Color Homogenization in 3D Reconstruction of Castles from Heterogeneous Imaging Sources

Angeliki Oraiozili Gyftodimou[1], Pierre Grussenmeyer[2], Maria Pateraki [1], Etienne Sommer[2], Mathieu Koehl[2]

[1] Laboratory of Photogrammetry, School of Rural, Surveying and Geoinformatics Engineering, National Technical University of Athens, Greece - angiegyftodimou@gmail.com - mpateraki@mail.ntua.gr
[2] Université de Strasbourg, INSA Strasbourg, CNRS, ICube Laboratory UMR 7357, 67084 Strasbourg, France
pierre.grussenmeyer@insa-strasbourg.fr, etienne.sommer@insa-strasbourg.fr, mathieu.koehl@insa-strasbourg.fr

Keywords: Colour Transformation, Image Equalization, 3d Reconstruction.

Abstract

This study, conducted as part of the Rhineland Castles 3D Modeling Project, addresses the problem of color inconsistencies in 3D reconstructions generated from diverse imagery sources, including drone footage and terrestrial photography. Variations in lighting conditions, camera sensors and acquisition settings often result in photometric discrepancies that degrade the quality of the textured 3D model.
Two datasets from Ramstein and Oedenbourg Castles in France were analysed, each presenting characterised by distinct photometric conditions. The Ramstein dataset exhibited relatively uniform lighting, allowing the application of statistical color transfer methods. Among the methods evaluated, Mean-Lab Transfer produced the most photometrically consistent and visually faithful results. In contrast, the Oedenbourg dataset presented significant photometric variability, characterised by extreme variations in saturation, luminosity, and contrast. To effectively manage these challenging conditions Here, a two-step enhancement approach was adopted, combining gamma correction (in HSV color space) and with CLAHE to balance brightness and preserve color identity. This paper details the applied methodologies, evaluates their effectiveness in achieving color consistency, and highlights the importance of adopting emphasizing the need for dataset-specific processing pipelines.

1. Introduction

In the context of cultural heritage documentation, 3D reconstruction plays a central role enabling detailed analysis, virtual access, and long-term preservation of cultural heritage sites. However, the integration of images acquired from different platforms, such as UAVs and DSLR cameras, poses significant challenges in terms of photometric consistency. Variations in lighting conditions, camera characteristics, and acquisition parameters often lead to inconsistent color representation, compromising both the visual realism and interpretability of 3D models. The monuments in question can be seen in Figure 1.

Figure 1. Ramstein castle (left) / Oedenbourg castle (right)

This study focuses on the problem of color homogenization in textured 3D reconstructions generated from heterogeneous image datasets. Specifically, it investigates the case of two historical castles, Ramstein and Oedenbourg, using terrestrial and aerial imagery.
The Ramstein Castle, was built in 1293 near Scherwiller in Alsace (in the Bas-Rhin region in France). It is a medieval fortress strategically placed opposite Ortenbourg Castle. Though largely destroyed in the 15th century, remains of its walls and a corridor likely used for catapults still stand. The Oedenbourg Castle, also called Petit-Koenigsbourg, is a 13th-century fortress near Orschwiller, close to Haut-Koenigsbourg Castle. Its name, meaning “abandoned castle,” appears in records from 1417. Built at 710 meters altitude, it served as a lookout and artillery site. Excavations revealed it was inhabit-ed in the early 15th century but abandoned after 1420. Its ruins reflect the region's medieval military architecture.
This work aims to assess and implement color correction techniques that restore consistent color representation across datasets, thereby enhancing the photorealism of reconstructed models. It explores statistical color transfer methods, gamma correction, and histogram equalization strategies tailored to the specific characteristics of each dataset.

2. Related work

Color transfer is a fundamental operation in digital image processing, aimed at harmonizing the chromatic appearance of images by modifying the color distribution of a target image to match that of a reference image. This is particularly critical in photogrammetric applications and 3D reconstruction workflows, where visual coherence across images captured under varying illumination, exposure settings, or sensor types is essential. According to Chenlei et al. (2023), color transfer methods are broadly classified into Statistical Color Transfer, Learning-Based Color Transfer, and Semantic-Based Color Transfer approaches. The present work focuses on the first category, which represents classic, well-established techniques relying on global or local image statistics.
Among the statistical methods applied are: (i) the Mean and Standard Deviation Transfer, which performs histogram matching in the RGB color space by aligning the first two statistical moments; (ii) Mean-Lab Transfer, which utilizes the CIELAB color space for better control over luminance and chromatic adaptation; and (iii) Probability Density Function

(PDF) Transfer, which attempts to align the full histograms of the source and target images, typically using cumulative distribution functions. These methods are widely used due to their computational simplicity, scalability, and independence from semantic understanding of the scene. However, they may introduce color distortions or fail under scenes with complex lighting variations, shadows, or non-linear camera responses.
As observed in this work, the effectiveness of statistical transfer methods depends heavily on the selection of a suitable reference image, as well as the photometric variability of the dataset. In relatively uniform datasets such as Ramstein, Mean-Lab transfer achieved satisfactory visual homogeneity. In more photometrically challenging datasets like Oedenbourg, however, additional contrast enhancement steps (e.g., gamma correction or CLAHE) were required to complement statistical color transfer.
Another inspiration was Welch et al. (2005), a study that addresses the challenge of achieving color consistency in multi-camera systems, where variations in camera sensor responses lead to significant color discrepancies, even under uniform illumination. The proposed framework introduces a calibration method based on a reference image approach, wherein one image is treated as the standard, and all others are transformed to match its color profile.

3. Ramstein Castle Dataset: Image Analysis & Color Correction

The Ramstein Castle dataset consists of 300 images captured using a DJI Phantom 4 drone as well as nearly 4000 terrestrial photos taken with a Canon Mirrorless Camera (model: EOS R5). Relatively uniform lighting conditions have been observed, alongside with some variations in brightness and color fidelity. The drone images were taken on May 4th, 2020, in the morning, whereas the terrestrial images were taken on November 23rd 2023, similarly in the morning.

3.1 Image quality issues

Unique challenges in terms of image quality have been presented in the dataset in question, that is comprised of both aerial and terrestrial imagery.
Aerial Imagery (Drone): The aerial imagery exhibits significant luminosity and color inconsistencies. A critical issue that most of the aerial imagery faces are the problematic exposure levels, as these images are underexposed. As a result, the initial visual outputs are darker than desired. In contrast, a subset of images suffers from overexposure, particularly scenes with extreme brightness and contrast. Another issue also is uneven image compositions, often with a disproportionate presence of soil dominating the frame and overpowering the visual focus on the castle itself.
In drone imagery especially, an unnatural yellow tone is prevalent in many cases, and more notably, the distinctive red hue of Ramstein Castle is almost entirely absent. This degradation in chromatic fidelity necessitates a more targeted approach to correction. Given the nature of these issues, a refined, nuanced approach is required, one that addresses both global and localized inconsistencies without introducing new artifacts. Figure 2 shows the problematic images captured from drone (aerial imagery).

Figure 2. Sample images captured from drone (aerial imagery)

Terrestrial Imagery (DSLR): Conversely, the colors of the terrestrial DSLR images are generally closer to real-life but minor enhancements are deemed necessary, nevertheless. The terrestrial images, while generally more faithful to the castle’s real-world appearance, have proper flaws of their own. The primary issue observed is underexposure, as the majority are very dark. Nevertheless, the colour balance and composition are more consistent and realistic compared to the aerial imagery. As a result, only minimal adjustments are necessary to bring these images to an analysis-ready state. The problematic images, captured by DSLR are seen on Figure 3.

Figure 3. Sample images captured from DSLR (terrestrial imagery)

3.2 Implemented Methods - Statistical Color Transfer

Modifications were applied to a specific wall of the Ramstein Castle, as this area exhibited the highest variance in luminosity and saturation across the dataset. The subsequent results can be found on Table 1. The unmodified wall of the castle of Ramstein can be seen in Figure 4.

Figure 4. Ramstein castle wall, without modifications (screenshot from Agisoft Metashape)

This subset of aerial images deviated most significantly from real-world appearance, thus presenting an ideal case study for testing color correction methods. The selected modifications focused solely on the aerial imagery, using a reference image sourced from the terrestrial dataset. This reference image was chosen for its accurate representation of the monument’s true colors, particularly the preservation of the characteristic red-pink tint of the castle’s walls.

To accommodate varying lighting conditions observed across the monument’s structure, the wall was subdivided into three horizontal zones: the top, middle, and bottom. Each of these segments presented distinct illumination challenges, which allowed for a granular evaluation of the effectiveness of different color transfer methods under diverse lighting conditions. To analyze lighting inconsistencies, the monument was divided into three sections: top, middle, and bottom.

The following Color Transfer methods were implemented, similarly using python programming language in Jupyter environment, and dedicated algorithms. The python libraries used were: PIL, skimage, glob, PIL.Image.

1. Mean-Std Transfer (RGB): This method aligns the mean and standard deviation of the source image with that of the reference image in the RGB color space. While computationally efficient, this method yielded suboptimal results. Notably, the processed output retained the original drone image’s undesired orange hue, and the monument’s red hue remained largely invisible.

2. *Probability Density Function (PDF) Transfer:* The second method employed, PDF-based color transfer, essentially aligns the distribution of pixel values between the source and reference images. Compared to the Mean-Std method, this approach resulted in sharpened contrast in the output images. It was also deemed unsuccessful in restoring the castle’s desired red hue and the tones across different sections of the monument remained inconsistent.

3. *Mean-Lab Transfer:* The final method converted images into the L*a*b* color space, applied the Mean-Std transfer, and then reconverted the output back to RGB. This approach successfully recovered the red-pink tint and homogenized the input images’ appearance. While it might be worth noting that the output appeared slightly darker than that produced by the RGB-based Mean-Std method, the overall color fidelity and contrast balance were superior.

Overall, Mean-Lab Transfer delivered the best results visually, maintaining color accuracy while minimizing contrast inconsistencies.

4. Oedenbourg Reconstruction Dataset: Image Enhancement

Figure 5. Oedenbourg castle wall, after the modifications (screenshot from Agisoft Metashape)

In this section, the image enhancement techniques applied to the Oedenbourg Castle dataset will be presented and the subsequent results are on Table 2. The modified wall of the castle of Oedenbourg can be seen in Figure 5. The dataset is, as well, comprised of both terrestrial and aerial imagery (specifically: 112 terrestrial images and 766 drone images). The drone images were taken on May 23th 2023, at noon, whereas the terrestrial images were taken on April 3rd, 2023, in the morning. The primary focus is placed on enhancing the existing images, which represent the most extreme cases in terms of contrast, saturation, and illumination inconsistency.

The aerial imagery (drone acquired images) in this dataset is characterized by overexposed or underexposed regions and severe color inaccuracies. These issues pose significant challenges for photogrammetric processing and 3D reconstruction. A key parameter that affects both contrast and brightness simultaneously, as identified in the literature (Singnoo et al. (2010)), is the gamma value (γ).

The **gamma (γ) parameter** plays a critical role in image processing, particularly in adjusting the overall luminance and contrast of an image. Unlike linear transformations, gamma correction applies a non-linear mapping to the pixel intensity values, allowing for finer control over the visual appearance of both dark and bright regions. Gamma correction does not rely on global brightness levels alone, as it affects the distribution of tones by compressing or expanding certain ranges within the intensity histogram. This is especially useful in datasets with challenging lighting conditions, such as those captured by UAVs, where shadows and overexposed areas often coexist. A gamma value greater than 1 compresses the darker tones and lightens the image, resulting in reduced contrast, while a value less than 1 expands the darker tones, increases contrast, and generally darkens the image. As such, gamma correction is frequently used as a pre-processing step in such workflows to restore visual balance and preserve spatial information in high-dynamic-range imagery. The mid-gray (“mid”) value is often set to 0.5 as a neutral reference point in calculating the transformation. Equation 1 is the formula of the gamma parameter. The mean value in Equation 1 refers to the average intensity of the image.

$$\gamma = \frac{\log(mid * 255)}{\log(mean\ of\ image)}$$

Equation 1. Calculation of gamma (γ)

To address these problems, a two-step enhancement workflow is proposed: gamma correction followed by Contrast Limited Adaptive Histogram Equalization (CLAHE). All processing steps were executed in Jupyter Notebook and python programming language, where white balance correction was implemented through a separate dedicated function. Additionally, a custom script was used to extract and transfer EXIF metadata from the original images to the processed outputs, ensuring data continuity.

Although basic color transfer was also tested, the presence of extreme shadowing and lighting conditions resulted in unsatisfactory output. The primary objective remains to enhance overall image uniformity and preserve spatial detail across shadowed and bright regions. Figure 6 shows the problematic images of the Oedenbourg castle dataset.

Figure 6. Oedenbourg castle dataset (sample images)

4.1 First step: Gamma Correction

Among the various parameters influencing image appearance, the gamma (γ) parameter stands out as a key factor affecting both brightness and contrast. Gamma correction is a non-linear transformation applied to the intensity values of individual pixels, aiming to adjust the luminance distribution of an image. This process is particularly effective in addressing issues arising from extreme lighting conditions, such as under- or overexposure, commonly encountered in aerial imagery.
The transformation operates according to the general formula, seen on Equation 2:

$$\gamma - correction = \ image^{\gamma}$$

Equation 2. gamma correction

Represent the normalized input and output pixel intensities, respectively.
The impact of gamma adjustment depends on the selected value:

- $\gamma > 1$ results in a lightening effect, compressing the dynamic range of dark areas and yielding an image with lower contrast and a narrower histogram.
- $\gamma < 1$ produces a darkening effect, enhancing contrast by stretching the dynamic range, particularly in brighter regions, and results in a broader histogram.

In this study, gamma correction was applied to grayscale (B&W) images to optimize spatial information recovery prior to histogram equalization. As suggested by Ergül & Eminoglu et al. (2020), gamma correction can be conceptualized as a logarithmic transformation that not only modifies brightness but also indirectly accounts for saturation and local pixel relationships. Figure 7 presents the transformation curve, highlighting the non-linear mapping behavior for different gamma values.

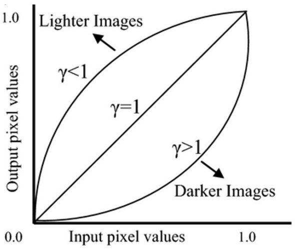

Figure 7. Ergül & Eminoglu (2020)

Gamma correction serves as a preparatory step in photometric preprocessing, effectively redistributing intensity values before more advanced enhancement techniques, such as CLAHE, are applied. Its integration ensures that saturated regions are minimized, and spatial information is preserved, particularly in UAV-acquired imagery with high dynamic range variations.

The first step of the proposed enhancement methodology is gamma correction. Gamma correction was tested RGB and HSV color spaces, to compare the output images. A representative image was selected as a test case due to the presence of highly overexposed regions, making it ideal for evaluating the effectiveness of each transformation.

4.1.1 Gamma Transformation in RGB Space

The application of gamma correction in the RGB color space resulted in more vivid colors and increased saturation. However, it also led to undesirable side effects. Specifically, due to overexposure, several pixels in the image were saturated, resulting in loss of spatial information. Histogram analysis revealed extreme spikes in the lower (0–20) and upper (>245) ranges of the R, G, and B channels, confirming the uneven tonal distribution and loss of detail.

4.1.2 Gamma Transformation in HSV Space

The HSV-based transformation produced a more balanced and perceptually accurate result. While the brightest areas remained luminous, the spatial information was better preserved, avoiding pixel burn-out. The overall color balance was closer to the true appearance of the monument. Nevertheless, even in the HSV-transformed output, histogram analysis continued to show peak concentrations at the extreme tonal values in the RGB channels, indicating persistent dynamic range issues. Figure 8 illustrates the results of the implementation of the gamma transformation in two different color spaces, RGB and HSV.

Figure 8. Gamma transformation results RGB (left) / HSV (right)

4.2 Second step: CLAHE (Contrast Limited Adaptive Histogram Equalization)

Following gamma correction, CLAHE was employed to address the pronounced tonal peaks observed in the image histograms. CLAHE was selected over other histogram equalization methods due to its suitability for such imagery, as highlighted by Manju et al. (2019).. Other methods such as BBHE (Brightness Preserving Bi-Histogram Equalization), first published in 1997 by Kim et al. (1997), that are effective for “normally” exposed images, are less appropriate for datasets such as this one, which contain severe exposure and saturation issues.
It is generally advised that CLAHE isn’t directly applied in the RGB color space, as RGB doesn’t take into consideration the luminocity factor. On the contrary, when is CLAHE applied on HSV, the results look much better. Thus, all input images were converted to either HSV or L*a*b* color spaces prior to processing. When CLAHE applied on HSV, the results look much better, and the contrast of the images also increased. The two key parameters, ClipLimit and TileGridSize, were systematically adjusted during experimentation. The ClipLimit determines the maximum contrast enhancement; higher values result in greater contrast but may also amplify noise and lead to histogram distortion. Conversely, excessively large values of

ClipLimit do not resolve saturated pixel issues and may introduce visual artifacts.
To conclude, although the combination of gamma correction and CLAHE improved contrast and preserved more detail in extreme lighting conditions, certain limitations remain. Specifically, despite enhancement, histogram analysis continued to reveal spikes in tonal extremes, and some residual noise persisted. Further research into adaptive parameter selection may yield improved results in future implementations.
CLAHE was chosen for its effectiveness in handling extreme histogram peaks. It was applied in HSV/lab color spaces as RGB was unsuitable. Despite the extreme histogram spikes that remained, the methodology produced satisfactory results. The castle’s red hue was preserved, as seen in Figure 9.

Figure 9. Effectiveness of CLAHE implementation, in terms of color accuracy (Screenshot from Agisoft Metashape)

4.3 Analytical Results – Oedenbourg Dataset
To evaluate the effectiveness of the implemented enhancement methods on the Oedenbourg dataset, a comparative analysis of the histograms and visual output was conducted for each transformation stage. Overall, this evaluation highlights how gamma correction and CLAHE affected pixel intensity distributions and overall image quality. An example of an image from the terrestrial dataset and its subsequent graphs will be utilized, since this image represents very well the dataset of the Oedenbourg castle.

Unedited Image

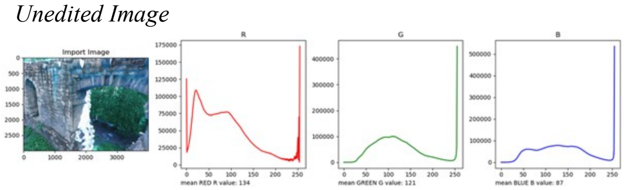

Figure 10 – R, G, B Histograms of unedited image (Unedited image with three histograms in red, green, and blue)

As shown in Figure 10, the unedited image displays strong peaks in the histograms of all three channels (R, G, B), with pixel values prominently clustering at the extremes. Specifically, on the “R” channel, there are values clustered on the (0–20). On all three channels (“R”, “G” and “B”) values are clustered in the (245–255). This is indicative of images with harsh lighting conditions and insufficient dynamic range. Notably, the mean values of the red and green channels (134 and 121, respectively) are significantly higher than that of the blue channel (87), indicating a general, slight chromatic imbalance.

Gamma Correction

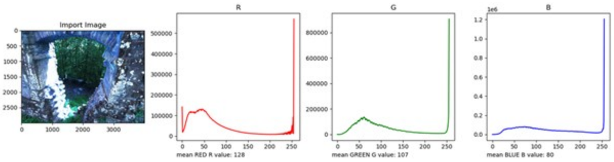

Figure 11 – R, G, B histograms before gamma correction (Original color image + 3 histograms)

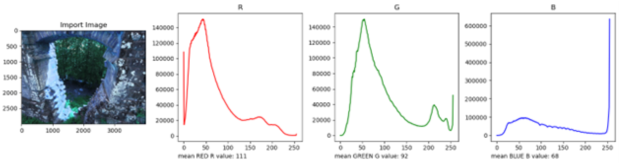

Figure 12 – R, G, B histograms after gamma correction (Gamma-corrected color image + 3 histograms)

Figures 11 and 12 demonstrate the image and histograms before and after gamma correction. The initial histograms (Figure 11) confirm the presence of numerous burned-out (overexposed) pixels, especially in the blue channel. After applying gamma correction (Figure 12), the histograms appear more spread out, indicating a broader dynamic range. The red channel, previously peaking near maximum intensity, now shifts leftwards, suggesting a better redistribution of values and a partial recovery of overexposed areas. The green channel similarly flattens, with more mid-range values emerging, while the blue histogram shows improved balance.
Overall, this transformation proves particularly beneficial for spatial information recovery in overexposed zones, as seen in the comparison of image brightness and contrast before and after the operation.

Gamma Transformation on B&W Image

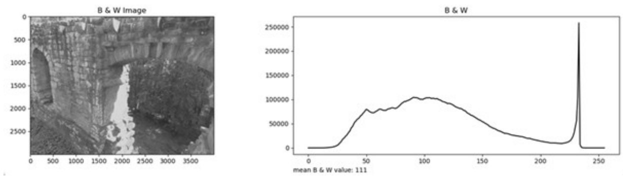

Figure 13 – Histogram of B&W image after gamma transformation but before CLAHE (Grayscale image + single grayscale histogram)

In the grayscale domain, the histogram of the gamma-corrected image, as seen in Figure 13, further illustrates this improvement. Compared to the flat, right-skewed histogram of the original image, gamma correction results in a more balanced distribution, with fewer clipped whites and a more even spread across mid-tone values. The mean grayscale value drops to 111, a sign of reduced overexposure and enhanced detail retention in previously saturated regions.

CLAHE Histogram Equalization

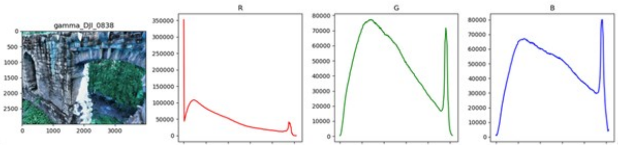

Figure 14 – R, G, B Histograms after CLAHE (CLAHE-enhanced color image + 3 histograms)

The final enhancement step involves applying CLAHE, and its impact will be analysed in this section. Specifically, in Figure 14, the histograms post-CLAHE show substantial contrast enhancement compared to earlier stages. The red, green, and blue channels now present smoother curves with visible peaks across a wider range of intensity values. This indicates a redistribution of pixel intensities that enhances local contrast without significantly amplifying noise.
The CLAHE-adjusted image reveals significantly improved visual balance, especially in shadowed and high-exposure areas. However, despite CLAHE’s benefits, strong histogram spikes remain in extreme brightness areas, suggesting the limits of histogram equalization in the presence of permanently burned-out pixels.

5. Conclusion

This work demonstrates the effectiveness of customized photometric correction strategies for heterogeneous image datasets in 3D reconstruction. While Mean-Lab Transfer yielded the best results for the relatively uniform Ramstein dataset, the Oedenbourg dataset required gamma correction and CLAHE due to its challenging photometric conditions. The study underscores the importance of dataset-specific approaches and highlights promising avenues for automation and refinement in future research.
Future work should delve into additional methods, including semantic-based color transfer techniques, to refine results. For Ramstein, optimizing reference image selection criteria is essential. For Oedenbourg, determining optimal ClipLimit and TileGridSize values will minimize noise and, subsequently, produce the most homogenous results. In both cases, color transfer methods preserve the castles' unique hues (e.g., Ramstein’s red tones), as certain colors appear differently across cameras. In all cases, colour modifications should be chosen carefully, paying attention to the monuments’ original colors. Park et al. (2016) provides inspiration for some of the following, as it presents an efficient and robust solution for achieving color consistency across large, unstructured image datasets, typically sourced from community photo collections. often suffer from inconsistent color appearances due to differences in lighting, camera sensors, and post-processing. The proposed method relies on establishing pairwise color relationships between images using automatically detected matching points, which are identified using methods such as SIFT feature matching.
Image processing plays a crucial role in computer vision and 3D reconstruction, with several promising areas for further study:

1. **Image Clustering:** Automating clustering based on exterior orientation parameters (yaw, pitch, roll) to group images by viewpoint.
2. **Reference Image Optimization:** Determining the most effective reference image or set of images for consistent color transfer.
3. **Semantic-Based Color Transfer:** Leveraging semantic segmentation to apply color correction based on object-level understanding.
4. **Successive Color Transfer:** Applying iterative color correction across image sequences to propagate photometric consistency.

References

Chenlei Lv, Dan Zhang, Shengling Geng, Zhongke Wu, and Hui Huang. 2023. Color Transfer for Images: A Survey. ACM Trans. Multimedia Comput. Commun. Appl. Just Accepted (November 2023). https://doi.org/10.1145/3635152

Ergul, Engin & Eminoglu, Ilyas. (2020). Power-law fitness scaling on multi-objective evolutionary algorithms: interpretations of experimental results. Soft Computing. 24. doi: 10.1007/s00500-019-04242-6.

Ilie, Adrian & Welch, Greg. (2005). Ensuring Color Consistency across Multiple Cameras. Proceedings of the IEEE International Conference on Computer Vision. 2. 1268-1275. doi: 10.1109/ICCV.2005.88.

Kim, Y. T. (1997). Contrast enhancement using brightness preserving bi-histogram equalization. IEEE Transactions on Consumer Electronics, 43(1), 1–8.

Manju et al. (2019) "Improved method for enhancing dark images based on CLAHE and morphological reconstruction." Procedia Computer Science 165 (2019): 391-398.

Park, Jaesik & Tai, Yu-Wing & Sinha, Sudipta & Kweon, Inso. (2016). Efficient and Robust Color Consistency for Community Photo Collections. 430-438. doi: 10.1109/CVPR.2016.53.

Singnoo et al (2010). Understanding the Gamma Adjustment of Images.

(sources of images from the internet)

Ramstein: Image: https://www.laptitealsacienne.com/chateaux-ligne-bleue-vosges-episode-3/
Oedenbourg: Image: https://www.alsaceterredechateaux.com/chateaux-et-cites-fortifiees/ramstein/

Reference image				
	Unmodified	**Mean – std transfer**	**Probability Density Function (Pdf) transfer**	**Mean – lab transfer**
Whole Wall				
Top				
Middle				
Bottom				

Table 1. [Ramstein castle] Modifications

Oedenbourg castle

		

Table 2. [Oedenbourg castle] Modifications