Bound-Constrained Optimized Dynamic Range Compression

Dorian Yao Chan
dychan@andrew.cmu.edu
Carnegie Mellon University

James F. O'Brien
job@berkeley.edu
University of California, Berkeley

Figure 1: Our dynamic range compression method applied to different natural HDR images.

ABSTRACT

We present a new spatially-varying dynamic range compression algorithm for high dynamic range (HDR) images based on bound-constrained optimization using soft constraints. Rather than explicitly attenuating gradients as in previous work, we minimize an objective function to instead compute a globally optimal manipulation of input pixel differences. Our framework provides simple yet effective preservation of visually important image properties, such as order statistics and global consistency, that requires little to no parameter tuning. Our results are free of haloing, washout, and other artifacts, while retaining detail across the image's full range. The speed of our algorithm and flexibility of the constraint framework allows our method to be easily extended to video.

ACM Reference Format:
Dorian Yao Chan and James F. O'Brien. 2020. Bound-Constrained Optimized Dynamic Range Compression. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Posters (SIGGRAPH '20 Posters), August 17, 2020.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3388770.3407415

1 INTRODUCTION

Thanks to the ubiquity of high dynamic range (HDR) imaging, numerous dynamic range compression techniques have been developed over the years, in order to display this high dynamic range content on low dynamic range devices and software. These tone-mapping operators can be broadly grouped into two classes: global and local. Global operators work across the entire input image, mapping every input luminance to some output luminance according to a global adjustment curve at the cost of washed-out details. In contrast, in a local operator, the output intensity of each pixel depends only on its local neighborhood, resulting in strong detail preservation at the cost of spatial artifacting such as haloing and ringing. In this work, we propose a new technique that uses an optimization framework to maximize the level of detail preservation, while ensuring that the results remain artifact-free.

SIGGRAPH '20 Posters, August 17, 2020, Virtual Event, USA

ACM ISBN 978-1-4503-7973-1/20/08.
https://doi.org/10.1145/3388770.3407415

2 APPROACH

Dynamic range compression fundamentally seeks an image that maximally preserves the visual aspects of the input radiance map while using a smaller, specified dynamic range. If we quantify "visual aspects" in terms of pixel differences [Fattal et al. 2002; Mantiuk et al. 2006], we can write a penalty function $F(l_i - l_j, h_i - h_j, i, j)$ that represents the error for output log intensities l_i and l_j corresponding to input log intensities h_i and h_j at pixels i and j. Armed with this function, we can now formulate dynamic range compression as a bound-constrained optimization problem:

$$\begin{aligned} \underset{l}{\text{minimize}} \quad & \sum_{i \in I} \sum_{j \in N_i} F(l_i - l_j, h_i - h_j, i, j) \\ \text{subject to} \quad & L_{min} \leq l_i \leq L_{max}, \ i = 1, \ldots, N \end{aligned} \tag{1}$$

where I is the set of all pixels in the image and N_i denotes the chosen neighborhood of pixel i. The constraints bound the dynamic range of the output image l to $[L_{min}, L_{max}]$, ensuring that the output of our algorithm optimally preserves detail and contrast for the desired dynamic range. As a result, our method does not explicitly attenuate pixel differences, as in existing gradient-domain methods [Fattal et al. 2002; Mantiuk et al. 2006]. Instead, the process of bound-constrained optimization determines the optimal attenuation according to the selected error function.

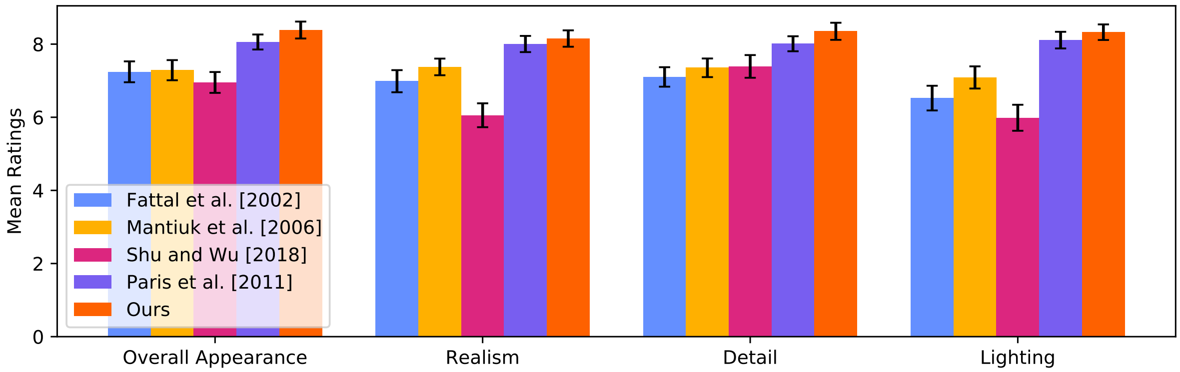

Figure 2: Average ratings with 95% confidence intervals from our subjective study. Our method achieved the highest average scores in every tested category.

We further decompose $F(l_i - l_j, h_i - h_j, i, j)$ as the sum of a pixel difference preservation term $P(l_i - l_j, h_i - h_j)$ and an order statistic term $O(l_i - l_j, h_i - h_j)$, scaled by a weighting function $W(h_i - h_j, i, j)$:

$$F(l_i - l_j, h_i - h_j, i, j) = W(h_i - h_j, i, j)\big(P(l_i - l_j, h_i - h_j) + O(l_i - l_j, h_i - h_j)\big) \tag{2}$$

We describe our exact choices for these functions in further detail in the supplement. This method can easily be applied to video by simply extending each pixel's neighborhood temporally, which we describe in the supplement.

The combination of the weighting function $W(h_i - h_j, i, j)$ and the preservation function $P(l_i - l_j, h_i - h_j)$ control the level at which each input pixel difference is preserved in the output image. Pixel differences with a larger weight value given by $W(\cdot)$ will be more likely to be preserved by the optimization process, while pixel differences corresponding to smaller $W(\cdot)$ will most likely be compressed. The choice of $P(\cdot)$ is also equally important. For example, using a squared error term would encourage many variations from the input pixel differences to avoid one large variation. In contrast, using an absolute error term would concentrate the variations from the inputs in as few pixel differences as possible. In essence, like how gradient-domain operators manipulate the input image using an attenuation function, our approach allows for manipulation through careful design of the penalty function $F(\cdot)$. However, unlike previous gradient-domain operators, our approach never unnecessarily attenuates an input gradient, potentially increasing the level of detail preservation.

The order statistic term $O(l_i - l_j, h_i - h_j)$ penalizes changes in luminance ordering between the input and output images, important for preventing spatial artifacting [Shu and Wu 2018]. Global operators, thanks to their monotonic tone curve design, fundamentally preserve order statistics. Therefore, if N_i is the entire image I and $O(\cdot) = \infty$ for violations of the order statistic, our optimization is a perfect global operator. In contrast, if N_i is just the immediate neighborhood of pixel i and $O(\cdot) = 0$ for violations, our optimization is a perfect local operator. Thus, $O(\cdot)$ and N_i provide a sliding knob between an ideal local operator and an ideal global operator.

3 RESULTS

In order to solve the optimization problem given in section 1, we rely on a hierarchical gradient descent solver with momentum. The hierarchical solver builds a Gaussian pyramid, solves the optimization for each level, and initializes the next level with the output of the previous one. To handle bound constraints, we simply clip invalid intensities at every iteration. With this setup, we achieve convergence within a hundred iterations at all but the coarsest level in every tested case. We find that our method compares favorably to previous state-of-the-art operators [Fattal et al. 2002; Mantiuk et al. 2006; Paris et al. 2011; Shu and Wu 2018] in a user study in figure 2. We show further results in the supplemental.

An important effect of our formulation is its improved handling of wash-out, as shown in figure 3. Our formulation explicitly uses the output dynamic range as part of the optimization process, removing the need for clipping. Simultaneously, thanks to its attenuate-as-needed approach, many of the high luminance details are not pointlessly reduced under our formulation.

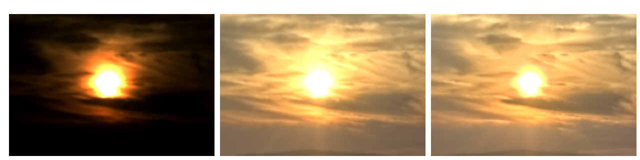

(a) Clipped Input **(b) Gradient-domain** **(c) Our method**

Figure 3: Without special care, the details in the clouds near the sun are washed-out after compression.

REFERENCES

Raanan Fattal, Dani Lischinski, and Michael Werman. 2002. Gradient Domain High Dynamic Range Compression. *ACM Trans. Graph.* 21, 3 (July 2002), 249–256.

Rafal Mantiuk, Karol Myszkowski, and Hans-Peter Seidel. 2006. A Perceptual Framework for Contrast Processing of High Dynamic Range Images. *ACM Trans. Appl. Percept.* 3, 3 (July 2006), 286–308.

Sylvain Paris, Samuel W. Hasinoff, and Jan Kautz. 2011. Local Laplacian Filters: Edge-aware Image Processing with a Laplacian Pyramid. *ACM Trans. Graph.* 30, 4, Article 68 (July 2011), 12 pages.

Xiao Shu and Xiaolin Wu. 2018. Locally Adaptive Rank-Constrained Optimal Tone Mapping. *ACM Trans. Graph.* 37, 3, Article 38 (July 2018), 10 pages.