User-assisted Image Compositing for Photographic Lighting

Ivaylo Boyadzhiev
Cornell University

Sylvain Paris
Adobe

Kavita Bala
Cornell University

(a) Input images of a large indoor scene under different lighting (showing 16 of 112 images).

(b) Average image

(c) Lighting design by a novice user, exploring our system for 5 minutes

Figure 1: *Cafe: (a) A large indoor scene lit with a light in various positions. Photographers usually spend hours selecting and blending desirable parts from different images to produce a final image. Possible solutions, like an average image (b) produce unappealing results. (c) We propose a set of basis lights and modifiers based on common photography goals that allow users to produce final results in a few minutes.*

Abstract

Good lighting is crucial in photography and can make the difference between a great picture and a discarded image. Traditionally, professional photographers work in a studio with many light sources carefully set up, with the goal of getting a near-final image at exposure time, with post-processing mostly focusing on aspects orthogonal to lighting. Recently, a new workflow has emerged for architectural and commercial photography, where photographers capture several photos from a fixed viewpoint with a moving light source. The objective is not to produce the final result immediately, but rather to capture useful data that are later processed, often significantly, in photo editing software to create the final well-lit image.

This new workflow is flexible, requires less manual setup, and works well for time-constrained shots. But dealing with several tens of unorganized layers is painstaking, requiring hours to days of manual effort, as well as advanced photo editing skills. Our objective in this paper is to make the compositing step easier. We describe a set of optimizations to assemble the input images to create a few *basis lights* that correspond to common goals pursued by photographers, e.g., accentuating edges and curved regions. We also introduce *modifiers* that capture standard photographic tasks, e.g., to alter the lights to soften highlights and shadows, akin to umbrellas and soft boxes. Our experiments with novice and professional users show that our approach allows them to quickly create satisfying results, whereas working with unorganized images requires considerably more time. Casual users particularly benefit from our approach since coping with a large number of layers is daunting for them and requires significant experience.

CR Categories: I.3.7 [Computing Methodologies]: Computer Graphics—Image Processing And Computer Vision

Keywords: lighting design, light compositing

Links: DL PDF

1 Introduction

Lighting is a key component of photography, on an equal footing with other aspects such as composition and content. In many cases, photographers actively illuminate their subject with a variety of lights to obtain a desired look. Lighting a scene is a challenging task that is the topic of many courses and books, e.g. [Hunter et al. 2011]. Not only the notion of "good" lighting is elusive and heavily relies on one's subjectivity, but the traditional way to set up the lights itself is complex. Positioning and setting the power of each flash is a nontrivial and tedious task; further, most lights are accompanied by modifiers that also need to be adjusted, e.g., a snoot to restrict the lit area, or a diffuser to soften the shadows.

While post-processing the result in photo editing software is common, this step has almost no effect on the lighting which essentially remains the same as what was captured at exposure time. Recently, a few photographers have introduced a new workflow to control lighting that relies a lot more on the editing stage. Instead of a single photo with many light sources, they take many photos with a single light located at different locations each time. Then, they load all the images as layers in photo editing software and carefully composite the images to produce the final image. There are several advantages to this workflow accounting for its increasing popularity. First, capture sessions are shorter, easier to set up, and require less equipment. Second, this workflow permits considerable control by enabling arbitrary layering and post-exposure adjustment over all the lights, allowing room for experimentation. For instance, one can easily control the region affected by a light source with a mask and

ACM Reference Format
Boyadzhiev, I., Paris, S., Bala, K. 2013. User-assisted Image Compositing for Photographic Lighting. ACM Trans. Graph. 32, 4, Article 36 (July 2013), 11 pages. DOI = 10.1145/2461912.2461973 http://doi.acm.org/10.1145/2461912.2461973.

DOI: http://doi.acm.org/10.1145/2461912.2461973

set its intensity with a brightness adjustment. This new process is fundamentally different from the traditional one because the capture session is not concerned with directly producing a visually pleasing image, it only seeks to record useful data for the later editing step. From this perspective, it is related to recent work in computational photography such as coded apertures, e.g., [Levin et al. 2007; Bishop and Favaro 2011], and lightfield cameras, e.g., [Ng et al. 2005; Adelson and Bergen 1991], in which computation is an integral part of the image formation process. This motivates us to name this modern approach *computational lighting design*. See [Kelley 2011; Guanzon and Blake 2011] for examples of this workflow from professional photographers that inspired us (unaffiliated with the project). These examples demonstrate the use of this workflow in architectural photography, and product photography.

However, one major disadvantage of this workflow is that it is quite cumbersome even for experienced photographers. When the number of images grows to several tens, or even above a hundred, navigating the corresponding layer stack becomes impractical. Further, with large scenes, most images show the main subject mostly in the dark with only a small part lit (see Figure 1a). Finding the useful images in the stack, setting their relative intensities, and blending them together to get visually pleasing results are highly challenging tasks that require advanced photography and image editing skills.

Contributions. In this paper we introduce a new approach to assist photographers in this workflow. Our contributions are:

• We create *basis lights* that achieve specific effects like accentuating edges, enhancing the material color of an object, or adding fill light.
• We introduce *modifiers* that affect the basis lights, and achieve effects similar to standard lighting equipment, such as soft boxes to soften highlights and shadows, and snoots to restrict the light extent.
• We design these lights and modifiers by reasoning entirely in image space, thereby avoiding a computationally expensive, and potentially brittle 3D reconstruction of the scene. This follows the trend set by Fattal et al. [2007] for detail enhancement and extends it to a larger set of photographic operations.
• We implement an interactive interface and demonstrate that casual photographers can benefit greatly from this interface.

1.1 Related Work

We discuss related work in terms of computational lighting design, single-image editing, and 3D lighting design.

Computational Lighting Design. Debevec et al. [2000], Akers et al. [2003] and Agarwala et al. [2004] provide user interfaces to combine several images taken from the same viewpoint but under different lighting, thereby introducing the idea of computational lighting design. In Debevec et al. a scene can be realistically relit under novel illumination conditions, by treating the input images as a basis and fitting a lighting model per image pixel. However, their method needs to know the position of the light source in each image, which requires specialized acquisition equipment. For comparison, our technique is based on simple, widely available equipment, such as a single flash, and we do not need to know the light positions. Further, their custom device has been designed for medium scale scenes, such as human faces, whereas we are interested in large scale architectural scenes. In Akers et al. and Agarwala et al. users mark the regions of interest in the input images and the algorithm is in charge of producing a satisfying composite. While this is a reasonable approach when there are only a few input photos, it becomes intractable when there are a hundred of them. With such datasets, deciding which images to use, and which parts in them to combine, is a major challenge that is as difficult as producing the actual combination. In our work, we introduce *basis light sources* to organize the many input photos into a smaller, more manageable set of images that correspond to standard photography goals.

Raskar et al. [2004], Cohen et al. [2003], Fattal et al. [2007] and Mertens et al. [2007] combine several photos taken under different lighting to generate a better picture. Raskar et al. generate non-photorealistic results, whereas we seek to retain a photorealistic look. Cohen et al. and Fattal et al. focus on a single object and aim at revealing the object's details and removing shadows, while ignoring other effects such as shadows projected onto other objects. As we shall see, this becomes an issue in larger scenes, that are of interest in our work. Further, their approaches are mostly automatic, with a few presets offered to users, whereas we give more control so that users can make artistic choices. Mertens et al. target high-dynamic range scenes as their main applications, and they show that their technique can also apply to simple multi-light configurations. But this technique does not handle the more diverse lighting configurations of the computational lighting workflow well.

Single-Image Lighting Editing. Several techniques exist to manipulate the lighting in a single image. For instance, Carroll et al. [2011] describe how to alter the interreflections in a picture after the user makes some annotations to describe the lighting configuration. Bousseau et al. [2009] and Boyadhziev et al. [2012] use similar annotations to perform white balance. Mallick et al. [2006] control the intensity of specularities. Tone-mapping operators remap intensities to fit the dynamic range to a given display [Reinhard et al. 2012]. Photo editing software follows a similar approach to define adjustments that brighten shadows and decrease highlights. From our perspective, all these methods have in common that they are limited to a specific effect such as white balance or shadow brightening, but keep the spatial configuration unchanged, e.g., shadows cannot be altered. In comparison, we seek to produce a wider range of effects to enable more control over the achieved lighting configuration.

3D Lighting Design. Several techniques exist to edit lighting environments in the context of 3D rendering, e.g., [Schoeneman et al. 1993; Pellacini 2010; Bousseau et al. 2011]. Compared to our approach, these methods tackle the problem from a fundamentally different direction since they have access to a full geometric description of the scene and have total control over the light sources. In comparison, we have no a priori knowledge about the photographed scene and have access to only a limited number of observations of it. One could try to solve an inverse problem to infer a 3D description of the scene and its materials, but current 3D reconstruction techniques are often fragile, and cannot handle such large problems, especially with a fixed viewpoint.

2 Motivation and Approach

Our objective is to assist photographers with creating a compelling lighting environment for a scene. In this section, we first describe how photographers work with lighting. This motivates our approach to designing lighting in this workflow.

2.1 Computational Lighting Design

With the traditional approach to lighting, photographers set up lights so that they fire simultaneously to produce the desired lighting. Then, illumination is essentially left untouched during post-processing. In comparison, for the computational workflow in which we are interested, photographers capture many images under different illuminations. The setup is often as simple as a single flash light moved to a different location between each shot. This has numerous advantages in terms of cost and mobility compared to the many lights used for studio lighting. Most importantly, the goal of the capture session is different. Whereas the traditional approach is concerned with the final result, the new computational workflow is about capturing useful "building blocks." Each photo aims to illuminate a portion of the scene in an interesting way that may be used later in combination with other images. The main objective is good coverage of the scene.

After the capture session, all the images are loaded as layers into image editing software. For each region, photographers select the desired appearance by editing the image alpha channels and ordering the layers appropriately. In parallel, each layer can be edited, for instance, to increase its brightness, which is equivalent to having a more powerful light source, but with all the advantages of using editing software, such as instant feedback and unlimited undo. Only when this process is done, that is, after all the adjustments and after combining the layers according to their alpha channels, is the final image produced. Further editing may occur, for instance to remove unwanted elements, but this is not in the scope of our work. See [Kelley 2011; Guanzon and Blake 2011] for examples.

2.2 Objectives of Photographic Lighting

There are many ways to illuminate a scene in photography. We found a few recurring trends based on interviews with professionals [Kelley 2012] and field reports [Kelley 2011; Guanzon and Blake 2011].

Discrete Reasoning. Photographers think of lighting as the discrete combination of a few standard configurations. For instance, the *key light* illuminates the main subject, and the *fill light* is aimed at shadows to control how dark they are. While the exact setups depend on each photographer and scene, decomposing illumination into a small number of objectives is standard practice. We match this approach with our *basis lights* that address a few well-defined goals, e.g., controlling shadow darkness.

Curves and Lines. Photographers identify a few important geometric features of the scene that they seek to accentuate in the final result. These features are typically occluding contours that separate two objects at different depth, surface discontinuities such as creases, and curved regions that are characteristic of the object's shape. To emphasize these features, photographers set the illumination up so that a highlight falls on one side and a shadow on the other. To help users accentuate scene edges and curved regions, we first analyze the input images to identify these features and then propose an energy function that favors high contrast around them. This defines what we call the *edge light*.

Light Quality. Photographers seek a "good light". While the concept is elusive, a few properties stand out. Harsh highlights and hard shadow boundaries are undesirable because they tend to distract from the scene. Overly dark shadows are also to be avoided because they hide the scene content. We propose a few options to mimic photographers' solutions to these issues. We simulate area light sources with several point sources to soften the highlight and shadows. We offer a *fill light* to control the darkness of the shadows, and let users restrict the extent of a light, which can be useful to prevent undesirable highlights and shadows.

2.3 Our Approach

We propose an approach inspired by the photographers' workflow described in the previous two sections. First, we describe the input data. Then, building upon our observations about the types of lights used, and lighting practices, we propose a set of basis light sources and controls that assist users in achieving powerful effects.

Input Photos. We use input data similar to what photographers capture, that is, a few tens or more photos taken from a fixed viewpoint and with a different lighting configuration each time, typically using a single flash light. We assume that the light sources are approximately uniformly distributed in the scene and that their power is approximately constant. Further, we assume that the input set of images is white balanced with respect to the color of the flash light, i.e., the light source appears white in the input photos. For the datasets that we acquired ourselves, we used a remotely triggered flash unit and moved it at a different position after each shot, covering the entire scene in about 100 pictures. This is a rather mechanical process, where the main goal is to get a good coverage of the whole scene, with no particular planning. For a single data set, we spent around 20 minutes photographing it. We put a camera on a tripod and walked around with a remotely triggered flash. Exposure was fixed so that the flash dominated the other lights. For the Library scene, (Fig. 8), we also took a few images with longer exposure so that the outside is visible. We tried to keep ourselves out of the shots, but in the occasional pictures where the equipment/photographer was visible, we manually masked it out, so that those regions are not considered later. This pre-processing step has to be done once.

Basis Lights. We propose an *edge light* that emphasizes edges and curved regions in the scene, a *diffuse color light* that emphasizes the underlying colors, and a *fill light* that provides more even illumination. For each of these lights, we formulate an energy function that models the objective, e.g., large gradients that align with the main scene features for the *edge light*. Minimizing the energy gives us a set of coefficients that we use to combine the input images.

Modifiers. We also introduce controls to mimic the effects of standard light modifiers. For instance, we embed the input images into a weighted graph based on their similarity and apply a diffusion process on this graph to modify a given set of coefficients to approximate the shadow-softening effect of an umbrella or a soft box. Other modifiers include *per-object* and *regional modifiers* that let us control the lighting for objects, like using a snoot, or to change the relative lighting of foreground vs. background.

The User Process. The user starts from a base image, for example the average of the image stack or a single image from the stack, and then they edit the stack using basis lights and modifiers. They first try to arrive at some globally reasonable solution, and then further refine the scene to identify either objects or parts of the scene that need more attention through edge enhancement, more light, or other lighting effects. See the supplementary for example sessions.

In allowing the user to pick individual objects, and applying an optimization of lighting for that particular object, we introduce inconsistent lighting in a scene. However, this is acceptable based on perceptual research about human insensitivity to lighting inconsistencies [Ostrovsky et al. 2005].

We evaluate our approach on several test cases demonstrating that it enables the design of sophisticated lighting environments with a small set of meaningful degrees of freedom, instead of a complex physical setup or the tedious manipulation of tens of layers. We further demonstrate the ease of lighting for novices and professionals using our basis lights in this new workflow.

3 Basis Lights

We propose a set of *basis* lights, which we relate to standard photography practices. Some of those lights correspond to actual lighting scenarios, commonly used in photography. The *fill light*, directly corresponds to lights used by photographers. Basis lights, like the *edge light* and the *diffuse color light*, address standard objectives such as emphasizing edges and curved areas, and revealing the diffuse color of objects, respectively.

We find the basis lights through an optimization scheme, that looks for the best linear combination of the input images, such that a certain objective is minimized. We first introduce some notation, and then describe the objective functions for our three basis lights: *fill*, *edge*, and *diffuse color*.

Standard Definitions and Notation. Throughout the paper, we use $\mathbf{I}_i(\mathbf{p})$ to denote the RGB components of a pixel $\mathbf{p}$ in the i^{th} input image. We work with sRGB values that are not gamma compressed, i.e. we apply inverse gamma correction by assuming 2.2 gamma. We refer to the intensity of a pixel as $\bar{I}_i(\mathbf{p}) = \text{dot}(\mathbf{I}_i(\mathbf{p}), (0.2990, 0.5870, 0.1140))$, defined as a weighted average of its RGB channels. We name N the number of input images. We use $\mathbf{W} = (1, 1, 1)^{\mathsf{T}}$ for the white color. We rely on angles between RGB vectors to reason about color saturation. We use the notation $\angle(\mathbf{C}_1, \mathbf{C}_2) = \arccos(\text{dot}(\mathbf{C}_1/\|\mathbf{C}_1\|, \mathbf{C}_2/\|\mathbf{C}_2\|))$ for the angle between the two colors $\mathbf{C}_1$ and $\mathbf{C}_2$. In several instances, we use a weighting function $w_i(\mathbf{p}) = \bar{I}_i(\mathbf{p})/(\bar{I}_i(\mathbf{p}) + \epsilon)$ that varies between 0 for low values of $\bar{I}_i(\mathbf{p})$ and 1 for high values. This function is useful to reduce the influence of dark pixels that are more noisy. In all our experiments, we use $\epsilon = 0.01$, assuming that the RGB channels range between 0 and 1.

3.1 Fill Light

The role of the fill light is to provide ambient lighting that gives approximately even illumination everywhere. This is the light that illuminates the shadows, i.e., it controls how dark they are. Since we assume that the input lights are roughly uniformly distributed, we could use the average of all the input images, $\frac{1}{N}\sum_i \mathbf{I}_i$. However, since the light distribution is not perfectly uniform, the average may exhibit some large intensity variations. We improve over this simple average by giving more importance to bright pixels using the w_i weights, which reduces the influence of dark noisy pixels:

$$\mathbf{I}_{\text{fill}}(\mathbf{p}) = \frac{\sum_i w_i(\mathbf{p})\, \mathbf{I}_i(\mathbf{p})}{\sum_i w_i(\mathbf{p})} \qquad (1)$$

where i is the index over all images. Figure 2 compares the simple average image to our actual fill light $\mathbf{I}_{\text{fill}}$ using the weighted average.

3.2 Edge Light

As discussed in Section 2.2, photographers often seek to emphasize the main edges and curved areas in the scene. A common approach is to position the lighting such that it creates a tonal variation around those regions of interest. In particular, highlights and shadows are among the main building blocks through which photographers achieve this [Hunter et al. 2011]. We define the *edge light* to assist them with this task. We proceed in two steps; first, we analyze the input images to identify the features to accentuate, and then we linearly combine the input images to emphasize the detected features, the mixture coefficients being a solution to an optimization problem that we define.

(a) Average image (b) Weighted average

Figure 2: *Compared to the average (a), the weighted average (b) produces more even illumination , which we use as a fill light.*

We define the features that we want to emphasize as edges in the input images that look persistent under the changing lighting conditions. Those can be due to geometric discontinuities in the scene, or more persistent highlights and shadows, which generate discontinuities in the observed images. However, occasional hard shadows and sharp highlights also produce image discontinuities but we are not interested in them, since creating highlights and shadows that compete with the main features of a scene is something photographers try to avoid [Hunter et al. 2011].

The key observation of our approach is that main edges of the scene are always located at the same place in the image, whereas discontinuities due to occasional highlights and shadows move depending on where the light source is. By observing the statistics at a given location, we can differentiate between a persistent scene feature and ephemeral edges due to occasional illumination effects. The former appears consistently in all images while the latter is only present once or a few times, i.e., it is an outlier. Our approach builds upon this observation and uses robust statistics to extract a map of the main scene features. We tested a few options such as computing the robust max or the median gradient at each pixel. However, we found that the solution that we present below performs better for our goal of emphasizing persistent edges. We compare against the robust max and median gradients in Figure 3.

Our approach uses the fact that edges due to occasional highlights and shadows have an inconsistent orientation. To exploit this phenomenon, at each pixel, we build the histogram of the gradient orientations. In practice, we use bins that span $5°$. To prevent dark noisy pixels from perturbing the process, we weight the contribution of each pixel using its w_i weight. Also, to differentiate between flat regions and vertical edges, small gradients of magnitudes less than 0.001 are handled separately. Then, within the largest bin, we pick the gradient of maximum amplitude. Intuitively, this process selects the strongest gradient that aligns with the most persistent orientation at every pixel. This gives us a target gradient map $\mathbf{G}$. We seek the *edge light* as a linear combination of the input images: $\mathbf{I}_{\text{edge}} = \sum_i \lambda_i\, \mathbf{I}_i$. To find the mixture coefficients λ_i, we minimize the following weighted least-squares energy function:

$$\arg\min_{\{\lambda_i\}} \sum_{\mathbf{p}} h(\mathbf{p}) \Big\| \nabla\Big(\sum_i \lambda_i\, \mathbf{I}_i(\mathbf{p})\Big) - \mathbf{G}(\mathbf{p}) \Big\|^2 \qquad (2)$$

where, the per-pixel weights $h(\mathbf{p})$ give more influence to pixels that have a peaked orientation histogram, that is, pixels that have a well-defined orientation. We define h by normalizing the histograms to 1 so that we can compare them across pixels, and picking the value of the largest bin at each pixel.

Discussion. The effect of our *edge light* is not to avoid all shadows and highlights, which would be undesirable from a photographic point of view. By optimizing for lighting that maximizes gradients that align with the main scene features, it favors highlights and shadows that align with them (Fig. 5a. 6). This behavior is reminiscent of the line drawing technique of Judd et al. [2007] who motivate their approach by characterizing the lines worth drawing as the ones that appear across multiple lighting configurations. From this perspective, our *edge light* seeks to produce an image in which discontinuities would be a good line drawing of the scene.

3.3 Diffuse Color Light

The objective of the *diffuse color light* is to emphasize the base color of objects. To reason about scene colors, we use a simple diffuse+specular image formation model in which the diffuse color can be arbitrary, and the specular color is the same as the light color.

First, we consider the case of a colorful object. We seek to design an energy function that favors images in which the diffuse component

(a) Robust max (b) Median (c) Ours

Figure 3: *We compare our gradient map against two other alternatives for emphasizing edges. Using the robust max gradients can produce results that emphasize distracting elements, like the shadows on the seat of the chair (a). The median gradients are more robust to occasional shadow boundaries, but other edges are also de-emphasized (b). In comparison, our proposed gradients better capture the main scene features and their orientations.*

is strong compared to the specular reflection. Similar to Tan et al. [2004] we observe that because the specular component is white, the stronger it is, the less saturated the observed color is. Formally, we consider a pixel $\mathbf{I} = d\mathbf{D} + s\mathbf{W}$ where $\mathbf{D}$ is the diffuse color and d its intensity, $\mathbf{W} = (1, 1, 1)^\intercal$ the white color, and s the specular intensity. We characterize the saturation by the angle $\angle(\mathbf{I}, \mathbf{W})$ between the observed color $\mathbf{I}$ and the white color $\mathbf{W}$. For a fixed d value, this angle decreases when s increases. This motivates the following energy term:

$$\arg\max_{\{\lambda_i\}} \sum_{\mathbf{p}} \hat{w}(\mathbf{p}) \angle\Big(\sum_i \lambda_i \mathbf{I}_i(\mathbf{p}), \mathbf{W}\Big) \quad (3)$$

where, $\hat{w}(\mathbf{p}) = \sum_i \lambda_i \mathbf{I}_i(\mathbf{p}) / (\sum_i \lambda_i \mathbf{I}_i(\mathbf{p}) + \epsilon)$ is a term that prevents selection of linear combination of lights that produce dark pixels, that tend to be noisier. With Equation 3, we seek a linear combination of input images that maximizes the angle with the white color, while preventing the selection of dark pixels.

This approach works well for colorful objects, that is, when $\angle(\mathbf{D}, \mathbf{W}) \gg 0$. However, this term is less effective for objects of neutral color, i.e., when $\mathbf{D} \approx \mathbf{W}$. For neutral colored objects, changes in specular intensity create only small angle variations. And most importantly, the optimization becomes sensitive to colored inter-reflections. For such neutral objects, even the small change of saturation generated by light reflecting off nearby colored objects has a significant impact on the energy value. In our experiments, using the previous energy term alone produced images in which gray objects have strong colored inter-reflections, which looked unpleasant. We address this issue with a second energy term based on the observation that the average image lowers the contribution of rare illumination effects, such as strong inter-reflections and highlights, which are undesirable features based on our *diffuse color light* definition. We design an energy term that encourages similarity between the average image and our result:

$$\arg\min_{\{\lambda_i\}} \sum_{\mathbf{p}} \angle(\sum_i \lambda_i \mathbf{I}_i, \frac{1}{N} \sum_i \mathbf{I}_i) \quad (4)$$

Since we seek to use this term only for neutral colors, else the solution will tend towards the average, we use a balancing term that equals 1 only for neutral colors and has lower values otherwise:

$$\alpha(\mathbf{p}) = \exp(-\angle(\frac{1}{N} \sum_i \mathbf{I}_i(\mathbf{p}), \mathbf{W})^2 / 2\sigma^2)) \quad (5)$$

with $\sigma = 0.5$. Figure 4 shows the significance of this term.

Putting the two terms together, and realizing that the goal is to maximize Equation 3, but minimize Equation 4, we obtain the coefficients of the *diffuse color light* $\mathbf{I}_{\text{diffuse}}$ by minimizing:

$$\arg\min_{\{\lambda_i\}} \sum_{\mathbf{p}} \Big[\alpha(\mathbf{p}) \angle(\sum_i \lambda_i \mathbf{I}_i(\mathbf{p}), \frac{1}{N} \sum_i \mathbf{I}_i(\mathbf{p})) - (1 - \alpha(\mathbf{p}))\ \hat{w}(\mathbf{p}) \angle(\sum_i \lambda_i \mathbf{I}_i(\mathbf{p}), \mathbf{W})\Big] \quad (6)$$

We minimize this function using an interior point method with finite differences to approximate the gradients (see Figures 4 and 5b).

(a) Saturation term only (b) Both terms

Figure 4: *The color of a neutral object, like the white ceiling, can be dominated by interreflections from nearby objects. We propose per-pixel weights that encourage more similarity between the average image and our result for pixels that look neutral on average.*

Summary. In summary, we have designed energy terms for each of the three basis lights: *fill light*, *edge light* and *diffuse color lights*. We solve for the linear combination of images, and their corresponding weights, that minimize the energy terms.

(a) Edge light (b) Diffuse color light (c) Fill light

Figure 5: *Our* edge light*, optimized for the red chair puts more emphasis on the main edges (a), whereas the* diffuse color light *reveals more of the deep red colors (b). We use the weighted average image as a* fill light *that provides more even illumination (c).*

4 Modifiers

In addition to the basis lights described in the previous section, we also introduce *modifiers* that alter the properties of these lights in ways that mimic standard practices in photography. The *per-object lighting modifier* restricts the light's extent, the *regional lighting modifier* balances the illumination intensity between different regions of the scene, and the *soft lighting modifier* modifies the lights so that they produce softer shadows and highlights.

4.1 Per-Object Lighting Modifier

This is the simplest of our modifiers. It is inspired by equipment like snoots that photographers use to control the spread of lights. We let users select objects in the image. Then, we compute the *fill*, *edge*, and *diffuse color lights* as described in the previous section. The only difference is that we only consider the pixels within the selected object. Users can then locally mix these three lights. To ensure smooth blending with the rest of the image, we apply a cross bilateral filter [Eisemann and Durand 2004; Petschnigg et al. 2004] to the binary selection, using the intensities of the current result as the guiding image. We use the fast cross bilateral filtering by Paris et al. [2009] to transform the binary selection into weighting masks that respect the edges for each of the basis lights. Then, in our interactive interface we approximate the weighting mask of the current combination of basis lights by linearly blending their corresponding masks. This produces a continuous mask that "snaps" at the main scene edges, which yields satisfying results. We also experimented with simple Gaussian blur but this generated severe halos, and also with multiscale blending [Burt and Adelson 1983] but color artifacts appeared.

4.2 Soft Lighting Modifier

This modifier aims for an effect akin to umbrellas and soft boxes, that is, simulating area light sources that produce soft shadows and highlights. Our strategy is to approximate an area light source by a set of nearby point sources. However, in our context, the position of the light sources is unknown a priori.

We address this problem with an approach inspired by Winnemöller et al. [2005] who showed that for two images taken from the same viewpoint with two different point lights, the spatial distance between the lights is correlated to the difference between the observed images: close light sources generate similar looking images and distant sources create different images. They demonstrate that this can be used to recover the positions of lights on a sphere. However, they mention that more general configurations are challenging.

For our *soft lighting modifier*, we build upon the same correlation between light position and image appearance, and sidestep the difficulties stemming from general light configurations by directly modifying the light mixture coefficients without explicitly recovering the light positions. We implicitly embed the input images into a weighted graph based on their similarity and apply a diffusion process on this graph to modify a given set of mixture coefficients $\{\lambda_i\}$ to approximate the effect of soft box lighting. We define a $N \times N$ matrix $\mathbf{S}$ with coefficients:

$$S_{ij} = \exp(-\|\mathbf{I}_i - \mathbf{I}_j\|^2 / 2\sigma_s{}^2) \tag{7}$$

and a vector $\Lambda = (\lambda_1, \ldots, \lambda_N)^\intercal$. Intuitively, multiplying Λ by $\mathbf{S}$ spreads the contribution of each light to the nearby sources using the image similarity $||\mathbf{I}_i - \mathbf{I}_j||$ as a proxy for the spatial distance. The σ_s parameter controls how far the intensity of each light is diffused. As is, this approach does not preserve the overall illumination intensity. We experimented with a few options and found that a simple global rescaling works well and is computationally inexpensive. To summarize, our *soft lighting modifier* is defined as:

$$soft_{\sigma_s}(\Lambda) = \frac{\|\Lambda\|}{\|\mathbf{S}\Lambda\|} \mathbf{S}\Lambda \tag{8}$$

To gain intuition, we observe two extreme σ_s settings. For $\sigma_s \to 0$, the modifier does nothing as one would expect, that is, point light sources remain as is. And for $\sigma_s \to \infty$, the output coefficients are all equal, i.e., the output is the average of the input images, which is in some sense the largest area light that we can simulate with our data. Other values of σ_s provide approximations to area light sources of intermediate sizes as shown in Figure 6c.

4.3 Regional Lighting Modifier

Photographers carefully balance light intensities in a scene to either emphasize a specific region or to do the opposite. We propose our *regional lighting modifier* to assist this process. We seek to provide a simple way to balance the lighting across the scene at a coarse level. Fine-grain adjustments can be made with our *per-object modifier*. Our observation is that the PCA decomposition of the input images extracts the main modes of variation of the illumination. In particular, for scenes that can be decomposed into "regions" illuminated independently of each other, e.g., foreground versus background, or left versus right, the first PCA component captures this structure well.

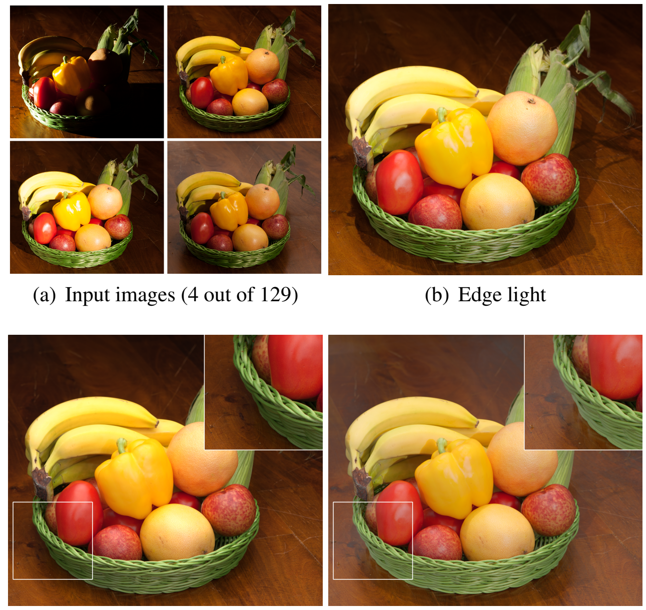

(a) Input images (4 out of 129) (b) Edge light

(c) With the soft lighting modifier (d) Adding fill light

Figure 6: *Basket. The* edge light *emphasizes the main edges of the scene, which tends to produce sharp highlights and deep hard shadows (b). Applying the* soft lighting modifier *softens the highlights and the shadow boundaries and keeps the shadows dark (c). The fill light has a complementary effect; it brightens the shadows and keeps their boundaries and the highlights sharp (d).*

We build our modifier upon this observation. Since PCA assumes an additive mode of variation, and light interaction with materials is multiplicative, we work in the log domain. Further, because we seek to only modulate pixel intensities without altering their color, we work with the intensity images $\{\bar{I}_i\}$. First, we estimate the first PCA component P of the log intensities $\{\ln(\bar{I}_i)\}$. To avoid perturbing the overall image intensity, we enforce a zero mean onto P by defining:

$$\hat{P} = P - \frac{1}{N} \sum_{\mathbf{p}} P(\mathbf{p}) \tag{9}$$

As is, $\hat{P}$ often exhibits undesirable shadow boundaries. We remove them by applying a cross bilateral filter [Paris and Durand 2009] to $\hat{P}$ with the current global result (a.k.a. the smooth blending of all locally and globally optimized basis lights) as the guiding image. Finally, we create a map $M = \exp(\beta \hat{P})$ where β is a user parameter controlling the magnitude of the effect: $\beta = 0$ does not alter the result, $\beta > 0$ emphasizes the regions where $\hat{P} > 0$ by making them brighter and the rest darker, and $\beta < 0$ has the opposite effect, i.e., it emphasizes the $\hat{P} < 0$ regions. The M map is multiplied pixel-wise to the current result to obtain the final result. The effect of this modifier is shown in Figure 7.

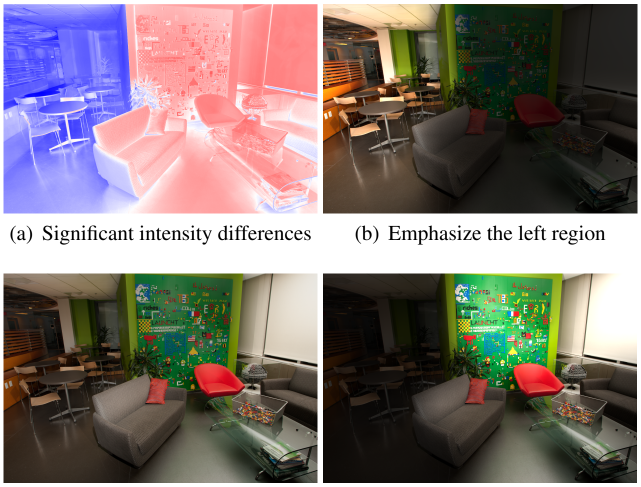

(a) Significant intensity differences (b) Emphasize the left region

(c) Original lighting ($\beta = 0$) (d) Emphasize the right region

Figure 7: *Our* regional lighting modifier *can be used to move the emphasis between the two regions that show significant intensities differences in the input images. We detect these regions automatically, by looking at the first PCA vector of the input intensities (a).*

5 Results

We now describe our implementation and the user interface of our prototype. Then, we demonstrate our approach on a variety of scenes, and show comparisons with related work.

5.1 Implementation

We use a combination of C++ and Matlab in our prototype system. The part that optimizes the basis lights is an offline process, which we implemented in Matlab, since it was less time critical. We use Matlab to solve the constrained linear and non-linear optimizations problems that correspond to our basis-light objectives. The timing for this step depends on the size of the regions, but it can take from a few seconds up to 10 minutes, when the region is the whole image. However, this offline process can be done in parallel for all pre-segmented objects and basis lights. In our user interface, in order to achieve interactive speeds, we do the image blending on the GPU.

In Table 1 we describe each scene, the resolution of the input images, the number of objects that we segmented, and the time for optimization, both per object and per image. We pre-compute the basis lights for the set of pre-segmented objects and the whole image so that users do not have to wait. Last two columns are max-per-object and full image timings.

5.2 User Interface

We now briefly describe our prototype interactive interface where users can explore a variety of lighting designs. For every object, the pre-computed basis lights (*edge*, *diffuse color*, *fill*) can be mixed by changing their relative contributions with sliders. Depending on the user's preferences, this can be done in two ways: (1) by preserving the overall intensity using normalized weights that add to one, or (2) by not normalizing. This is controlled by the checkbox "Keep intensities constant", visible in the supplementary video. In addition to that, we also let users control a simple exposure slider.

For every local object (selected by clicking on the object, and using the checkbox "Show local lighting"), the sliders control the object-based lighting. Segmenting the image into regions is beyond the scope of this paper, and we assume it is done by the user. When a region is first selected and the local lighting enabled, the initial exposure is set to match that of the global lights.

Users can also interactively change the strength of the *soft lighting modifier* to control the softness of the shadows and highlights (Fig. 6c). To ensure a consistent behavior across scenes, we normalize the image differences in Equation 7 so that the largest one is 1. To enable interactive editing, for each light, we precompute the effect of the modifier for 100 regularly spaced values of σ_s between 0.01 and 10. At run-time, we linearly interpolate the values of the two samples that are closest to the requested parameter.

Our last slider controls the regional modifier. This allows users to interactively change the emphasis in the scene, by smoothly modifying the per-pixel exposure, through the parameter β (Fig. 7).

5.3 Image Results

We now demonstrate our results on a range of scenes.

Cafe. In Figure 1, we show results for a larger interior scene that has a variety of objects with different shapes and materials. For example, the red chair has strong glossy components, and at the same time, a deep red color. Our *edge light* better reveals the shape of the chair, by emphasizing highlights (Fig. 5a). Our *diffuse color light* shows more of the deep red color of the chair (Fig. 5b). Our system allows novice users to easily explore different lighting designs, by mixing the basis lights in various proportions, globally or per pre-segmented objects (Fig. 1c). In the supplemental material we show 6 more results, which demonstrates that even novice users can produce non-trivial variations using our system.

Library. The Library in Figure 8 shows an example of an indoor room, where an outside view is also visible through a window. We gave our data set to a professional photographer, with the instructions to create a result that looks good to him. Figure 8b shows his result, achieved in 20 minutes, using the full power of Photoshop. We gave the same scene to novice users, who were able to explore various lighting designs, using our system. For example, in Figure 8c our user used more *diffuse color light* and locally decreased the exposure to better show the black color of the sofa. *Diffuse color light* was also used to emphasize the deep red of the armchair. Other results available in the supplemental material show that even novice users are able to produce nontrivial variations.

House. Figure 9, third row, shows an example of a big outdoor scene. Lighting this scene by walking around with a single light is particularly useful and one of the few options at this time of the day, when the ambient lighting is low. In the supplemental material we show the two regions found by our *regional lighting modifier*. For this scene, the regions roughly correspond to objects closer to the ground, and objects closer to the sky, which exhibit significant intensity differences in the input data set. Users of our system can use this for an artistic control to add more contrast between the lighting in those two regions. In the supplemental video we show all other steps used to generate this result.

Basket. Figure 6b shows the result of our *edge light*, applied to the entire scene. One of the main edges in this scene separate the foreground objects from the background table, and our *edge light* emphasized those further (Fig. 6b). However, in this cluttered scene, occasional shadows can arise for many lighting positions, producing distracting shadows in the *edge light* solution. In our system, users can interactively apply the *soft lighting modifier* to simulate a larger area light source. This can be used to soften the hard shadows on the table, cast by the fruits. The *soft lighting modifier* can also be used to soften the highlights on the tomato (Fig. 6c). A *fill light* can be used to add even illumination to the whole scene, which provides more details in dark regions, including shadows (Fig. 6d). We want to stress the difference between the *soft lighting modifier* and the *fill light* Although at their extreme values they both produce some version of the average image, their intermediate effect is different. The *soft lighting modifier* simulates a gradual increase of the current light source, by combining nearby lights and weighting them appropriately, whereas the fill light cross fades between the current illumination and the average image.

Los Feliz. In Figure 10, we show results for a stack of images that we received from a professional photographer (naive to our research). This is a set that was not captured or processed by us. The photographer also gave us his preferred final result for that scene. We demonstrate in the supplemental video that we were able to achieve a similar effect in a few seconds, rather than half an hour. In particular, our *edge light* applied to the whole scene has captured the gist of the result produced by the professional.

5.4 Dependency on the Input Images

In Figures 11 and 12 , we evaluate the dependency of our system on the number and quality of the input images. Figure 11, row 1 shows the results of our *edge light* optimized for 11 different segments, such as the red chair, the central sofa, the green wall, etc. In the supplemental material we show all pre-segmented regions. Figure 11, row 2 shows the results of our *diffuse color light*, optimized for the same set of segments. We conducted two tests based on the image selection. In the first test, we randomly select 5 and then 15 images from the original data set, (Fig. 11b,d). In the second test, 5 and then 15 images were carefully selected, so that they contain useful features such as lighting that emphasizes the edges, or lighting that reveals the underlying material, (Fig. 11c,e). First, our evaluation suggests that a small number of random shots are insufficient to produce good features across all parts of this large scale scene, (Fig. 11b,d). Second, even if carefully chosen, if the number is too small (5) it is not enough for a scene of this size, (Fig. 11c). Finally, post-hoc it was possible for us to find 15 images that would produce reasonable basis lights, (Fig. 11e). So, it might be possible for a person with a lot of experience to make use of our basis lights with a smaller number of carefully planned shots. However, even experienced photographers use the earlier work flow (capturing many images) because they are worried they could miss something and do not want to take chances. Further, in Figure 12 we show that the quality of our *soft lighting modifier* is more closely related to the number of input images. The reason is that a more uniform sampling of the lighting in the scene produces more close-by lights. These are needed for the gradual simulation of large area lights, computed by our *soft lighting modifier*.

5.5 Comparison with Related Work

In Figure 9, we compare with two other systems that share some common goals with ours. Exposure Fusion [Mertens et al. 2007] expects a sequence of images with different exposures, and produces a single, well-exposed image. Their system was primarily designed for scenes with constant lighting and high dynamic range. In comparison, our input image sequences with dynamic lighting moving around the scene introduce new challenges. When applied to our datasets, Exposure Fusion produces unsatisfactory results (Fig. 9, second column). Further, their results appear flat since they seek to expose the entire scene, including the shadow areas, equally well.

Figure 9, third column, shows results from MLIC [Fattal et al. 2007], which is more closely related to our goals, as they work on similar input data: a static scene under dynamic lighting. However, they propose an automatic system, that emphasizes details over different scales. As a result, they also tend to flatten the look of the image, and decrease its realism by mixing information from different scales, under different lighting. While this works well on single objects as demonstrated in the original paper, it is less successful on large scenes. In comparison, we generate plausible pleasant lights, based on common photography practices.

Scene	Size	Images	Objects	Max (min)	Full (min)
Cafe	$1.5MP$	112	11	4	9
Library	$1.0MP$	83	13	3	8
Basket	$1.2MP$	129	9	1	9
House	$1.5MP$	149	6	3	10
Sofas	$1.5MP$	32	7	0.5	2
Kitchen	$1.5MP$	127	7	4	10

Table 1: *Number of regions and time for optimizing our set of basis lights on those regions and the full scene.*

5.6 User Validation

To validate our contribution we evaluated our idea with novice users and expert photographers.

For the expert evaluation, we asked 3 experienced photographers to comment on the merits of our basis lights. We sent them two Photoshop projects: one containing the full stack of input images, and a second project with a reduced set of images that represent our basis lights, optimized for a few pre-segmented objects. We asked each of them to spend some time working with both image stacks. The goal of this study was to see whether our basis lights could give them a good starting point, while saving them time wasted in finding and blending features from the original images. They were all enthusiastic about the lights, and reported that working with the reduced image stack made their workflow more efficient, compared to going through tens of input images. They also reported that our basis lights would be useful in their workflow, but sometimes they had to apply additional adjustments, like color balance and light levels to achieve a desired effect. However, this type of adjustments is orthogonal to our work. See the supplemental material for their results using the full image stack and our reduced image stack.

We also tested our system with 7 novice users who had little or no experience with photography. The goal of this study was to show that our prototype system can allow ordinary users to explore different options and achieve sophisticated lighting designs in a couple of minutes. Figure 8b shows the result of a professional photographer using Photoshop with our reduced image stack. Figure 8c shows results on the same scene generated by a novice user interacting with our system for the first time. The task was not to match the solution of the professional, as Photoshop allows them to apply many adjustments, like nonlinear curves to increase contrast, which is orthogonal to our project. However, we showed the professional's result to our users, just for a minute, as an example of a good lighting design. We then asked them to explore solutions in an open-ended manner. Our system allowed users to rapidly explore different lighting designs and produce visually pleasing results (Fig. 8c). We show results of 6 other users in the supplemental material. On average, users spend 15 minutes on this data set.

In the supplementary material, we show another evaluation on the "Cafe" scene, where we gave users 5 to 10 minutes to explore different lighting designs, using our system. Then, we asked them to switch to Photoshop and spend the same amount of time, using the full, unprocessed image stack. They were instructed to look for similar features as the one they produced using our system. Our experiment showed that for novice users the results were poor; searching and blending features from the full stack of input images is a nontrivial task which prevents them from producing good results.

(a) Library, 4 out of 83 images (b) Result of a professional (20min in Photoshop) (c) Novice user, using our system for 16min

Figure 8: *Our evaluation with novice users shows that our system allows them to explore non-trivial lighting designs (c) in a short amount of time, comparable in quality to what a professional photographer would achieve, using the full power of Photoshop in 20 minutes (b).*

5.7 Discussion and Limitations

Our approach is a user-driven creation process meant to help users create compelling images. However, not all slider configurations produce such images, e.g., if one uses only the fill light to illuminate the scene, the result will look dull. That said, our experiments show that even novice users are able to generate quality results.

In general, our results do not correspond to a physical setup. For instance, our *regional lighting* and *per-object modifiers* do not alter the illumination in a physical way. However, they are close to what could be produced using blockers, and our results look plausible. In addition, the core of our approach is based on linear combinations of the input lights, which corresponds to actually turning on the lights at the same time with the appropriate intensity. This further contributes to generating plausible images.

Finally, image editing software offers virtually infinite control over the result whereas our approach covers a smaller design space. However, we argue that for most users, unbounded editing capability is actually a hindrance more than a help since it requires advanced skills and a lot of time, which is confirmed by our experiments. For the few users with mastery of advanced editing tools, we do not offer a complete set, but we envision that they would first use our approach to quickly obtain a satisfying result and, if needed, they would later refine it with standard photo editing software.

6 Conclusion and Future Work

Lighting is critical to good photography. There is a new workflow emerging for lighting static scenes, where photographers capture many images of the scene with a single (or small set of) light(s) in different locations to create a set of input images that serve as data to a later compositing stage. This workflow gives the photographers a lot of flexibility in post-process, and is faster in time-constrained shots, but dealing with the many images after that is difficult.

We have introduced a set of optimizations to help the photographers assemble all these images into a small set of basis lights and modifiers that are easy and fast to user in an interactive editing session. The photographer can use these lights to achieve common photography goals like accentuating highlights, filling shadows, and emphasizing object color. Our studies with both novice and professional users shows that this approach is a significant improvement over the traditional workflow.

There are multiple areas of future work. One possibility is to explore other types of basis lights related to common photography practices, such as rim lighting that emphasizes contours, or lighting that better reveals the glossy behavior of objects. Another interesting avenue for future work would be the development of an interactive system that could assist the acquisition process, by guiding the placement of lights that achieve a desired effect. Finally, we would like to explore better optimization techniques for the basis light objectives, like multi-grid on GPU, which could make this step interactive.

Acknowledgments. We would like to thank the anonymous reviewers for their constructive comments. We would like to acknowledge our funding agencies NSF IIS 1161645 and NSF 1011919, and funding from Adobe. We thank the experienced photographers Michael Kelley, Daniela Steinsapir and Michaël Gharbi, and all the participants of our user study.

References

ADELSON, E. H., AND BERGEN, J. R. 1991. The plenoptic function and the elements of early vision. In *Computational Models of Visual Processing*, MIT Press.

AGARWALA, A., DONTCHEVA, M., AGRAWALA, M., DRUCKER, S., COLBURN, A., CURLESS, B., SALESIN, D., AND COHEN, M. 2004. Interactive digital photomontage. *ACM Trans. Graph.*.

AKERS, D., LOSASSO, F., KLINGNER, J., AGRAWALA, M., RICK, J., AND HANRAHAN, P. 2003. Conveying shape and features with image-based relighting. In *Proceedings of the 14th IEEE Visualization 2003 (VIS'03)*, IEEE Computer Society.

BISHOP, T., AND FAVARO, P. 2011. The light field camera: Extended depth of field, aliasing and super-resolution. *IEEE Trans. Pattern. Anal. Mach. Intell.*.

BOUSSEAU, A., PARIS, S., AND DURAND, F. 2009. User-assisted intrinsic images. *ACM Trans. Graph. 28*, 5 (Dec.).

BOUSSEAU, A., CHAPOULIE, E., RAMAMOORTHI, R., AND AGRAWALA, M. 2011. Optimizing environment maps for material depiction. In *CGF*, Eurographics Association, EGSR'11.

BOYADZHIEV, I., BALA, K., PARIS, S., AND DURAND, F. 2012. User-guided white balance for mixed lighting conditions. *ACM Trans. Graph. 31*, 6 (Nov.).

BURT, P. J., AND ADELSON, E. H. 1983. A multiresolution spline with application to image mosaics. *ACM Trans. Graph. 2*, 4.

CARROLL, R., RAMAMOORTHI, R., AND AGRAWALA, M. 2011. Illumination decomposition for material recoloring with consistent interreflections. *ACM Trans. Graph. 30*, 4 (July).

COHEN, M. F., COLBURN, R. A., AND DRUCKER, S. 2003. Image stacks. Tech. rep., Microsoft Research. MSR-TR-2003-40.

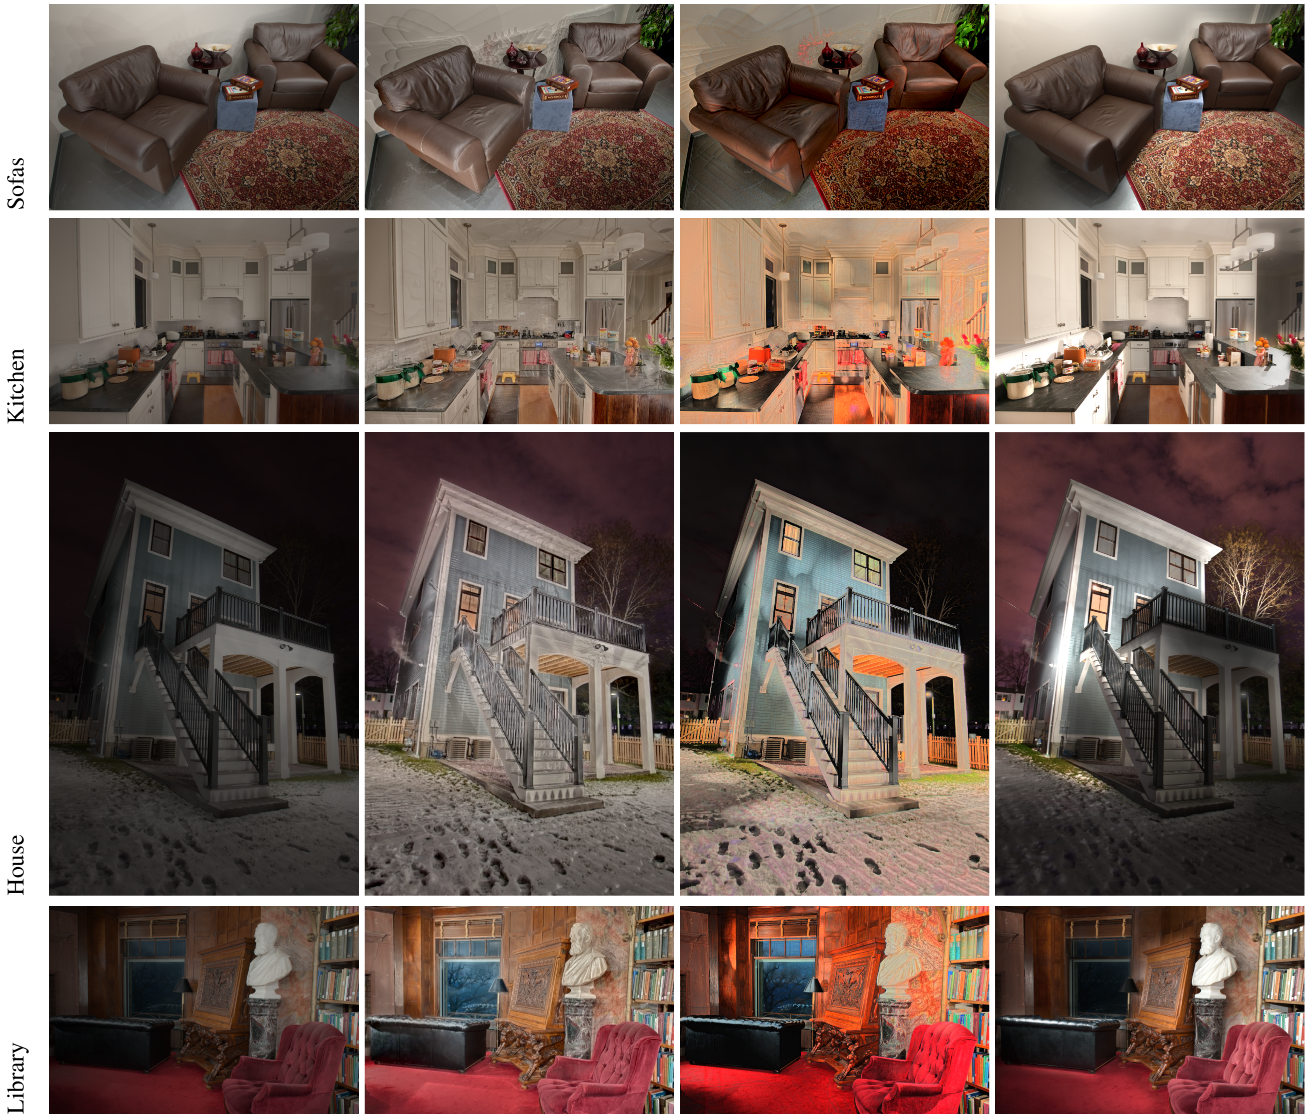

(a) Average (b) Exposure fusion (c) MLIC (d) Ours

Figure 9: *Comparison of our user-driven system to two automatic systems, that share some goals with ours, shows that those systems produce less satisfying results for the type of scenes in which we are interested. Further, the goal of our system is to provide a simple set of controls, through which users can explore a variety of solutions, whereas these systems are mostly automatic.*

(a) 4 out of 30 images (b) Result by professional (30min in Photoshop) (c) Result, using our system for 30sec

Figure 10: *On a data set provided by a professional photographer (© Michael Kelley), our system provides a quick way to explore a reasonable solution, that captures the overall look and feel of the professional result. In a supplemental video we show that combination of our global* edge *and* fill lights *capture many of the desired features, like background that is well lit and fill light on the table and chairs.*

(a) All images (b) 5 random (c) 5 planned (d) 15 random (e) 15 planned

Figure 11: *Our* diffuse color *and* edge lights *optimized for 11 regions, using input data with different properties. In (a) we used the original data set (all images). In (b) and (d) we show the quality of the results that could be obtained with 5/15 randomly selected images. In (c) and (e) we demonstrate the results of our system with 5/15 carefully selected pictures. First, note that a small number of randomly selected images produces unsatisfactory results for many parts of the scene, (b) and (d). Second, even if carefully chosen, if the number is too small (5) it is not enough for a scene of this size, (c). In (e) we show that with 15 carefully selected images our basis lights can produce reasonable results.*

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 ACM SIGGRAPH 2000*.

EISEMANN, E., AND DURAND, F. 2004. Flash photography enhancement via intrinsic relighting. *ACM Trans. Graph. 23*, 3.

FATTAL, R., AGRAWALA, M., AND RUSINKIEWICZ, S. 2007. Multiscale shape and detail enhancement from multi-light image collections. In *ACM SIGGRAPH 2007 papers*.

GUANZON, J., AND BLAKE, M. 2011. Video of computational design workflow (https://http://vimeo.com/30363913). Tech. rep.

HUNTER, F., FUQUA, P., AND BIVER, S. 2011. *Light Science and Magic 4/e*. Elsevier Science.

JUDD, T., DURAND, F., AND ADELSON, E. 2007. Apparent ridges for line drawing. *ACM Transactions on Graphics 26*, 3.

KELLEY, M. P. 2011. Video of computational design workflow (https://www.youtube.com/watch?v=J-exuHchmSk). Tech. rep.

KELLEY, M. P. 2012. Private communication with proffesional photographer. Tech. rep., Private Photographer.

LEVIN, A., FERGUS, R., DURAND, F., AND FREEMAN, W. T. 2007. Image and depth from a conventional camera with a coded aperture. In *ACM SIGGRAPH 2007 papers*, ACM.

MALLICK, S. P., ZICKLER, T., BELHUMEUR, P. N., AND KRIEGMAN, D. J. 2006. Specularity removal in images and videos: a pde approach. In *Proceedings of the 9th European conference on Computer Vision - Volume Part I*, Springer-Verlag, ECCV'06.

MERTENS, T., KAUTZ, J., AND REETH, F. V. 2007. Exposure fusion. In *Proceedings of the 15th Pacific Conference on Computer Graphics and Applications*, IEEE Computer Society.

NG, R., LEVOY, M., BRÉDIF, M., DUVAL, G., HOROWITZ, M., AND HANRAHAN, P. 2005. Light Field Photography with a Hand-Held Plenoptic Camera. Tech. rep., Apr.

OSTROVSKY, Y., CAVANAGH, P., AND SINHA, P. 2005. Perceiving illumination inconsistencies in scenes. *Perception 34*.

PARIS, S., AND DURAND, F. 2009. A fast approximation of the bilateral filter using a signal processing approach. *International Journal of Computer Vision 81*, 1.

PELLACINI, F. 2010. Envylight: an interface for editing natural illumination. *ACM Trans. Graph. 29* (July).

PETSCHNIGG, G., SZELISKI, R., AGRAWALA, M., COHEN, M., HOPPE, H., AND TOYAMA, K. 2004. Digital photography with flash and no-flash image pairs. In *ACM SIGGRAPH 2004 Papers*.

RASKAR, R., TAN, K.-H., FERIS, R., YU, J., AND TURK, M. 2004. Non-photorealistic camera: depth edge detection and stylized rendering using multi-flash imaging. In *ACM SIGGRAPH 2004 Papers*, ACM, New York, NY, USA, SIGGRAPH '04.

REINHARD, E., POULI, T., KUNKEL, T., LONG, B., BALLESTAD, A., AND DAMBERG, G. 2012. Calibrated image appearance reproduction. *ACM Trans. Graph. 31*, 6 (Nov.).

SCHOENEMAN, C., DORSEY, J., SMITS, B., ARVO, J., AND GREENBERG, D. 1993. Painting with light. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, SIGGRAPH '93.

TAN, R., NISHINO, K., AND IKEUCHI, K. 2004. Separating reflection components based on chromaticity and noise analysis. *IEEE Trans. Pattern. Anal. Mach. Intell.*.

WINNEMÖELLER, H., MOHAN, A., TUMBLIN, J., AND GOOCH, B. 2005. Light waving: Estimating light positions from photographs alone. *Computer Graphics Forum 24*, 3.

(a) Initial lighting (b) Using 15 images (c) Using all images

Figure 12: *Our* soft lighting modifier, *applied to the lighting in (a) shows that a small number of images (15), even if carefully selected, do not provide enough close-by lights, needed for the gradual simulation of large area lights, (b). In (c) we show that the effect of our* soft lighting modifier *improves when used with the full image set.*