DOI: 10.58895/ksp/1000150865 -9 erschienen in:
Forum Bildverarbeitung 2022 | Image Pocessing Forum 2022
DOI: 10.5445/KSP/1000150865 | https://www.ksp.kit.edu/site/books/m/10.58895/ksp/1000150865/

Real-time multi-image vignetting and exposure correction for image stitching

Christian Kinzig[1], Guanzhi Feng[1], Miguel Granero[2], and Christoph Stiller[1]

[1] Karlsruhe Institute of Technology (KIT),
Institute of Measurement and Control Systems,
Engler-Bunte-Ring 21, 76131 Karlsruhe, Germany
[2] University of Seville (US),
C. San Fernando 4, 41004 Seville, Spain

Abstract Seamless image stitching depends not only on the accurate alignments of camera images, but also on the compensation of illumination inconsistencies. Even if two images are aligned perfectly, the seam is still visible if the images have a distinct vignetting or different exposure. Image stitching is used to expand the field of view, but a visible seam can lead to significant errors in subsequent visual perception tasks. As a result, we present a straightforward and accurate method for vignetting and exposure correction for stitched images. Firstly, we estimate the camera response function that maps irradiance to intensity. Then, the vignetting model is determined, which is applied to the irradiance images. After that, the exposure of the stitched images is corrected with the irradiance values at the seam. Finally, the irradiance is converted back into intensity using the camera response function. Our approach is evaluated using data recorded by our experimental vehicle and the public nuScenes dataset. Thereby, we test the performance of our method using the IoU of the histograms as well as the mean absolute error of the intensity values in the overlapping image regions. Furthermore, we demonstrate the real-time capability of our approach.

Keywords Autonomous driving, panorama, image stitching, vignetting, exposure, illumination

(a) Sensor setup prototype for UNICAR*agil* project.

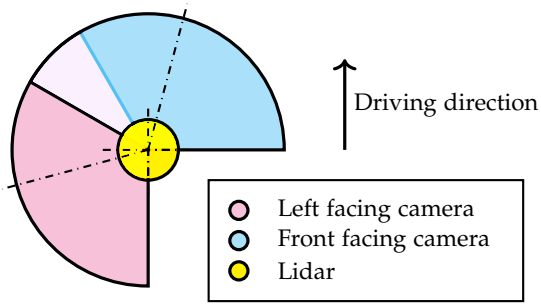

(b) Schematic top view of the UNICAR*agil* sensor setup to visualize sensor coverage of color cameras and lidar.

Figure 1: The images from the two lower color cameras of the UNICAR*agil* sensor modul are stitched to a 270° horizontal panoramic image to improve object detection and other perception tasks.

1 Introduction

Autonomous vehicles heavily depend on camera sensors to perceive their surroundings. Object detection, visual localization and mapping are fundamental challenges in automated driving based on camera images. Instead of performing perception tasks for each individual image, the images can be fused to a panorama beforehand [1]. Thus, the horizontal viewing angle can be significantly expanded using image stitching. This facilitates object detection, especially when an object is cut off at the image boundaries by a limited field of view. Image stitching precisely aligns individual images based on image features or lidar measurements. However, the seam is still visible due to vignetting and different exposure times. On the one hand, as shown in [2], vignetting is caused by a radial falloff in irradiance at the image boundaries, while on the other hand, the cameras adjust the exposure time to the current lighting conditions. As a result, the seam between stitched images can lead to false features in object detection and other processing tasks.

In this paper, we propose a straightforward method for compensating vignetting and correcting exposure for multiple stitched images in a time-critical environment. This distinguishes our method from many approaches that aim to compensate for vignetting in individual images using more complex models [3–5]. Hence, we estimate the camera response function (CRF) and the vignetting model before runtime. After the images are stitched, the vignetting model is applied and the exposure is corrected. Our approach is tested on the sensor setup prototype built as part of the publicly funded project UNICAR*agil* [6,7], as well as on the public nuScenes dataset [8]. The prototype of the sensor setup mounted on a vehicle of the Karlsruhe Institute of Technology is shown in Fig. 1(a). In this setup, the camera images from the front-facing camera and the left-facing camera are stitched together to create a panorama. The sensor coverage of the two cameras and the lidar, which allows better alignment of the images, is shown in Fig. 1(b).

2 Related Work

Since cameras are widely used, inexpensive sensors, many articles have already addressed vignetting and exposure correction. Goldman and

Chen provide a good overview of the causes of vignetting and suggest a general approach by modeling the vignetting model as 6th order even polynomial in [2]. In addition, the approaches of Zheng et al. in [3,4] and Cho, Lee, and Lee in [5] focus on vignetting correction for single images. Furthermore, the approach of Kordecki, Palus, and Bal propose the use of a non-radial vignetting model in [9].

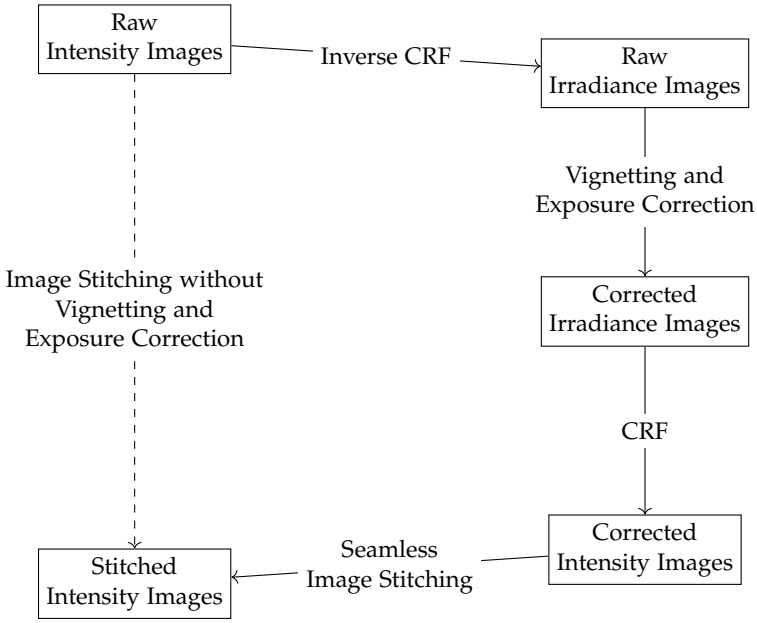

Figure 2: Workflow for vignetting compensation and exposure correction in image stitching.

3 Vignetting and Exposure Correction

Our approach consists of four individual steps. The workflow of our approach is depicted in Fig. 2. First, the camera response functions of all cameras are estimated. Then, the vignetting model is generated and applied. Afterwards, the correction of the exposure between the stitched camera images is performed. In the last step, the corrected

(a) Image stitching without vignetting and exposure correction.

(b) Image stitching with vignetting compensation but without exposure correction.

(c) Image stitching with vignetting and exposure correction.

Figure 3: Comparison of image stitching with and without vignetting compensation and exposure correction. The images are recorded with the UNICAR*agil* sensor setup prototype in Fig. 1.

irradiance values are converted back into intensities. The quantitative effect of our approach is shown in Fig. 3 for an exemplary pair of stitched images.

3.1 Calibration of the Camera Response Function

Both the vignetting compensation and the exposure correction are performed based on the irradiance, which is calculated from the intensities using the non-linear camera response function. Therefore, we determine the camera response function of our sensor setup before the actual runtime. The camera response functions are estimated by exposure series in a static scene with known exposure times as in [10]. For each of the color cameras in Fig. 1, we obtain three response functions for the three color channels. However, we found that the camera response functions are approximately identical for the cameras and all color channels. The qualitative evaluation of UNICAR*agil* sensor data shows decent results for vignetting and exposure correction using the approximated camera response function. For this reason, we store only the approximated camera response function for the entire panorama, which is shown in normalized form in Fig. 4(a). After vignetting and exposure correction in 3.2 and 3.3 the camera response function is used to convert the irradiance values back to intensity values.

3.2 Estimation of the Vignetting Model

To compensate for vignetting, we found that in our case a model can be sufficiently created by approximating the vignetting by the cosine-fourth-power law. This estimates the radial irradiance falloff at the boundaries of the camera images. To get a better result for the panoramic image we use a spherical camera model in our approach, that is described in more detail in [1]. As with the pinhole camera model, the intrinsic parameters can be specified in the matrix $\mathbf{A}$ as in Eq. 1, where f denotes the focal length and (u_0, v_0) describes the principal point. Another advantage of the spherical camera model is that the pixel coordinate is proportional to the angle of incidence. This results in Eq. 2, which models the vignette as a function of the distance to the principal point r and the focal length f. In addition, the variables a and b are used to fit the vignetting model. The values are determined

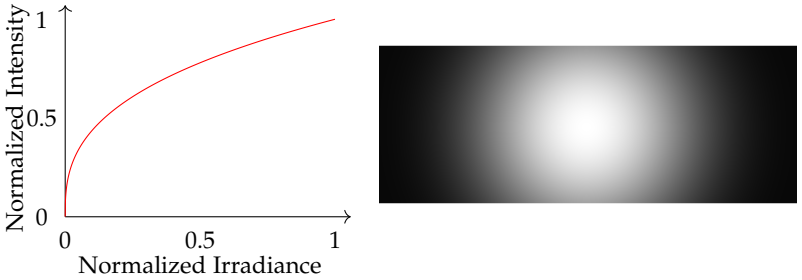

(a) Approximated and normalized camera response function of the UNICAR*agil* sensor setup.

(b) Normalized vignetting model in the spherical camera frame.

Figure 4: Both the camera response function (a) and the vignetting model (b) are computed before the actual runtime, to be applied to the camera images afterwards.

empirically based on a sequence recorded with the UNICAR*agil* sensor module and result in our case in $a = 3.4$ and $b = 0.1$. In Fig. 4(b) the normalized vignetting model in the spherical camera frame is shown. Just like the camera response function, the vignetting model is also created before runtime to achieve real-time capability.

$$\mathbf{A} = \begin{pmatrix} f & 0 & u_0 \\ 0 & f & v_0 \\ 0 & 0 & 1 \end{pmatrix} \tag{1}$$

$$\begin{aligned} g(u,v) &= a\,cos^4(r/f) + b, \\ r &= \sqrt{(u-u_0)^2 + (v-v_0)^2},\ a,\ b \in \mathbb{R}_0^+ \end{aligned} \tag{2}$$

3.3 Exposure Correction

After the vignetting model is applied, the brightness between the stitched images is not fully adjusted due to a difference in exposure, as can be seen in Fig. 3(b). To perform exposure correction, an individual exposure compensation factor c is determined for each image. For

this purpose, we calculate the quotient c_{rel} from the cumulative irradiances E along the vertical seam in the overlapping region for a pair of stitched images in Eq. 3. Based on the quotient c_{rel}, we derive in the following step the individual exposure correction factors for the single images in Eq. 4 by an additional constraint to ensures that the average of the factors is equal to one. If a panorama consists of more than one seam n, several quotients $c_{rel,n}$ are obtained, from which the individual exposure factors can be calculated. For multiple stitched images with vertical seams n, we use the transitional condition $c_{right,n} = c_{left,n+1}$. In Fig. 3(c) the stitching result after exposure correction is shown.

$$c_{rel} = \sum_{v=0}^{\text{height}} \frac{E_{left}(u(v), v)}{E_{right}(u(v), v)} \tag{3}$$

$$c_{left} = \frac{2}{c_{rel} + 1}, \; c_{right} = \frac{2c_{rel}}{c_{rel} + 1} \tag{4}$$

4 Experimental Results and Evaluation

We evaluate the presented approach on the UNICAR*agil* sensor setup shown in Fig. 1 as well as on the sequences 1 to 10 of the nuScenes dataset [8]. In the latter case, we use the images from the front-facing camera and the front-left-facing camera to create a panorama. Since we do not know what kind of cameras are used in the nuScenes setup and we cannot reconstruct the camera response function from the available data, we assume a linear response function as an approximation. Besides qualitative results in Fig. 3, we show the performance of our method using two different metrics in 4.1 and 4.2. First, the intersection over union of the histograms in the overlapping region of the stitched images is used. The second metric used is the mean absolute error of intensity differences in the overlapping region. Furthermore, we analyze in 4.3 the runtime of the vignetting compensation and exposure correction for stitched images and show its real-time capability.

4.1 Intersection over Union of Histograms

To measure the accuracy of image stitching, we compare the histograms of the two overlapping regions of the single images. This is done before and after vignetting and exposure correction to show the improvement due to our approach. For a better comparison, the images are converted to 8-bit grayscale so that the histogram values are between 0 and 255 with a bin size of 1. The similarity of two histograms H_i can be measured by calculating the intersection over union (IoU). The IoU between two histograms is calculated according to Equation 5. To prevent the resulting panorama from being extremely over- or underexposed, only pixels intensities with values unequal 0 and 255 are considered for evaluation. Table 1 shows the average IoU values of the histograms before and after vignetting and exposure correction on the recording with our UNICAR*agil* sensor setup and on the nuScenes dataset. The increasing IoU using our approach shows that the histograms of the two overlapping regions are better aligned than without our approach.

$$IoU = \frac{\sum_{n=0}^{255} min(H_{left}(n), H_{right}(n))}{\sum_{m=0}^{255} max(H_{left}(m), H_{right}(m))} \tag{5}$$

Table 1: Comparison of the average IoU of the histograms from the overlapping areas of the stitched images. The stitching quality is compared between using only raw images to processed images using our approach for vignetting and exposure correction for the two different image sequences.

	Raw images	Processed images
UNICAR*agil*	46.29 %	55.94 %
nuScenes	38.43 %	46.45 %

4.2 Mean Absolute Error

Since identical histograms can be derived from different images, we additionally evaluate the local similarity between pixel intensities with the mean absolute error. Compared to Zheng et al. we do not measure the difference to a ground truth vignetting function [3]. Instead, we also use the overlapping regions of the single images and calculate

the mean absolute error as another similarity measure to evaluate the image stitching performance. Thereby, we convert the images to 8-bit grayscale and calculate the mean of the absolute differences pixelwise, as shown in Equation 6. Similar to 4.1, we use only pixel pairs with values unequal 0 and 255 for evaluation. The mean absolute error is calculated by dividing by the number of pixel pairs and averaging it by the number of samples in the sequences. The results before and after vignetting and exposure correction are depicted in Table 2 for the sequence recorded with the UNICAR*agil* sensor setup as well as for the nuScenes dataset. The evaluation clearly shows that the mean absolute error decreases if our approach is applied to the images.

$$MAE = \frac{\sum_{u=0}^{width} \sum_{v=0}^{height} |g_{left}(u,v) - g_{right}(u,v)|}{u \cdot v} \tag{6}$$

Table 2: Comparison of the average mean absolute error of the pixel intensities from the overlapping areas of the stitched images. This allows the comparison between using only raw images to processed images using our approach as in 1.

	Raw images	Processed images
UNICAR*agil*	21.58	7.71
nuScenes	37.56	10.99

4.3 Runtime Analysis

In addition to the metrics, which show an improvement in accuracy, we evaluated the real-time capability of our approach. To optimize our approach in terms of its execution time, we run the processing operations directly on the graphics card. This is an option as soon as the entire image processing in an autonomous vehicle is performed on the graphics card since copying data to and from the graphics card takes a lot of time. This offers further advantages, for example, for object detection with machine learning. The improved runtime is an exceptional feature of our simplified approach to vignetting and exposure correction compared to the approaches in [3,5]. In Table 3, we compare the average runtimes of vignetting compensation and exposure correction on

CPU and GPU. On the computer for evaluation, we use *Ubuntu 18.04.6 LTS* as operating system. As CPU an *Intel® Xeon® Prozessor E5-2640 v3* running at 2.6 GHz with 64 GB of RAM and as GPU a *NVIDIA GeForce RTX 2080 Ti* are installed. The table clearly shows that we achieve real-time capability at a frame rate of 10 Hz with an average processing time of 31.36 ms by using the GPU. Further improvements can be expected on the latest generation of *NVIDIA* graphics cards.

Table 3: Comparison of the average runtimes of our approach on vignetting and exposure correction on CPU and GPU.

	CPU	GPU
Runtime in ms	155.35	31.36

5 Conclusion

In this paper we presented a straightforward and effective method on vignetting and exposure correction for multiple camera images and image stitching. Our approach relies on a known camera response function and a previously estimated vignetting model that are applied on the images to be stitched. First, the irradiance is calculated from intensity using the inverse camera response function and our vignetting model is applied. Then, the optimal exposure correction factors for the single images are estimated from the pixels at the seam to improve the quality of the panorama. After vignetting and exposure correction, the intensities are obtained from the modified irradiance values. In summary, the vignetting of the single images is compensated and the transition at the seam of the panorama due to different exposure is corrected. We evaluated our approach by calculating the IoU between the histograms of the overlapping regions of the stitched images before and after vignetting and exposure correction and have clearly demonstrated that the IoU increases significantly after applying our approach. In addition, we have shown that the mean absolute error of the overlapping regions after vignetting and exposure correction also decreases strongly. Both quantitative results confirm the significant improvement in image stitching quality after using our approach. This can lead to higher precision in object detection and other perception tasks. Finally,

we have also shown that our approach can be executed on a graphics card in real-time. To further extend our approach, we plan to integrate joint optimization of exposure correction factors for multiple seams of a full 360° horizontal panoramic image in the future.

Acknowledgment

This research is accomplished within the project UNICAR*agil* (FKZ 16EMO0287). We acknowledge the financial support for the project by the Federal Ministry of Education and Research of Germany (BMBF).

References

1. C. Kinzig *et al.*, "Real-time seamless image stitching in autonomous driving," in *25th International Conference on Information Fusion (FUSION)*, 2022, pp. 1–8.
2. D. B. Goldman and J.-H. Chen, "Vignette and exposure calibration and compensation," in *IEEE International Conference on Computer Vision*, 2005, pp. 899–906.
3. Y. Zheng *et al.*, "Single-image vignetting correction using radial gradient symmetry," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2008, pp. 1–8.
4. ——, "Single-image vignetting correction," *IEEE transactions on pattern analysis and machine intelligence*, vol. 31, pp. 2243–56, 2009.
5. H. Cho, H. Lee, and S. Lee, "Radial bright channel prior for single image vignetting correction," in *ECCV*, 2014.
6. T. Woopen *et al.*, "UNICARagil - Disruptive Modular Architectures for Agile, Automated Vehicle Concepts," *27th Aachen Colloquium Automobile and Engine Technology*, pp. 663–694, 2018.
7. M. Buchholz *et al.*, "Automation of the UNICARagil vehicles," in *29th Aachen Colloquium Sustainable Mobility*, 2020, pp. 1531–1560.
8. H. Caesar *et al.*, "nuscenes: A multimodal dataset for autonomous driving," *arXiv preprint arXiv:1903.11027*, 2019.
9. A. Kordecki, H. Palus, and A. Bal, "Practical vignetting correction method for digital camera with measurement of surface luminance distribution," *Signal, Image and Video Processing*, vol. 10, 11 2016.

10. P. E. Debevec and J. Malik, "Recovering high dynamic range radiance maps from photographs," in *24th Annual Conference on Computer Graphics and Interactive Techniques*, 1997, p. 369–378.