

Homework assignment 3

15-463, 15-663, 15-862 Computational Photography, Fall 2026
Carnegie Mellon University

Due Friday, Oct. 9, at 11:59pm ET

The purpose of this assignment is to explore photography with flash/no-flash pairs. In general, getting the lighting in your scene right is one of the most important considerations one needs to take care of when taking a photograph. This can be particularly challenging when the scene you are trying to photograph has a large dynamic range, including both very dark and very bright objects. Previously, we saw a *passive* technique for photographing such difficult scenes, using HDR imaging. Here we will investigate an alternative *active* technique, where we insert new light into the scene using our camera’s flash.

As we discussed in class, using flash has advantages and disadvantages. The extra light reduces noise, allows us to use larger shutter speeds to reduce motion blur, and can help illuminate the scene more uniformly. On the flip side, the use of the harsh flash lighting can ruin the scene’s ambiance, produce highlights, and overall result in unappealing appearance. These issues commonly arise when using simple point flash lights, e.g., the one on the class DSLR or at the back of your cell phone. A photographer with access to some more advanced equipment would use either a pointable flash that they can bounce off of the ceiling, or they would use a diffuse display to make the light smoother. However, for the purposes of this assignment, we can assume we do not have access to such equipment.

Instead, we can attempt to produce images that combine the advantages of flash and ambient illumination by capturing two images, one with each illumination, and then *fusing* them into one. In class we saw two types of techniques for this kind of *flash/no-flash photography*, one based on bilateral filtering, and another based on gradient-domain processing. In this homework, you will implement both of these techniques, and apply them to photos provided by us, as well as photos you capture on your own.

Throughout the assignment, we refer to a number of key papers that were also discussed in class. Even though the assignment and class slides describe most of the steps you need to perform, we highly recommended that you read the associated papers.

Towards the end of this document, you will find a “Deliverables” section describing what you need to submit. Throughout the writeup, we also mark in **red** questions you should answer in your submitted report. Lastly, there is a “Hints and Information” section at the end of this document that is likely to help. We strongly recommend that you read that section in full before you start to work on the assignment.

Notation. Throughout this assignment we use bold font to denote image-sized quantities, and regular font to denote scalars. We use regular multiplication, division, squaring, absolute value, etc., notation involving image-sized quantities to denote element-wise operations. We use $\langle \cdot, \cdot \rangle$ to indicate the scalar resulting from the vector dot product between *vectorized* versions of the image-sized quantities.

1. Bilateral filtering (100 points)



Figure 1. From left to right: Photos taken under ambient lighting and flash lighting; denoising result using basic bilateral filtering; and final denoising result.

In the first part of the homework, you will use the lamp flash/no-flash images from the paper by Petschnigg et al. [6], available in the `./data/lamp` directory of the homework ZIP archive. Your focus will be on *denoising*

the ambient image using bilateral filtering. Figure 1 shows representative results for this section. As we discussed in class, bilateral filtering is a form of “blurring” that attempts to respect sharp edges in the image. This is achieved by weighting neighboring pixels not just based on their distance from the pixel that is being filtered, but also based on their similarity to it. This can help reduce noise in an image, without overly blurring details. Further improvements can be obtained by using bilateral filtering techniques that combine information from both the flash and ambient illumination images, as proposed by Petschnigg et al. [6]. Figure 3 of that paper is particularly helpful for navigating this homework assignment.

Implement bilateral filtering (40 points). As a first attempt at denoising, you can simply apply the bilateral filter on the ambient image. Using the notation of Petschnigg et al. [6], we will denote the ambient image $\mathbf{A}$ and the denoised image $\mathbf{A}^{\text{Base}}$. You will implement bilateral filter as described in the “PiecewiseBilateral” algorithm of Section 5 of the paper by Durand and Dorsey [3].

Implement joint-bilateral filtering (30 points). You will now attempt to get a better result by using *joint bilateral filtering*, as described in Section 4.1 of Petschnigg et al. [6], to produce the image $\mathbf{A}^{\text{NR}}$ in the paper’s notation. As we discussed in class, joint bilateral filtering involves computing the intensity kernel using the flash image—see also Equation (4) of Petschnigg et al. [6].

Implement detail transfer (20 points). While joint bilateral filtering can help avoid overblurring the ambient image, it cannot *add* new detail to it. You will enhance the detail of your result by implementing the flash-to-ambient detail transfer procedure described in Section 4.2 of Petschnigg et al. [6]. For this, you will need to apply basic bilateral filtering to the flash image. Following Petschnigg et al. [6], we denote the flash image $\mathbf{F}$ and the result of filtering $\mathbf{F}^{\text{Base}}$. Then, you can form a new estimate for the ambient image by combining $\mathbf{F}$, $\mathbf{F}^{\text{Base}}$, and the image $\mathbf{A}^{\text{NR}}$ from before, as:

$$\mathbf{A}^{\text{Detail}} = \mathbf{A}^{\text{NR}} \frac{\mathbf{F} + \epsilon}{\mathbf{F}^{\text{Base}} + \epsilon}, \quad (1)$$

where ϵ is some small value you can select on your own.

Implement shadow and specular masking (10 points). Finally, it is time to create the mask $\mathbf{M}$ described in Section 4.3 of Petschnigg et al. [6]. You will use this mask to detect regions in the flash image that have shadows or specularities not present in the ambient image. You can generate the mask using the simple thresholding operations the paper describes. Once you have the mask, form a final estimate $\mathbf{A}^{\text{Final}}$ of the ambient image as

$$\mathbf{A}^{\text{Final}} = (1 - \mathbf{M}) \cdot \mathbf{A}^{\text{Detail}} + \mathbf{M} \cdot \mathbf{A}^{\text{Base}}. \quad (2)$$

Implementation details. These details apply to all four of the previous steps. As we discussed in class, the bilateral filter takes as input two parameters, the standard deviation σ_s of the spatial Gaussian kernel, and the standard deviation σ_r of the intensity Gaussian kernel. You will need to experiment with different values of these parameters to find which ones give the most aesthetically pleasing results. A good range to experiment with is $\sigma_s \in [1, 64]$ and $\sigma_r \in [0.05, 0.25]$. Note that these values assume that the maximum value in your image is 1, so make sure to normalize first.

In general, the best parameters will be different for each of the four types of filtering you will implement: basic bilateral filtering, joint bilateral filtering, denoising with detail transfer, and mask-based merging. **In your report, compare the results of the four algorithms for different sets of parameter values. For these comparisons, make sure to show difference images $\text{image}_1 - \text{image}_2$, which should help highlight where in the images the techniques perform differently. Report the parameter values you chose as best for the four types of filtering, and discuss the advantages and disadvantages of each type.**

2. Gradient-domain processing (100 points)

In the second part of the homework, you will use the museum flash/no-flash images from the paper by Agrawal et al. [1], available in the `./data/museum` directory of the homework ZIP archive. Figure 2(a-b)



Figure 2. Combining photographs taken under ambient lighting (left) and flash lighting (middle), to produce a new image that combines the best of both worlds (right).

shows the two images. As you can see in the ambient image, the background painting is nicely illuminated from the ambient lighting. However, the person in the foreground is very dimly lit. We can attempt to address this using the camera's flash, which results in the flash photograph. Now the person in the foreground is well-illuminated, but our photograph suffers from a different problem: The sharp lighting of our flash has created a specular highlight on the painting. In this part of the homework, you will focus on fusing the flash and ambient images, in order to produce uniform illumination without the specularity and other artifacts. Figure 2(c) shows the end result of the fusion. You will do this fusion using gradient-domain processing, as proposed by Agrawal et al. [1]. As usual, we strongly encourage you to read through the paper, to find many implementation details and insights about the various parameters.

Differentiate and then re-integrate an image (75 points). This task will not give you any new images, but it should help you write and test the Poisson solver you will use next to create fused flash/no-flash images.

Take any image you want (e.g., the ambient image from the museum pair), which we will denote as $\mathbf{I}$. The gradient field $\nabla \mathbf{I}$ of this image will be a two-channel image, where one channel is the per-pixel partial- x derivatives $\mathbf{I}_x$, and the other channel is the per-pixel partial- y derivatives $\mathbf{I}_y$. You will attempt to integrate this gradient field to form a new image, $\mathbf{I}^*$, that should ideally be the same as the original image $\mathbf{I}$.

For this, you should first compute the divergence of the gradient field $\nabla \mathbf{I}$,

$$\text{div}(\nabla \mathbf{I}) = \nabla \cdot \nabla \mathbf{I} = \mathbf{I}_{xx} + \mathbf{I}_{yy}, \quad (3)$$

which you can do by using Laplacian filtering on the original image $\mathbf{I}$. Then, you will use this divergence as input to your implementation of the Poisson solver. You will implement the conjugate gradient descent (CGD) algorithm we discussed in class. We recommend that you read through Sections 8 and B2 of Shewchuk's [7] excellent (and funny) discussion of CGD. Algorithm 1 shows the CGD procedure you should implement. In the algorithm, `LaplacianFiltering` is convolution with the Laplacian kernel, $\mathbf{B}$ is a binary image-sized mask that equals 0 on the boundary and 1 elsewhere, and $\mathbf{I}_{\text{boundary}}^*$ is an image equal at the boundary to the values used by the boundary conditions.

For initialization, set $\mathbf{I}_{\text{init}}^*$ equal to the zero image. For Dirichlet boundary conditions, set the outermost pixels of $\mathbf{I}_{\text{boundary}}^*$, at all four image edges, to equal the corresponding pixels of the original image $\mathbf{I}$. You can experiment with different values for the parameters ϵ and N controlling the convergence of your algorithm. If your implementation is correct, then the final image $\mathbf{I}^*$ should be equal to the original image $\mathbf{I}$.

Create the fused gradient field (25 points). In this part, we will follow the notation of Agrawal et al. [1], to make it easier to compare with the paper. Note that some quantities reuse symbols from the bilateral filtering question, but *are not necessarily the same*. We will denote the ambient image as $\mathbf{a}$, the flash image as Φ' , and the fused image as Φ^* . The first step in the fusion process requires forming the gradient field that you will use to create Φ^* . In turn, you will create this gradient field using the gradient fields $\nabla \mathbf{a}$ and $\nabla \Phi'$ of the ambient and flash image, respectively.

ALGORITHM 1: Gradient field integration with conjugate gradient descent.

Input: target divergence $\mathbf{D}$, initialization $\mathbf{I}_{\text{init}}^*$, boundary mask $\mathbf{B}$, boundary values $\mathbf{I}_{\text{boundary}}^*$, convergence parameter $\epsilon > 0$, iteration number N .

Output: integrated image $\mathbf{I}^*$.

```

/* Initialization */
1  $\mathbf{I}^* \leftarrow \mathbf{B} \cdot \mathbf{I}_{\text{init}}^* + (1 - \mathbf{B}) \cdot \mathbf{I}_{\text{boundary}}^*$ ;
2  $\mathbf{r} \leftarrow \mathbf{B} \cdot (\mathbf{D} - \text{LaplacianFiltering}(\mathbf{I}^*))$ ;
3  $\mathbf{d} \leftarrow \mathbf{r}$ ;
4  $\delta_{\text{new}} \leftarrow \langle \mathbf{r}, \mathbf{r} \rangle$ ;
5  $n \leftarrow 0$ ;
/* Conjugate gradient descent iteration */
6 while  $\sqrt{\langle \mathbf{r}, \mathbf{r} \rangle} > \epsilon$  and  $n < N$  do
7    $\mathbf{q} \leftarrow \text{LaplacianFiltering}(\mathbf{d})$ ;
8    $\eta \leftarrow \frac{\delta_{\text{new}}}{\langle \mathbf{d}, \mathbf{q} \rangle}$ ;
9    $\mathbf{I}^* \leftarrow \mathbf{I}^* + \mathbf{B} \cdot (\eta \cdot \mathbf{d})$ ;
10   $\mathbf{r} \leftarrow \mathbf{B} \cdot (\mathbf{r} - \eta \cdot \mathbf{q})$ ;
11   $\delta_{\text{old}} \leftarrow \delta_{\text{new}}$ ;
12   $\delta_{\text{new}} \leftarrow \langle \mathbf{r}, \mathbf{r} \rangle$ ;
13   $\beta \leftarrow \frac{\delta_{\text{new}}}{\delta_{\text{old}}}$ ;
14   $\mathbf{d} \leftarrow \mathbf{r} + \beta \cdot \mathbf{d}$ ;
15   $n \leftarrow n + 1$ ;
16 end

```

First, compute $\nabla \mathbf{a}$ and $\nabla \Phi'$. As with $\nabla \mathbf{I}$, these are two-channel images, with one channel corresponding to partial- x and the other partial- y derivatives. Then, use these gradients to implement the following steps:

1. Compute the *gradient orientation coherency map* $\mathbf{M}$ using Equation (5) of Agrawal et al. [1],

$$\mathbf{M} = \frac{|\nabla \Phi' \cdot \nabla \mathbf{a}|}{\|\nabla \Phi'\| \|\nabla \mathbf{a}\|} = \frac{|\Phi'_x \cdot \mathbf{a}_x + \Phi'_y \cdot \mathbf{a}_y|}{\sqrt{\Phi'^2_x + \Phi'^2_y} \sqrt{\mathbf{a}^2_x + \mathbf{a}^2_y}}. \quad (4)$$

Note that, in this equation, both the numerator and denominator are computed pixel-wise, and the mask $\mathbf{M}$ is image-sized.

2. Compute the pixel-wise *saturation weight map* $\mathbf{w}_s$ from the flash image Φ' using Equation (11) of Agrawal et al. [1],

$$\mathbf{w}_s = \tanh(\sigma \cdot (\Phi' - \tau_s)). \quad (5)$$

Normalize the saturation map to be in the range $[0, 1]$.

3. Compute a new gradient field (two-channel image) for the fused image Φ^* using Equation (12) of Agrawal et al. [1],

$$\nabla \Phi^* = \mathbf{w}_s \cdot \nabla \mathbf{a} + (1 - \mathbf{w}_s)(\mathbf{M} \cdot \nabla \Phi' + (1 - \mathbf{M}) \cdot \nabla \mathbf{a}), \quad (6)$$

or equivalently,

$$\begin{bmatrix} \Phi^*_x \\ \Phi^*_y \end{bmatrix} = \begin{bmatrix} \mathbf{w}_s \cdot \mathbf{a}_x + (1 - \mathbf{w}_s)(\mathbf{M} \cdot \Phi'_x + (1 - \mathbf{M}) \cdot \mathbf{a}_x) \\ \mathbf{w}_s \cdot \mathbf{a}_y + (1 - \mathbf{w}_s)(\mathbf{M} \cdot \Phi'_y + (1 - \mathbf{M}) \cdot \mathbf{a}_y) \end{bmatrix}. \quad (7)$$

Experiment with different values for the scalar parameters σ and τ_s , starting with the values recommended in the paper. Show the gradient fields $\nabla \mathbf{a}$, $\nabla \Phi'$, and $\nabla \Phi^*$, and mention the parameter values you used.

Finally, integrate the gradient field $\nabla\Phi^*$, by computing its divergence and providing it as input to the Poisson solver you implemented earlier. The result will be a final fused image Φ^* . We recommend using Dirichlet boundary conditions: Set the outermost pixels of Φ^* , at all four image edges, to equal the corresponding pixels of either the ambient image, or the flash image, or their average. Likewise, you can use different initializations, corresponding to either the ambient image, or the flash image, or the average of the two, or even the all-zero image. **Experiment with different boundary conditions and initializations, and show the final fused image you obtain.**

3. Capture your own flash/no-flash pairs (100 points)

Now it is time to apply what you implemented above to your own pictures. You should capture two flash/no-flash pairs. The first pair should be suitable for applying the denoising techniques based on bilateral filtering. Good examples include dimly lit environments, e.g., a room illuminated by just a desk lamp, or an outdoors scene at night illuminated by only a street lamp. You can look at Petschnigg et al. [6] for inspiration.

The second pair should be suitable for applying the fusion algorithm based on gradient-domain processing. To create results which are clearly better than any single exposure, you should take pictures of a scene where flash affects some parts of the image, while leaving others relatively unaffected. Good examples include dark scenes that have both matte and specular objects. You can look at Agrawal et al. [1] for inspiration.

The total number of points you will get for this part will depend on how visually compelling the final fused images you create are.

4. Bonus: Reflection removal in flash/no-flash photography (50 points)

A different potential problem when using flash is that it can create reflections off of glass surfaces. See Figure 6 in Agrawal et al. [1]. Implement the fusion algorithm described in Section 4.1 (“Observation Deck”) of the paper, and then use it on a flash/no-flash pair you capture with your own camera. We suggest using a scene with a window for demonstrating this part, as in the paper.

5. Bonus: Other types of gradient-domain processing (up to 50 points)

Gradient-domain processing is a general image processing paradigm that can be used for a broad set of applications, including blending, tone-mapping, colorization, converting to grayscale, edge enhancement, image abstraction and non-photorealistic rendering, and so on.

[GradientShop](#) [2] is a paper and associated app that provide a unifying framework for expressing all these tasks as gradient-domain processing algorithms. Read this paper, and implement whichever of the applications described within appeal to you the most (the instructor’s personal favorite is non-photorealistic rendering, but it won’t reflect badly on your grade if you go for something else). For additional inspiration, you can take a look at the results in the Poisson image editing paper [5].

Your implementation does not need to follow GradientShop exactly. For instance, instead of implementing the “long-edge detector” described in the paper, you can come up with your own detector of continuous edges. We also recommend *not* trying the video-based applications.

In your submitted report, you should make sure to describe in detail exactly what you did. This includes what task you chose to implement, where you followed and where you deviated from GradientShop, and in the latter case what you did differently. You should also show at least three examples showcasing the effects of your gradient-domain processing pipeline.

Deliverables

When submitting your solution, make sure to follow the [homework submission guidelines](#) available on the course website. Your submitted solution should include the following:

- A PDF report explaining what you did for each problem, including answers to all questions asked throughout Problems 1, 2, and 3, as well as any of the bonus problems you choose to do. The report

should include any figures and intermediate results that you think may help. Make sure to include explanations of any issues that you may have run into that prevented you from fully solving the assignment, as this will help us determine partial credit. The report should also explain any additional image files you include in your solution (see below).

- All of your Python code, including code for the bonus problems, as well as a `README` file explaining how to use the code.
- Image files for Problems 1 and 2, showing the various fusion results you are asked to generate.
- Image files that you create in Problem 3, showing the original flash and no-flash pairs you capture with your camera, and the fused images you generate from them. You can also include additional image files for various experiments (e.g., fusion with different values) other than your final ones, if you think they show something important.
- If you do Bonus Problem 4: The original flash/no-flash images you captured and the corresponding fused image files you produced.
- If you do Bonus Problem 5: The original and corresponding processed image files you used.
- For the photography competition: Submit one of the fused photographs you produced for Problem 3, named as `competition_entry.png`.

Hints and Information

Color. When dealing with color images, you should apply all the algorithms you implement to each color channel separately. (Though note the hint below about the mask M .)

Gaussian filtering. In your bilateral filtering implementation, you do not need to implement Gaussian filtering on your own. Instead, you can use OpenCV's `gaussianblur` function.

Interpolation weights. The “PiecewiseBilateral” algorithm of Durand and Dorsey [3] includes an InterpolationWeight step, which the paper implements using linear interpolation. We recommend implementing this as follows: First, accumulate the images J^j for all j values in a `NB_SEGMENTS+1`-channel image, where the j -th channel equals J^j . Then, use `scipy`'s `interpolate.interpn` to interpolate the channels and form a final single-channel image.

Additionally, when creating the segments, you should make sure that the first value i^j is smaller than the smallest value in your image (typically 0), and the last value i^j is larger than the largest value in your images (typically 1). You can do this by defining $\text{minI} \equiv \min(I) - \lambda$ and $\text{maxI} \equiv \max(I) + \lambda$, where λ is a small constant (e.g., $\lambda = 0.01$, should be smaller than σ_r). Then, use these adjusted values of minI and maxI to set $\text{NB_SEGMENTS} = \text{ceil}\left(\frac{\text{maxI} - \text{minI}}{\sigma_r}\right)$ and $i_j = \text{minI} + j \cdot \frac{\text{maxI} - \text{minI}}{\text{NB_SEGMENTS}}$.

Lastly, it is better to compute minI and maxI by computing a min and max (respectively) intensities across the three color channels, rather than per-channel. In joint bilateral filtering, make sure to compute these values using the flash (i.e., guide) image.

Debugging bilateral filter. When debugging your bilateral filter implementation of the bilateral filter, it can be helpful to compare your output with that of OpenCV's `bilateralFilter`. However, you should *not* use the OpenCV implementation to produce any results.

Bilateral filtering function design. To make code management easier, we recommend having a single function for bilateral filtering, that takes as input two images. When doing regular bilateral filtering, both of these input images should be the ambient image. When doing joint bilateral filtering, one input image should be the ambient image, and the other should be the flash image. Organizing your code this way helps avoid code duplication issues (e.g., you need to fix any bugs you discover in only one place instead of two).

Shadow and specular mask. When computing the mask $\mathbf{M}$ for shadow and specular masking in Equation (2), you will get better results by using *linear* images. However, the images you are provided with are non-linear, and you do not have an exposure stack for radiometric calibration. As a default, you can apply the gamma correction operator of the *sRGB standard* [4]. This is the inverse of the gamma encoding operator you used in homework assignments 1 and 2, and corresponds to the following transformation:

$$C_{\text{linear}} = \begin{cases} \frac{C_{\text{non-linear}}}{12.92}, & C_{\text{non-linear}} \leq 0.0404482, \\ \left(\frac{C_{\text{non-linear}} + 0.055}{1.055} \right)^{2.4}, & C_{\text{non-linear}} > 0.0404482, \end{cases} \quad (8)$$

Additionally, when comparing the linearized flash and ambient images to compute $\mathbf{M}$, you will need to make sure that their intensities correspond to the same exposure settings. For this, you will need to normalize the linearized ambient image so that its ISO is the same as that of the linearized flash image. See Equation (1) of Petschnigg et al. [6]. You can find the ISO values for the lamp image file pair we provide you with in the metadata of the image files. You can assume that the exposure time for both images is the same.

The mask $\mathbf{M}$ may turn out to have a lot of small disconnected areas, or large connected areas with small gaps. These can result in strong artifacts in your final image $\mathbf{A}^{\text{Final}}$. You can get better results by using *morphological filtering* operations (dilation, erosion, opening, and closing) to post-process the individual shadow and specular masks, before combining them to form the mask $\mathbf{M}$.

Finally, even though you should apply the various bilateral filtering operations separately on each channel, you will likely get better results by computing a mask $\mathbf{M}$ from the luminance of the ambient and flash images, and using this same mask for all color channels. You can experiment with all these options and see what gives you better results.

Gradient operators. To solve this assignment, you will need to implement the gradient, divergence, and Laplacian differential operators. We strongly recommend creating three subroutines, `gradient`, `divergence`, `laplacian`, that you reuse throughout the assignment whenever you need to compute these quantities. To compute the first-order derivatives in the gradient and divergence operators, we recommend using `numpy`'s `diff` function *with appropriate padding*. To compute the Laplacian, we recommend using `scipy`'s `signal.convolve2d` function to convolve with the Laplacian kernel,

$$\Delta = \begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix}, \quad (9)$$

with arguments `mode='same'`, `boundary='fill'`, `fillvalue=0`.

The gradient and divergence implementations require different padding; figuring out exactly how to pad can be tricky and will require some trial-and-error. To test your implementation, you should ensure that for any arbitrary image $\mathbf{I}$, the following holds

$$\text{divergence}(\text{gradient}(\mathbf{I})) = \text{Laplacian}(\mathbf{I}), \quad (10)$$

at all pixels *except* at the boundary (first and last row, first and last column). Note that the incorrect boundary values is the reason we use the mask $\mathbf{B}$ when computing the residual in Algorithm 1.

RAW versus rendered. The algorithms in this assignment can work with both RAW and rendered (e.g., PNG or JPEG) images. Typically, working with RAW will give you better looking results, but in that case you will also need to use `dcraw` and tonemapping. We leave it up to you to decide what images to use.

Clipping. If you write images using the `skimage` input-output functions, these functions automatically scale all intensities in the $[0, 1]$ range. As a result, if your filtered image happens to have a very bright pixel with intensity greater than 1, your output image file will look very dark, even if the filtered image looks correct. To avoid this issue, you can clip (not scale) all image intensities to the $[0, 1]$ range before writing to a file.

Credits

Some inspiration for this assignment and the write-up came from Oliver Cossairt’s computational photography course at Northwestern University. The lamp and museum flash/no-flash pairs came from Petschnigg et al. [6] and Agrawal et al. [1], respectively.

References

- [1] A. Agrawal, R. Raskar, S. K. Nayar, and Y. Li. Removing photography artifacts using gradient projection and flash-exposure sampling. *ACM Transactions on Graphics (TOG)*, 24(3):828–835, 2005.
- [2] P. Bhat, C. L. Zitnick, M. Cohen, and B. Curless. Gradientshop: A gradient-domain optimization framework for image and video filtering. *ACM Transactions on Graphics (TOG)*, 29(2):10, 2010.
- [3] F. Durand and J. Dorsey. Fast bilateral filtering for the display of high-dynamic-range images. In *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH ’02*, pages 257–266, New York, NY, USA, 2002. ACM.
- [4] International Electrotechnical Commission and others. IEC 61966-2-1:1999. *Multimedia systems and equipment—Colour measurements and management—Part 2-1: Colour management—Default RGB colour space—sRGB*, 1999.
- [5] P. Pérez, M. Gangnet, and A. Blake. Poisson image editing. In *ACM Transactions on graphics (TOG)*, volume 22, pages 313–318. ACM, 2003.
- [6] G. Petschnigg, R. Szeliski, M. Agrawala, M. Cohen, H. Hoppe, and K. Toyama. Digital photography with flash and no-flash image pairs. *ACM transactions on graphics (TOG)*, 23(3):664–672, 2004.
- [7] J. R. Shewchuk. An introduction to the conjugate gradient method without the agonizing pain. school of computer science. *Carnegie Mellon University, Pittsburgh, PA*, 15213:10, 1994.