CEIG – Spanish Computer Graphics Conference (2021)
A. Chica and L. Ortega (Editors)

Intensity-Guided Exposure Correction for Indoor LiDAR Scans

M. Comino[1], C. Andújar[2], C. Bosch[3], A. Chica[2] and I. Munoz-Pandiella[4]

[1]Universidad Rey Juan Carlos. Spain
[2]Universitat Politècnica de Catalunya. Spain
[3]Universitat de Vic - Universitat Central de Catalunya, Spain
[4]Universitat de Barcelona, Spain

Abstract

Terrestrial Laser Scanners, also known as LiDAR, are often equipped with color cameras so that both infrared and RGB values are measured for each point sample. High-end scanners also provide panoramic High Dynamic Range (HDR) images. Rendering such HDR colors on conventional displays requires a tone-mapping operator, and getting a suitable exposure everywhere on the image can be challenging for 360° indoor scenes with a variety of rooms and illumination sources. In this paper we present a simple-to-implement tone mapping algorithm for HDR panoramas captured by LiDAR equipment. The key idea is to choose, on a per-pixel basis, an exposure correction factor based on the local intensity (infrared reflectivity). Since LiDAR intensity values for indoor scenes are nearly independent from the external illumination, we show that intensity-guided exposure correction often outperforms state-of-the-art tone-mapping operators on this kind of scenes.

1. Introduction

High-end LiDAR scanners, in particular those targeting architecture and cultural heritage applications, often provide two types of outputs: a point cloud dataset and a 360°(usually equirectangular) High Dynamic Range (HDR) image.

Concerning the point cloud, we assume each point sample is described by a tuple (x,y,z,ir,r,g,b), where (x,y,z) are the spatial coordinates of the point, *ir* is the intensity value, typically in the infrared part of the spectrum (around 1550 nm), and (r,g,b) is the color captured by the scanner-mounted camera(s). We do not consider additional attributes (such as number of returns) also captured by LiDAR equipment. The intensity or *ir* value records the amount of energy of the laser beam reflected by the surface. In indoor environments, such value is nearly independent from the scene illumination conditions. Since the radiant flux of LiDAR laser pulses often exceeds 3 kW, this results in irradiance values several orders of magnitude above those achieved by artificial illumination.

Regarding HDR color data, these are typically available as a separate HDR image (e.g. in EXR format). HDR images are captured either by HDR sensors, or by combining multiple images with different exposure values, in the case of LDR sensors. HDR images are required to capture adequately the wide range of luminance levels that exists in real-world environments. However, HDR colors cannot be shown directly on displays supporting a limited range of values. A tone mapping operator is then needed to compress the wide range of HDR images to the limited range of such display devices.

Directly applying a tone mapping operator to a scanner-based HDR panorama results in LDR images that, although sometimes look right, they are too dependent on the existing lighting conditions. This is unsuitable in some applications. For example, cultural heritage users might wish to see the colors closer to what they will look if uniformly illuminated with a standard illuminant. Another problem with traditional tone mapping operators is that the same surface area might show very different LDR colors depending on the scan position the panorama was acquired from. This results in severe color artifacts when multiple scans are combined.

In this paper we propose a simple, local, intensity-guided tone mapping operator for converting HDR images into LDR images. Besides compressing the range of color values, we wish to reduce the effect of the specific (often imperfect) illumination conditions at the time the scan was captured. In this sense, our approach partially shares the goals of Intrinsic Image Decomposition methods, which attempt to separate the reflectance (albedo) and shading (illumination) components of an image. A key observation is that *ir* values are nearly illumination-independent in indoor scenes, and for many materials, near-infrared reflectance correlates reasonably well with luminance [BPMTMML17] (see also Figure 1). Our approach benefits from this fact, and uses *ir* values to guide the selection of a suitable exposure correction factor.

Figure 1: *LDR color and intensity (ir) from a LiDAR scan.*

DOI: 10.2312/ceig.20211357

2. Previous Work

Laser Scan Radiometry In addition to geometry, most LiDAR systems record light intensity, which depends on parameters such as range, incidence angle, environmental conditions and surface reflectance [TC15]. With an appropriate correction, the measured intensity can be directly associated with material reflectance, which is beneficial for object classification and segmentation [KOPW15]. Some studies also found a good agreement between intensity values and measured luminance [BPMTMML17], which has been used in shadow removal [YC14] or laser intensity colorization [OK14].

Color information captured by LiDAR systems has received much less attention. Julin et al. evaluate the quality of color images acquired by different laser scanners, including the Leica RTC360 used in our tests [JKR*20]. By using test charts in lab conditions, they measure color accuracy, sharpness and noise for both LDR and HDR images. However, images captured in such conditions do not show common artifacts appearing in field conditions.

Color reconstruction Some techniques combine multiple RGB [CCCS08] or RGB-D [ZSG*18] images to reconstruct the color of a captured model. The input images are combined according to geometric, topological and colorimetric criteria, but they are not prepared to deal with the extreme lighting defects that tend to be present when capturing models on the field. In structure-from-motion pipelines, additional problems need to be considered, like the large number of input images and their drastically varying characteristics [WMG14]. Improving color balance and equalizing exposure improves the quality of the results [GABR17].

Tone mapping Tone mapping is an active area of research, as demonstrated by the large number of books and surveys on the subject [RHD*10, DLCMM16, OAT*21]. *Global operators* apply a single tone reproduction curve for the whole image [RSSF02, MMM*10]. The general approach is to extract the luminance L_{in} from the input HDR image, compute per-image global statistics, and use these to output a compressed luminance L_{out}. *Local operators* represent an adaptive version of global operators, that apply different mappings depending on luminance values on the neighborhood of the input pixel [Ash02, DMAC03, FBPC11, LFUS06, MDK08, RD05]. Local operators often include additional filters to reduce edge, halo, glare and noise artifacts [OAT*21].

3. Our approach

Given a LiDAR scan consisting of a point cloud P with (x,y,z,ir,r,g,b) values per sample, and a (registered) HDR image $\mathcal{C}_{in}$, we wish to compute an LDR version $\mathcal{C}_{out}$ of it so that it can be shown on conventional displays. In addition, the effect of varying, imperfect lighting conditions is considerably reduced.

Conceptually, our approach has two major steps, described in detail below. First, we precompute a collection of LDR images $\mathcal{C}^0_{out},\ldots\mathcal{C}^k_{out}$ by applying a user-defined tone mapping operator with increasing luminance levels. Many tone mapping algorithms include a parameter to adjust the output luminance, to get lighter or darker images. For example, Reinhard and Devlin approach [RD05] uses a brightness parameter f for this task. Other methods with less intuitive parameters can be adapted by just multiplying the output luminance L_{out} with a varying gain factor g_i. Second, we choose the final output color C_{out} on a per-pixel basis, by choosing the color whose luminance closely matches the *ir* reflectance value.

3.1. Generating LDR images with increasing brightness levels

Although our method can use any tone mapping operator, in our experiments we used a very simple global operator that mainly compresses the high luminances [RSSF02]. For the pixels in $\mathcal{C}^i_{out}$, the output luminance is computed as

$$L^i_{out} = g_i \frac{L_{in}}{L_{in}+1},$$

using a collection of increasing gain values g_i. Then, we use Schlick's color ratios [FLW02] for color reproduction, $C^i_{out} = \left(\frac{C_{in}}{L_{in}}\right)^s L^i_{out}$, where C_{in} and C_{out} are the input and output colors, and s is a saturation parameter (s=1.3 in our experiments). The output is a collection of LDR images with varying brightness levels (Figure 2). Notice that, for a sufficiently large range of gain values, every pixel of the image has an appropriate exposure in at least one of the images. Due to its simplicity, $\mathcal{C}^i_{out}$ images are not precomputed nor stored; their colors C^i_{out} are computed on-the-fly on the GPU. We compute the input luminance L_{in} as the Y component after the standard RGB-to-XYZ conversion, i.e. $L_{in} = Y = 0.2126r + 0.7152g + 0.0722b$.

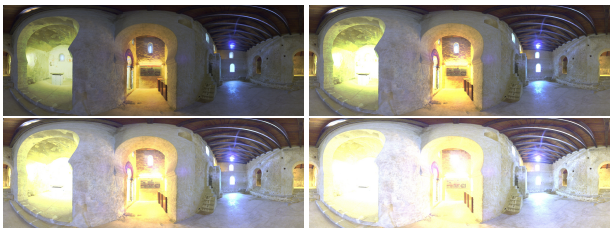

Figure 2: *LDR images with increasing gain factors (1, 2.5, 5, 10).*

3.2. Choosing per-pixel colors

Regardless of the tone mapping operator adopted to generate the collection of LDR images, for each pixel we choose the LDR color C^i_{out} that minimizes the difference between the luminance of C_{out} and the *ir* reflectance. As a final step, we use a standard linear-to-sRGB conversion (gamma correction for $\gamma = 2.4$) to display the LDR images. The rationale of using the *ir* reflectance to chose the final exposure is that, as stated above, LiDAR *ir* values are nearly illumination independent on indoor scenes, and for many materials, *ir* correlates with the luminance of the material's albedo. Figure 3 shows one example. Notice that the LDR image fails to get a correct exposure everywhere, whereas our output seems reasonably close to the surface albedo. Since albedo images naturally look flat, we also show a relighted version of the output.

Figure 3: *Intensity image, one possible LDR image, and the output of our method (as is, and after applying basic Lambert shading with respect to a virtual light placed at the camera).*

4. Results

We implemented our approach on the GPU using a simple fragment shader that runs in real-time on NVidia GeForce RTX 3090 hardware. The test dataset was a collection of scans of St. Quirze de Pedret, captured by a Leyca RTC360 scanner. The dataset is quite challenging because of the extreme variation in the contribution of natural and artificial illumination. Figure 4 shows the results of our method. Our approach succeeds in getting an appropriate exposure for all the pixels across the image. Notice that surface details are visible in all rooms, regardless of their illumination conditions. Despite *ir* values are used to find the proper output luminance, color reproduction is reasonably accurate: all output colors are taken from the collection of LDR images, and thus at least one tone mapping operator produces this color. Note however that color balance is out of the scope of the paper. In the test dataset, illumination profiles were extremely varied, which caused the output images to exhibit clear color biases depending on the surface. Figure 5 compares our approach with state-of-the-art tone mapping operators with default parameters. Although many tone mapping operators produced visually appealing images, none of them managed to get a correct exposure for the darkest areas. In addition, shadows (e.g. scanner tripod) as well as the impact of spotlights on the walls are clearly visible.

5. Conclusions and Future Work

We have presented a simple-to-implement, GPU-based method for processing HDR images associated to LiDAR scans. The main idea is to use the infrared reflectance to choose an LDR color from those produced by tone mapping operators with varying luminance gains. We have shown that our method achieves an appropriate exposure across the image, removing shadows and major lighting effects besides computing an image suitable for conventional displays. Our method has limitations though. First, it assumes Lambertian surfaces (although mirror-specular materials are not supported by LiDAR scans either). Second, it assumes the scene mostly contains materials whose infrared reflectance correlates with luminance. Third, it does not attempt to get a suitable color balance. Actually, as future work, we plan to incorporate a local color balance correction, as well as more sophisticated tone mapping filters for removing edge, halo and glare artifacts.

Acknowledgments This work has been partially funded by the Spanish Ministry of Economy and Competitiveness and FEDER under grant TIN2017-88515-C2-1-R, the *Romanesque Pyrenees, Space of Artistic Confluences II* (PRECA II) project (HAR2017-84451-P, UB) and by EU Horizon 2020, JPICH Conservation, Protection and Use initiative (JPICH-0127) and the Spanish Agencia Estatal de Invesigación (grant PCI2020-111979). The authors would like to thank Prof. Xavier Pueyo for his helpful comments.

References

[Ash02] ASHIKHMIN M.: A tone mapping algorithm for high contrast images. In *13th eurographics workshop on rendering: Pisa, Italy.* (2002), Citeseer. 2, 4

[BPMTMML17] BALAGUER-PUIG M., MOLADA-TEBAR A., MARQUÉS-MATEU A., LERMA J.: Characterisation of intensity values on terrestrial laser scanning for recording enhancement. *International Archives of the Photogrammetry, Remote Sensing & Spatial Information Sciences 42* (2017). 1, 2

[CCCS08] CALLIERI M., CIGNONI P., CORSINI M., SCOPIGNO R.: Masked photo blending: Mapping dense photographic data set on high-resolution sampled 3d models. *Computers & Graphics 32*, 4 (2008), 464–473. 2

[DLCMM16] DUFAUX F., LE CALLET P., MANTIUK R., MRAK M.: *High dynamic range video: from acquisition, to display and applications*. Academic Press, 2016. 2

[DMAC03] DRAGO F., MYSZKOWSKI K., ANNEN T., CHIBA N.: Adaptive logarithmic mapping for displaying high contrast scenes. *Computer Graphics Forum 22*, 3 (2003), 419–426. 2, 4

[FBPC11] FERRADANS S., BERTALMIO M., PROVENZI E., CASELLES V.: An analysis of visual adaptation and contrast perception for tone mapping. *IEEE Transactions on Pattern Analysis and Machine Intelligence 33*, 10 (2011), 2002–2012. 2, 4

[FLW02] FATTAL R., LISCHINSKI D., WERMAN M.: Gradient domain high dynamic range compression. In *Proc. of the 29th annual conference on Computer graphics and interactive techniques* (2002), pp. 249–256. 2, 4

[GABR17] GAIANI M., APOLLONIO F. I., BALLABENI A., REMONDINO F.: Securing color fidelity in 3d architectural heritage scenarios. *Sensors 17*, 11 (2017), 2437. 2

[JKR*20] JULIN A., KURKELA M., RANTANEN T., VIRTANEN J.-P., MAKSIMAINEN M., KUKKO A., KAARTINEN H., VAAJA M. T., HYYPPÄ J., HYYPPÄ H.: Evaluating the quality of TLS point cloud colorization. *Remote Sensing 12*, 17 (2020), 2748. 2

[KOPW15] KASHANI A. G., OLSEN M. J., PARRISH C. E., WILSON N.: A review of lidar radiometric processing: From ad hoc intensity correction to rigorous radiometric calibration. *Sensors 15*, 11 (2015), 28099–28128. 2

[LFUS06] LISCHINSKI D., FARBMAN Z., UYTTENDAELE M., SZELISKI R.: Interactive local adjustment of tonal values. *ACM Transactions on Graphics (TOG) 25*, 3 (2006), 646–653. 2, 4

[MDK08] MANTIUK R., DALY S., KEROFSKY L.: Display adaptive tone mapping. In *ACM SIGGRAPH 2008*. 2008, pp. 1–10. 2, 4

[MMM*10] MAI Z., MANSOUR H., MANTIUK R., NASIOPOULOS P., WARD R., HEIDRICH W.: Optimizing a tone curve for backward-compatible high dynamic range image and video compression. *IEEE transactions on image processing 20*, 6 (2010), 1558–1571. 2, 4

[OAT*21] OU Y., AMBALATHANKANDY P., TAKAMAEDA S., MOTOMURA M., ASAI T., IKEBE M.: Real-time tone mapping: a survey and cross-implementation hardware benchmark. *IEEE Transactions on Circuits and Systems for Video Technology* (2021), 1–1. 2

[OK14] OISHI S., KURAZUME R.: Manual/automatic colorization for three-dimensional geometric models utilizing laser reflectivity. *Advanced Robotics 28*, 24 (2014), 1637–1651. 2

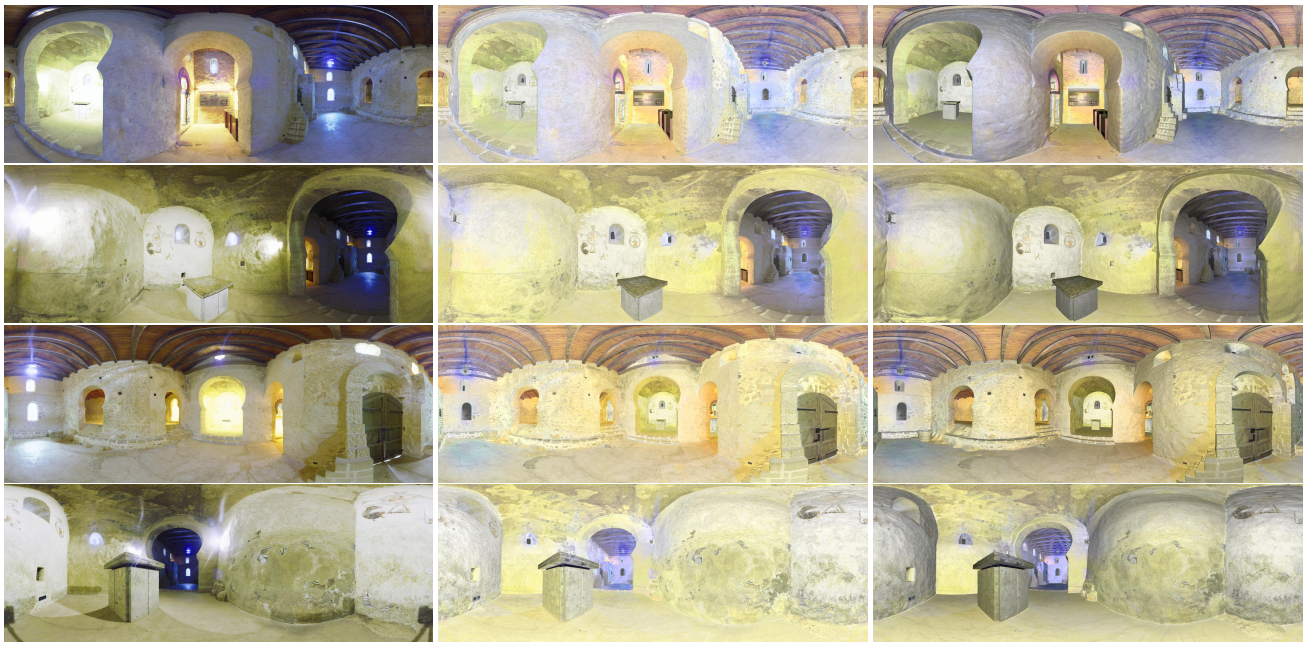

Figure 4: *Results with several scans. Each row shows one LDR image and the output of our approach (unlit/relighted).*

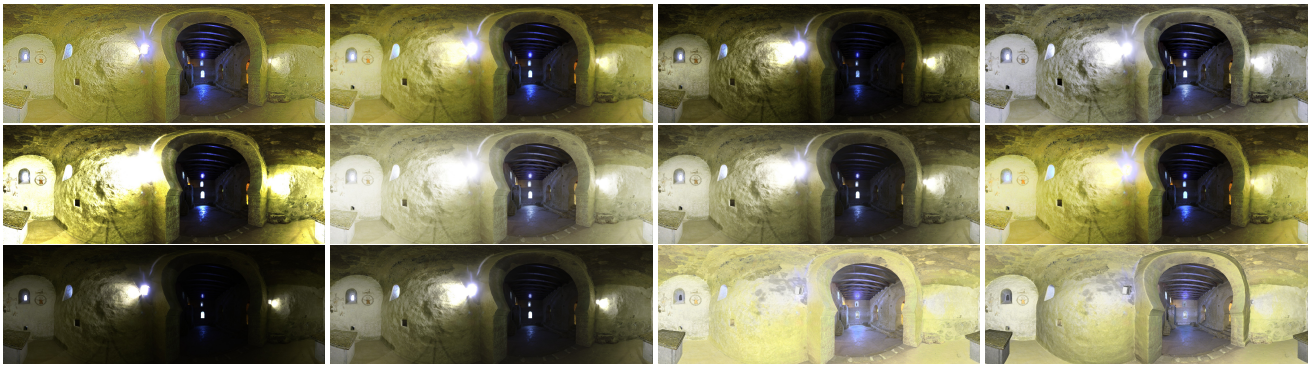

Figure 5: *Comparison with common tone mapping operators. We used the default parameters recommended in Luminance HDR 2.6.0 software. First row: Ashikhmin [Ash02], Drago et al. [DMAC03], Fattal et al. [FLW02] and Ferradance et al. [FBPC11]. Second row: Lischinski et al. [LFUS06], Mai et al. [MMM*10], Mantiuk et al. [MDK08] and Reinhard et al. [RSSF02]. Third row: Reinhard et al. [RD05], Van Hateren [VH06] and our approach (unlit/relit).*

[RD05] REINHARD E., DEVLIN K.: Dynamic range reduction inspired by photoreceptor physiology. *IEEE transactions on visualization and computer graphics 11*, 1 (2005), 13–24. 2, 4

[RHD*10] REINHARD E., HEIDRICH W., DEBEVEC P., PATTANAIK S., WARD G., MYSZKOWSKI K.: *High dynamic range imaging: acquisition, display, and image-based lighting*. Morgan Kaufmann, 2010. 2

[RSSF02] REINHARD E., STARK M., SHIRLEY P., FERWERDA J.: Photographic tone reproduction for digital images. In *Proc. of the 29th annual conference on Computer graphics and interactive techniques* (2002), pp. 267–276. 2, 4

[TC15] TAN K., CHENG X.: Intensity data correction based on incidence angle and distance for terrestrial laser scanner. *Journal of Applied Remote Sensing 9*, 1 (2015), 094094. 2

[VH06] VAN HATEREN J. H.: Encoding of high dynamic range video with a model of human cones. *ACM Transactions on Graphics (TOG) 25*, 4 (2006), 1380–1399. 4

[WMG14] WAECHTER M., MOEHRLE N., GOESELE M.: Let there be color! large-scale texturing of 3d reconstructions. In *European conference on computer vision* (2014), Springer, pp. 836–850. 2

[YC14] YAGISHITA D., CHIKATSU H.: Generation of effective orthophotos for road surfaces using MMS. *International Archives of the Photogrammetry, Remote Sensing & Spatial Information Sciences 45* (2014). 2

[ZSG*18] ZOLLHÖFER M., STOTKO P., GÖRLITZ A., THEOBALT C., NIESSNER M., KLEIN R., KOLB A.: State of the art on 3d reconstruction with rgb-d cameras. In *Computer graphics forum* (2018), vol. 37(2), Wiley Online Library, pp. 625–652. 2