




Figure 3. Visualization of patches learned from albedo (left), shading (center) and geometric bumps maps. Notice how they differ from each other: albedo patches have large pieces of constant value with sharp boundaries; shading patches have smooth intensity change; and the geometric detail patches have more complex fast changing patterns.

3. Applications

We show four applications of our decomposition: *geometric detail removal*, *albedo estimation*, *material editing*, and *material recognition*.

3.1. Geometric detail removal

Geometric detail in images can operate as a form of noise that impedes matching and tracking. For example, geometric detail on clothing can create real difficulty scoring the match between image regions in a tracker (*e.g.* Fig. 11.6 of [11]). It is beneficial to be able to remove geometric detail from images. We do so by applying our patch-based filter directly on images. We learn patch bases from images that do not contain geometric details and use them to reconstruct images. The output image is a geometry-detail free image, while albedo and shading are well preserved. Figure 4 shows a few examples of our geometric detail removal operation. We use dictionary size = 500, patch size = 12×12 .

Edge preservation The over-sampling step in reconstruction tends to create an image with blurred edges even when the median of sampled values is taken. This is because channels with small values near an edge tend to have large reconstruction error. We fix this problem with a *color chan-*




Figure 4. Geometric detail removal on MIT intrinsic image dataset. Top row: images with geometric detail; Bottom row: detail removal results. The first three columns are on original examples in the MIT intrinsic images dataset. In the last two examples, we add external bumps to the image by bump mapping and re-render on a recovered shape using a physically-based renderer. Local views of column 1, 2 and 3 are displayed with scaled intensity for better viewing. Notice how effectively the various types of geometric detail are removed, while the shading and albedo edges are preserved.

nel correction: for every pixel in the reconstructed image, we consult the corresponding pixel in the source image, compute the ratios of the two smaller color channels to the larger and adjust the channels in the output image so the color ratio is preserved.

3.2. Albedo estimation

Geometric detail causes problems for albedo estimation algorithms. Barron and Malik show strong albedo reconstruction results on the MIT intrinsic image dataset [1] for two illumination conditions: laboratory illumination, and “natural” illumination. Laboratory illumination involves actual images of physical objects, but this case has relatively few illumination fields. “Natural” illumination images are obtained by reshaping the ground truth shape with a richer set of illumination fields. Table 1 of [1] shows an improvement of near an order of magnitude in albedo recovery MSE for “natural” illumination over laboratory illumination. We hypothesize this occurs because “natural” illumination images tend to suppress the contribution of geometric detail, because it is not preserved in the ground truth shape (estimated from photometric stereo method); but geometric detail appears in the laboratory images, and tends to appear in their recovered albedo fields. We confirm this by applying their method to the original (laboratory illumination) images, producing examples like Fig 5-a. Notice how the albedo field contains high frequency shading terms caused by geometric detail. These are relatively easily removed by our patch-based filtering process (Fig 5-b).

We evaluate our refined albedo estimates on the MIT intrinsic image dataset quantitatively and show significant improvement. We compare our result (obtained by refining the Barron and Malik albedo estimate with our patch-based filtering) with two baseline methods: color Retinex [14] and




Figure 5. Our albedo refinement improves albedo estimates by suppressing vestiges of albedo detail. (a): Albedo estimate from Barron and Malik [1]; notice the geometry detail. (b): Our refined albedo by our patch-based filtering; Notice the geometric detail is suppressed.

Barron and Malik [1]. We split the MIT intrinsic image dataset into two sets, learn the dictionary using the ground-truth albedo of one set (10 images) for our filter, and use the other set for evaluation. Table 1 shows the MSE of the four results. Color channel correction is always applied to the reconstruction outputs to preserve albedo edges. Dictionary size is 500; patch size is 12×12 .

Our albedo estimates have slight improvements over Barron and Malik in MSE, but MSE is not an accurate metric for albedo error. Small geometric detail signals, as in Fig. 5-(a) generate small MSE values but cause undesirable visual artifacts. Similarly, small contrast changes in the in-

Method	MSE
Retinex	0.0522
Tappen et al.	0.0379
Shen et al.	0.0458
Barron&Malik	0.0152
Ours	0.0146

Table 1. Albedo estimation error comparison on the MIT intrinsic images dataset (Laboratory illumination). Numbers of Tappen et al. and Shen et al. are excerpted from [1] (Table 1) for reference.




Figure 6. Edge-based albedo estimation error comparison. The x-axis is the gradient threshold of edge pixels, i.e. pixels whose gradient magnitude is great than the threshold are taken as edge pixels.

ferred albedo may be insignificant in practice, but produce large MSE values. One way to compare the results is to compute the differences in gradient field for pixels with “large” (greater than some threshold) gradients (edges). Fig 6 shows the result for a range of thresholds. These results strongly suggest that our method and color retinex preserve the location of gradients (though color retinex cannot correctly estimate albedos—see the high MSE in table 1); and that Barron and Malik produce small but significant errors in the locations of large gradients, most likely due to geometric detail as in Fig. 5.

3.3. Material editing

Human eyes are insensitive to fine-scale shading inconsistencies as long as long-scale directional shading is correct, according to [24]. This allows us to apply image-based editing on the shading detail map (Equation 4) without creating noticeable visual artifacts. To perform image-based detail editing we do not need to explicitly compute the surface geometry and material; nor do we need to re-render the object using a physically based renderer which is difficult and computationally expensive. The shading detail

map mostly represents the material property; by editing on it we modify object material appearance while preserving the original color and directional shading.

Detail smooth/enhancement is achieved simply by adding a coefficient to the shading detail map S_d in the decomposition and adjusting it:

$$I' = A(S_c + kS_d) \quad (4)$$

When $k = 1$, we get $I' = I$. When $k = 0$, we get a geometric detail free image. When $0 < k < 1$ we get a detail-dampened version of the image. By increasing k from 1 we get a detail-enhanced image. This yields a wide range of image variants; see Fig. 7 for two examples. The editing operation is extremely simple and fast. To demonstrate its convenience, we developed an easy-to-use interface which loads the three components of an image, and allows users to adjust the scale of shading detail magnitude interactively with a slider bar.

Material synthesis and transfer We can also apply the shading detail image of one object to another by texture synthesis. This results in transferring object material property (Fig. 8). We use the image analogy technique by Hertzmann *et al.* [15] to synthesize a detail image from a source detail. Image analogy is an image synthesis technique that create pairs of analogous images. Given an exemplar image, a filtered version of the exemplar image, and an input image, it synthesizes a “filtered” image of the input image that mimics the “filtering” function in the exemplar image pair. In our shading detail synthesis, we use the shading and shading detail image of the source object as the exemplar and exemplar filtered image pair. We then supply the shading of a target object. The synthesized image will inherit patterns from the source shading detail image, yet adapt to the local appearance of the target object’s shading. After we obtained the synthesized shading detail, we use Equation 4 and adjust the weight k by our interface to reach a desired material transfer effect.

3.4. Material recognition

The shading detail map is decoupled from image albedo and directional shading, amplifying the fine structure of object material. We demonstrate that it is a good representation for material discrimination.

We start by building a new material dataset. The Flickr Material Database (FMD) of [25] is a good candidate to start with. But it includes object scale material examples (each class has 50/100 images at object scale) which incorporate object context but lose fine-scale material detail. Our dataset consists of complex material images in local scale “taken in the wild”. Furthermore, it has a rich categorical structure. In particular, *bark*, *fabric*, *construction materials*, and *outer coat of animals* make up the first level. Each

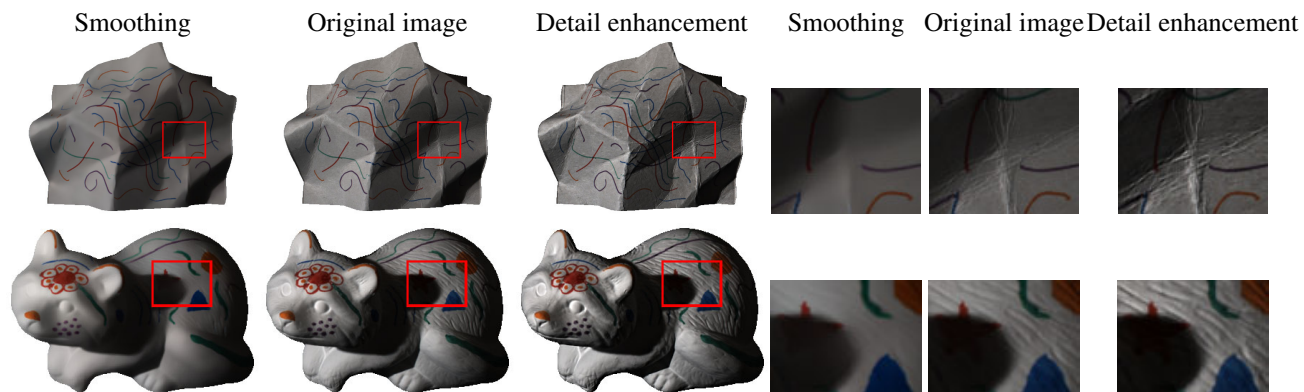


Figure 7. Geometric detail smoothing and sharpening. Notice how effectively we change the perception of surface geometry through image-based editing using the shading detail.




Figure 8. Material transfer by shading detail synthesis and editing. Left: Original images; Middle: New images with transferred material appearance. Right: Synthesized shading detail maps for composition (with detail source image in the corner). The detail images are translated and normalized to 0-1 for display.



Figure 9. Material image (column 1, 3) and the shading detail image (column 2, 4) examples.

has subcategories, up to a 3-level hierarchy.

We use a variant of Barron and Malik [1] in shading estimation for the shading detail extraction. Barron and Malik [1] has a strong shading smoothness prior (parametrized in shape and illumination) that pushes geometric shading detail into albedo, which is lost in shading detail extraction. Instead, we use a weaker shading prior:

$$f_s(S) = \sum_i (\nabla S_i)^2 \quad (5)$$

to replace their shape and illumination model. We also remove the absolute color term $g_a(R)$ in [1] (equation 6). This works very well in practice. See Fig. 9 for examples of the shading detail maps we get at the end of the pipeline.

We choose a standard classifier: bag of words model on PHOW [4] features trained and tested on linear SVM to highlight the quality of material representation from different sources: *color image*, the *micro-texture image* from [21] (a bilateral filtered image detail), the *bilateral filtered shading detail*, *our shading detail*, and the combination of the color image and the detail images.

Results We test the performance on our material dataset and FMD. Table 2 and Fig. 10 display the classification accuracy and confusion matrix. Combining image and detail features results in significantly improved performance on both datasets. Our shading detail, when combined with image, outperforms the other two detail images. It also shows: 1. PHOW is a reasonable choice because it beats Varma and Zisserman [28] on the FMD dataset. 2. The image is more discriminative than detail images alone, confirming that albedo is correlated to material. 3. Supplemental detail images help in material classification, because they capture complementary material information. 4. Our shading detail performs better than other detail images in combination with the color image. This suggests our shading detail best captures decoupled material features from the image. Although the micro-texture (Bilateral filtered detail map) from [21] has higher accuracy on its own, it provides little gain in accuracy when combined with image features, meaning that it carries mostly redundant material information.

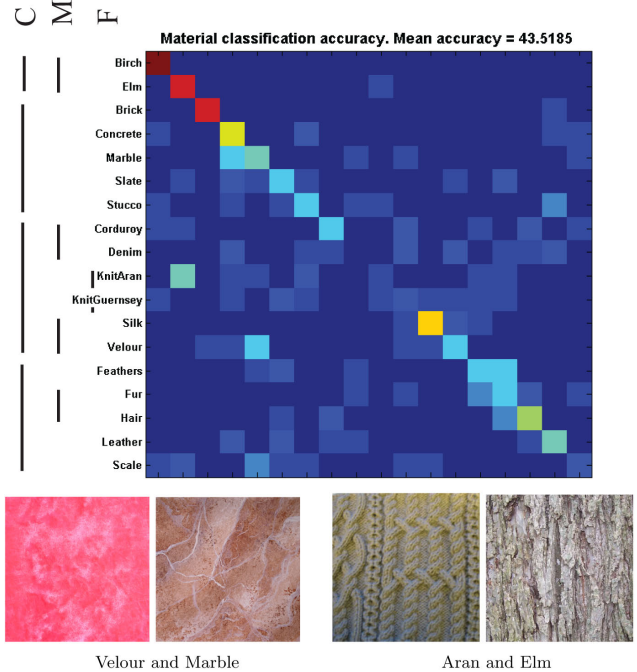


Figure 10. Classification confusion matrix on our 18-class material dataset by image feature + our shading detail feature. The bars on the left indicate the categorical similarity from coarse (C) to medium (M) to fine (F). We also show two unexpected confusions (Velour and Marble, Aran and Elm) which actually have visual similarity.

Feature source	FMD [25]	Out dataset
image	0.240	0.389
micro-texture	0.245	0.296
s_d bl	0.247	0.208
our s_d	0.243	0.241
image + micro-texture	0.287	0.394
image + s_d bl	0.272	0.407
image + our s_d	0.299	0.435

Table 2. Material classification with different material representations (feature source). Micro-texture is the micro-texture image from [21], “ s_d bl” is $\hat{S} - bilater(\hat{S})$, and s_d is our shading detail. Varma and Zisserman [28] have accuracy 0.238 on FMD. Liu et al. [21] have accuracy 0.388 on FMD; but they use specialized classifier with complex object-level features suited to the dataset.

4. Conclusions and Future work

In this paper, we separate geometric detail from intrinsic image components by a non-parametric patch-based filter, and propose a new image decomposition. Our method leads to quantitative improvement in albedo recovery and material discrimination. We also demonstrate interesting image editing activities, i.e., material enhancement and transfer.

We believe separating geometric detail from intrinsic image components has high application potential, especially in the direction of material representation.

While our two-step decomposition solution works well in practice, a natural extension is to explore a unified solution to the image decomposition problem of Equation 1. Another possibility is to learn discriminative dictionaries for albed-shading decomposition, or the separation of other pairs of intrinsic images, *e.g.* diffuse shading and specular-ity.

Acknowledgements: This work was supported in part by NSF Expeditions award IIS-1029035 and in part by ONR MURI award N000141010934.

References

- [1] J. T. Barron and J. Malik. Color constancy, intrinsic images, and shape estimation. In *ECCV*, 2012.
- [2] H. Barrow and J. Tenenbaum. Recovering intrinsic scene characteristics from images. In *Comp. Vision Sys.*, pages 3–26, 1978.
- [3] A. Blake. Boundary conditions for lightness computation in mondrian world. *Computer Vision, Graphics and Image Processing*, 32:314–327, 1985.
- [4] A. Bosch, A. Zisserman, and X. Muoz. Image classification using random forests and ferns. In *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*, pages 1–8. IEEE, 2007.
- [5] A. Bousseau, S. Paris, and F. Durand. User assisted intrinsic images. In *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2009)*, volume 28, 2009.
- [6] G. Brelstaff and A. Blake. Computing lightness. *Pattern Recognition Letters*, 5(2):129–38, 1987.
- [7] A. Buades, B. Coll, and J.-M. Morel. A non-local algorithm for image denoising. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 2, pages 60 – 65, june 2005.
- [8] K. Dana, B. Van-Ginneken, S. Nayar, and J. Koenderink. Reflectance and Texture of Real World Surfaces. *ACM Transactions on Graphics (TOG)*, 18(1):1–34, Jan 1999.
- [9] W. Dong, X. Li, L. Zhang, and G. Shi. Sparsity-based image denoising via dictionary learning and structural clustering. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 457–464, june 2011.
- [10] M. Farenzena and A. Fusiello. Recovering intrinsic images using an illumination invariant image. In *IEEE ICIP*, 2007.
- [11] D. Forsyth and J. Ponce. *Computer Vision: A Modern Approach(2E)*. Pearson, 2011.
- [12] F.Romeiro, Y. Vasilyev, and T. Zickler. Passive reflectometry. In *ECCV*, 2008.
- [13] B. Funt, M. Drew, and M. Brockington. Recovering shading from color images. In *ECCV*, 1992.
- [14] R. Grosse, M. K. Johnson, E. H. Adelson, and W. T. Freeman. Ground-truth dataset and baseline evaluations for intrinsic image algorithms. In *International Conference on Computer Vision*, pages 2335–2342, 2009.
- [15] A. Hertzmann, C. E. Jacobs, N. Oliver, B. Curless, and D. H. Salesin. Image analogies. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques, SIGGRAPH '01*, pages 327–340, New York, NY, USA, 2001. ACM.
- [16] B. K. P. Horn. Determining lightness from an image. *Computer Vision, Graphics and Image Processing*, 3:277–299, 1974.
- [17] C. Kervrann and J. Boulanger. Local adaptivity to variable smoothness for exemplar-based image regularization and representation. In *Int. J. Comput. Vision*, volume 79, pages 45–69, Hingham, MA, USA, Aug. 2008.
- [18] E. A. Khan, E. Reinhard, R. W. Fleming, and H. H. Bühlhoff. Image-based material editing. In *ACM SIGGRAPH 2006 Papers, SIGGRAPH '06*, pages 654–663, New York, NY, USA, 2006. ACM.
- [19] E. Land and J. McCann. Lightness and retinex theory. In *J. Opt. Soc. Am.*, volume 61, pages 1–11, 1971.
- [20] H. Lee, A. Battle, R. Raina, and A. Y. Ng. Efficient sparse coding algorithms. In *In NIPS*, pages 801–808. NIPS, 2007.
- [21] C. Liu, L. Sharan, E. H. Adelson, and R. Rosenholtz. Exploring features in a bayesian framework for material recognition. In *CVPR*, 2010.
- [22] J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman. Discriminative learned dictionaries for local image analysis. In *CVPR*, 2008.
- [23] J. Mairal, M. Leordeanu, F. Bach, M. Hebert, and J. Ponce. Discriminative sparse image models for class-specific edge detection and image interpretation. In *ECCV*, 2008.
- [24] Y. Ostrovsky, P. Cavanagh, and P. Sinha. Perceiving illumination inconsistencies. In *Perception*, volume 34, pages 1301–1314, 2005.
- [25] L. Sharan, R. Rosenholtz, and E. Adelson. Material perception: What can you see in a brief glance? *Journal of Vision*, 9(8):784–784, 2009.
- [26] M. Tappen, W. Freeman, and E. Adelson. Estimating intrinsic component images using non-linear regression. In *CVPR*, 2006.
- [27] M. Tappen, W. Freeman, and E. Adelson. Recovering intrinsic images from a single image. *PAMI*, 2006.
- [28] M. Varma and A. Zisserman. A statistical approach to texture classification from single images. *International Journal of Computer Vision*, 62(1):61–81, 2005.
- [29] M. Varma and A. Zisserman. A statistical approach to material classification using image patch exemplars. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(11):2032–2047, November 2009.
- [30] Y. Weiss. Deriving intrinsic images from image sequences. In *ICCV*, 2001.
- [31] R. White and D. Forsyth. Retexturing single views using texture and shading. In *ECCV*, pages 70–81, 2006.
- [32] L. Zhang. Patch based denoising withh andd without dictionary learning: a comparative study. In *Mini-symposium on Recent advances in patch-based image processing. SIAM IS*, 2012.