Contrast Use Metrics for Tone Mapping Images

Miguel Granados[1,2] Tunç Ozan Aydın[2] J. Rafael Tena[2] Jean-François Lalonde[3] Christian Theobalt[1]

[1]MPI for Informatics [2]Disney Research [3]Laval University

Abstract

Existing tone mapping operators (TMOs) provide good results in well-lit scenes, but often perform poorly on images in low light conditions. In these scenes, noise is prevalent and gets amplified by TMOs, as they confuse contrast created by noise with contrast created by the scene. This paper presents a principled approach to produce tone mapped images with less visible noise. For this purpose, we leverage established models of camera noise and human contrast perception to design two new quality scores: contrast waste *and* contrast loss*, which measure image quality as a function of contrast allocation. To produce tone mappings with less visible noise, we apply these scores in two ways: first, to automatically tune the parameters of existing TMOs to reduce the amount of noise they produce; and second, to propose a new noise-aware tone curve.*

(a) $\beta = 0.85, \alpha = 0.2$ (b) $\beta = 0.9, \alpha = 0.7$

Figure 1. HDR images of low-light scenes contain camera noise that can be amplified by TMOs. The amount of amplification depends on the TMO and its parameters. This HDR image was tone mapped using [3] with different parameters (β, α), but the relation between the parameters and noise amplification or detail loss is unknown to the users. We present new metrics that capture this relationship, allowing a user to intuitively browse the parameter space of a TMO and quickly choose a good combination.

1. Introduction

High dynamic range (HDR) images can easily be captured nowadays, even with consumer cameras. To properly view this HDR content on low dynamic range displays, one needs a tone mapping operator (TMO) to map it to the limited displayable range, while retaining as much of its original contrast as possible [15]. The sheer number of different algorithms proposed in the literature is testament to the complexity of this task: they must adapt to different displays, be free of visual artifacts, and provide intuitive artistic controls to allow users to achieve desired visual styles.

Despite these challenges, today's powerful tone mapping operators have been very successful and have found their way into a wide variety of consumer photography applications. However, while they work remarkably well on images taken under daylight conditions or in well-lit indoor scenes, they often produce very objectionable artifacts on images taken under low light conditions because these images might contain significant sensor noise (fig. 1). Noise gets amplified in ways that depend both on the particular TMO used and on the values of its parameters. This makes tone mapping a tedious, trial-and-error process, where the user must try several parameter settings individually to find the desired result.

In this paper, we introduce two new quantitative metrics that capture how effectively a tone mapping operator utilizes the available output (display) range to preserve the original contrast while keeping the noise visually imperceptible. To develop these measures, we leverage existing models of camera noise and of human perception. We demonstrate the usefulness of these metrics in two potential applications. First, we show how they can be used to automatically find a combination of parameters which will yield the best tone mapped result for a given noisy HDR input. Since manually exploring the space of possible tone mapped images for a given TMO can be laborious, our method provides an intuitive way to visualize the space of TMO parameters in a noise-aware way. Second, we can design noise-optimal tone curves which directly optimize these measures to create a tone-mapped image that best exploits the output range in the presence of noise.

2. Related work

Tone mapping has been an active research topic in computer graphics for nearly two decades [15]. Early work involved analyzing common practices of film development and applying them to the field of HDR imaging. Reinhard et al. [14] proposed applying a sigmoidal response curve globally and performing local operations to mimic photographic dodging and burning. While this operator comprises local components, its results are often a faithful reproduction of the original scene's contrast and colors as it would be experienced by a human observer. A different look with higher local contrast can be achieved using a bilateral-filtering-based tone mapping approach [2]. This method produces a *base layer* from the input HDR image through bilateral filtering. A corresponding *detail layer* is computed by the ratio of the original HDR and the base layer. Tone mapping is achieved by applying a compressive tone curve to the base layer and combining the result with the detail layer. Reinhard and Devlin's TMO [13] is inspired by the photoreceptors response in the human eye. The parameters simulate in part the behavior of the human visual system with respect to global and local adaptation to the overall luminance and particular chrominance in the image. While there are many tone mapping operators, in this work we focus on the Photographic TMO [14] and the Bilateral TMO [2] as two prominent representatives of global and local tone mapping operators.

Previous tone mapping work focused on simulating the visual perception of extremely dark and bright HDR scenes [4, 7, 12]. The main aim was to model the luminance adaptation mechanisms of the human visual system assuming an HDR image free of any artifacts, yet this assumption does not hold for low-light scenes where camera noise is significantly present. Another research area on low-light tone mapping explores the hue shifts that occur in dark scenes [10], but a solution for obtaining visually pleasant results in the presence of camera noise is not provided.

The noise properties of digital cameras have been studied in the field of optics and photonics [9]. The two principal noise sources are shot noise, associated with the process of light emission, and readout noise, which is an umbrella term for several sources that affect the image capturing process. These two sources affect each pixel individually. In this paper, we apply a simplified noise model (see Sec. 3.1) that takes into account these major sources and ignores other spatially dependent sources. The parameters of this model can be recovered from a set of calibration images [9] or from regularly captured images [5, 11]. In this work, we assume that a calibrated camera noise model is available. The next sections explain how this model can be used to measure the image quality of existing tone mapping operators, and how it enables noise-aware TMOs with greatly enhanced performance on low-light images.

3. Evaluation of contrast utilization in TMOs

We begin by describing an approach to measure the effectiveness of a TMO in allocating the available display contrast when tone mapping a high dynamic range image. This is a challenging task that becomes more difficult in situations where noise is dominant, such as low-light conditions. In these cases, existing tone mapping operators may inadvertently boost the noise in the image (see fig. 1). We argue that an effective use of the contrast range means succeeding at two potentially conflicting tasks: *preserving* the original contrast of the input image; while *preventing* the amplification of noise. In this section, we first describe the camera noise model and the visual perception model that are the foundation of our work. Based on these models, we then introduce two new quality measures to assess a TMO's performance: *i) contrast waste*, and *ii) contrast loss*.

3.1. A model of camera noise

By calibrating the noise parameters of a digital camera, it is possible to predict the noise distribution of the color values in the images it captures. To estimate image noise, we apply the calibration method in [5] to the raw output of digital cameras. This calibration needs to be performed once, offline, for a given camera model; it could also be provided by the manufacturer. Calibration yields a noise model defined by camera-dependent and photograph-dependent parameters. The four camera-dependent parameters are the camera gain at reference ISO level G_{100}, the black level $v_{\min}$, the saturation level $v_{\max}$, and the readout noise σ_R. The two photograph-dependent parameters are the ISO value S and the exposure time t. We can approximate the variance of the Gaussian probability distribution for a pixel p in an input image I at the ISO level S as [8]:

$$\sigma_I^2(p) \approx G_s(I(p) - v_{\min}) + \sigma_R^2, \tag{1}$$

where $G_S = \frac{S}{100} G_{100}$ is the camera gain at ISO level S.

This model predicts the noise distribution in raw images, which have a higher bit depth than standard 8-bit displays. It can also be used to predict the noise of HDR images obtained from averaging raw multi-exposure sequences. Let $\mathbf{I} = \{I_1, \ldots, I_n\}$ be a multi-exposure sequence with exposure times t_i and ISO values S_i. Each image I_i provides the irradiance estimate

$$X_i(p) = \frac{I(p) - v_{\min}}{G_{S_i} \cdot t_i}, \text{ with variance} \tag{2}$$

$$\sigma_{X_i}^2 \approx \frac{\sigma_{I_i}^2(p)}{(G_{S_i} \cdot t_i)^2}. \tag{3}$$

An irradiance map, or HDR image, X can be obtained from

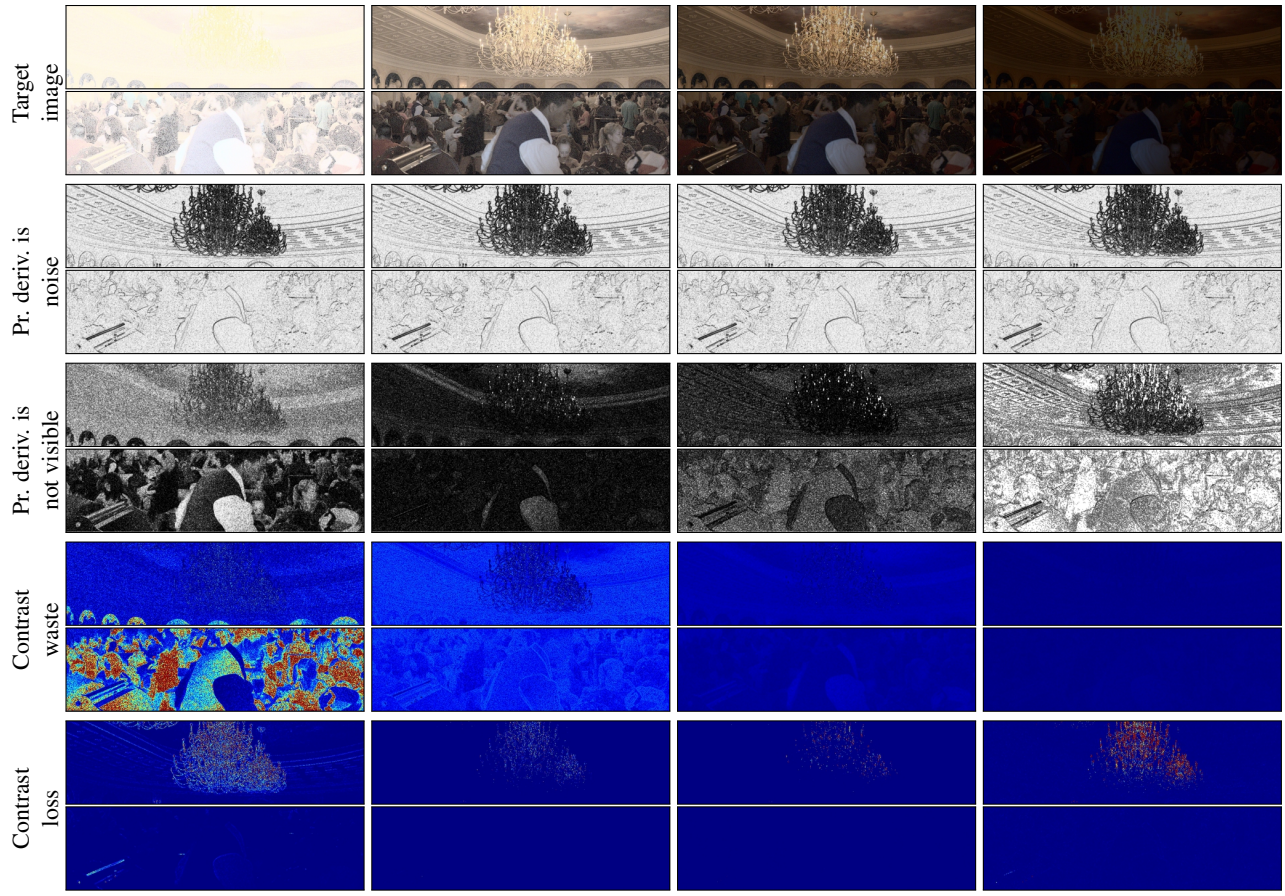

Figure 2. Examples of the contrast waste and contrast loss maps for four different settings of the TMO in [13]. Each row shows a top and bottom crop of the photograph shown in Fig. 3. High contrast loss (last row) occurs at pixel locations where the real image derivatives (darker pixels in 2nd row) are no longer perceivable (bright pixels in 3rd row). High contrast waste (4th row) occurs whenever derivatives attributable to noise are displayed in the tone mapped image (best seen in PDF).

the weighted average

$$X(p) = \frac{\sum_i w_i(p) X_i(p)}{\sum_i w_i(p)}, \text{ with variance} \quad (4)$$

$$\sigma_X^2(p) \approx \frac{\sum_i w_i(p)^2 \sigma_{X_i}^2(p)}{\left(\sum_i w_i(p)\right)^2}. \quad (5)$$

In the remainder of the paper, we assume that the input image I and its variance σ_I^2 are known or recovered using a similar calibration procedure. We discontinue the use of X and use only I instead.

3.2. Detection of image derivatives caused by noise

Let the input image be an HDR image $I : \Omega \to \mathbf{R}$ where each pixel p is an observation of a random variable that follows a Gaussian distribution with mean $\hat{I}(p)$ and standard deviation $\sigma_I(p)$ (estimated accd. to Sec. 3.1). Let $p, q \in \Omega$ be two adjacent pixel locations, and let $D(p,q) = I(p) - I(q)$ be an approximation of the image derivative at $I(p)$. $D(p,q)$ also follows a Gaussian distribution with mean $\hat{D}(p,q) = \hat{I}(p) - \hat{I}(q)$ and standard deviation $\sigma_D(p,q) = \sigma_I(p) + \sigma_I(q)$. Whenever the image is flat at $I(p)$, $\hat{I}(p) = \hat{I}(q)$, the mean of the derivative's distribution is zero. Therefore, to test whether the observed derivative is caused by noise we define the null hypothesis H_0 and the alternative hypothesis H_1 as

- H_0: The observed derivative $D(p,q)$ is generated by the distribution $\mathcal{N}(0, \sigma_D(p,q))$, and
- H_1: The observed derivative $D(p,q)$ is *not* generated by the distribution $\mathcal{N}(0, \sigma_D(p,q))$.

The probability of rejecting H_0 incorrectly (type I error) should be bounded by a confidence value α as

$$\Pr\left(\text{rejecting } H_0 | H_0 \text{ is true}\right) \equiv \Pr\left(Z > |z_D(p,q)|\right) < \alpha, \quad (6)$$

where Z is a random variable with normal distribution, and $z_D(p,q) = \frac{D(p,q)}{\sigma_D(p,q)}$ is the statistical standard score or z-value of the observed derivative. The probability in Eq. 6

captures the percentage of derivatives due to noise that are larger than $D(p,q)$. Since our goal is to misclassify as few derivatives due to noise as possible, the confidence value α is set to an arbitrary low value (e.g., 1%). If the probability of observing a derivative larger than $D(p,q)$ is larger than α, we reject the alternative hypothesis and accept that $D(p,q)$ is generated by the distribution of the image noise. The result of this test is encoded in a mask image

$$M(p,q) = \mathbf{1}_{\{\Pr(Z>|z_D(p,q)|)>\alpha\}}, \tag{7}$$

that assigns the value of 1 to derivatives $D(p,q)$ that are attributable to camera noise (see fig. 2, second row).

3.3. Detection of perceptible visual differences

Our visual perception model consists of a predictor that tests if two intensities are visually indistinguishable to the observer. Let I^t be a tone mapped version of the input image I. Assuming a standard display with sRGB response function, $\gamma \approx 2.2$, and luminance range $[L_{\min}, L_{\max}]$, we construct the image $I^t_L = (I^t/\max(I^t))^{\frac{1}{\gamma}} \cdot (L_{\max} - L_{\min}) + L_{\min}$ whose values approximate the luminance emitted by the display. For each value I^t_L, the contrast sensitivity function $\mathrm{csf}(L,\cdot) = \Delta L$ predicts the minimum luminance offset ΔL from an observed luminance L that is necessary for the difference to be perceivable in 75% of the cases under standard illumination and adaptation conditions. This threshold depends on the particular viewing conditions (e.g. the viewing distance and the screen's pixel size), and the frequency of the signal (measured in cycles per degree). Since we evaluate noise perception between pairs of adjacent pixels, our target frequency corresponds to half of the pixels per degree (i.e. a cycle is produced at every pair of adjacent pixels)[1]. Based on the contrast sensitivity function, the probability $V(p,q)$ that a user detects a luminance difference is

$$V(p,q) = 1 - \exp\left(-|\phi z(p,q)|^3\right), \text{where} \tag{8}$$

$$z(p,q) = \frac{I^t_L(p) - I^t_L(q)}{\max\{\mathrm{csf}(I^t_L(p)), \mathrm{csf}(I^t_L(q))\}}, \tag{9}$$

and $\phi = (-\log(1-0.75))^{\frac{1}{3}}$ [1].

3.4. Contrast waste score

The *contrast waste* score for a tone-mapped image I^t measures how many pairs of adjacent pixels, whose colors in the input image are indistinguishable under noise, are mapped to screen values whose luminance differences are likely to be detected by the user, in which case contrast is wasted. For an adjacent pixel pair p, q, it is defined as the normalized perceivable luminance difference between the pixels times the probability that both pixels measure the same luminance:

$$W(p,q) = M(p,q)V(p,q)\left|I^t_L(p) - I^t_L(q)\right|. \tag{10}$$

The aggregate waste score for the entire image is the average per-pixel waste score,

$$\bar{W} = \frac{1}{|\mathcal{N}(\Omega)|} \sum_{p,q\in\mathcal{N}(\Omega)} W(p,q), \tag{11}$$

where $\mathcal{N}$ is an 8-neighborhood system in the image domain Ω. Fig. 2 (4th row) illustrates the contrast waste produced by the same tone mapper [13] with different parameters.

3.5. Contrast loss score

The *contrast loss* score estimates how many luminance differences are missing in a tone-mapped version I^t of an image I. This loss of contrast occurs at image locations whose derivatives are not attributable to noise, but their corresponding tone-mapped values are visually indistinguishable. For a pair of pixels p, q the score is computed as the loss of perceivable luminance differences in the tone-mapped image with respect to a standard tone mapping procedure, such as a linear intensity mapping.

$$L(p,q) = (1 - M(p,q))\,(1 - V(p,q))\,|I^r(p) - I^r(q)|. \tag{12}$$

Here, I^r is a reference tone mapping of I, such as $I^r(p) = (I(p)/\max(I))$. The aggregate contrast loss of the image is the average of the per-pixel scores

$$\bar{L} = \frac{1}{|\mathcal{N}(\Omega)|} \sum_{(p,q)\in\mathcal{N}(\Omega)} L(p,q). \tag{13}$$

Fig. 2 (last row) shows the contrast loss produced by the TMO in [13] with different parameters.

3.6. Constrast misuse score

The waste and loss scores can guide the choice of TMO and parameters for a given scene (sec. 4.1), and can be used to define criteria for globally optimal tone curves (sec. 4.2). For these purposes, it is useful to define a single *contrast misuse score* that represents the contrast waste and contrast loss of a given image and additionally encodes the user-preference regarding the balance between these two types of artifacts

$$\bar{S} = (1-\beta)\bar{W} + \beta\bar{L}. \tag{14}$$

Here, $\beta \in [0,1]$ represents the relative importance of contrast waste and loss. For $\beta = 0$, optimal tone-mapped images will not contain visible noise artifacts but may suffer from detail loss. Conversely, at $\beta = 1$ the image will preserve the details of the input HDR image but will display noise artifacts. In our experiments, we set $\beta = 0.9$ to preserve details while allowing some noise artifacts.

[1] In our experiments, we used a 20 inch display at 1600×1200 resolution viewed at 25 inch distance, resulting into 45 pixels per degree, and 22.5 cycles per degree.

4. Results

We demonstrate the usefulness of our scores in two application scenarios.

4.1. Application I: visualizing TMO parameters

To tone map an HDR image, users must choose a particular TMO and values for its parameters. To a novice, this process is unintuitive, and may involve several iterations of trial and error. To complicate the situation, TMOs can produce noise artifacts for a wide range of parameter configurations (see fig. 1). The chosen operator may also not generalize well to other cameras or scenes. Accordingly, providing users with a quick and intuitive way to navigate the space of TMOs and their parameters would be beneficial. Our new quality scores can provide the user with information regarding the suitability of different values of TMO parameters, and even suggest noise-optimal values.

Fig. 3 illustrates how a user can easily explore the TMO parameter space in a noise sensitive way. This example explores the tone mapper from [13], which makes use of two main parameters: the "contrast" and "intensity" parameters. The 2D waste/loss plot in fig. 3 represents the space spanned by these two parameters. Contrast waste (assigned to the green channel) and loss (assigned to the red channel) scores are computed for a discrete set of parameter combinations regularly sampled over that space. By this means, the impact of parameter combinations can be predicted without having to scrutinize tone-mapped images directly: parameters that generate high contrast waste (bright green), high contrast loss (bright red), or noise-optimal results (black) can be identified at a glance. Fig. 3 shows four example locations where a user might click to observe the influence of tone-mapping parameters on the quality of the results. The best result is obtained when the sum of contrast waste and loss scores is minimized (bottom right).

By design, our scores assess effective contrast preservation and noise suppression in an image, which are both results of complex and highly subjective cognitive processes. As such, formulating metrics that cover their every aspect is highly challenging, if at all possible. That said, practical metrics that achieve even some level of correlation with these complex tasks are useful in practice, an example being the SSIM metric for image quality assessment [16]. Similarly, our measures provide a useful practical estimate correlating with a highly challenging task.

Experiments To empirically test the use of the contrast waste score for TMO parameter selection, we acquired a set of 11 photographs in low light conditions using a Canon EOS 5D Mark III with calibrated noise model (see sec. 3.1). The photographs were taken either indoors or at nighttime, and without a flash. All images were stored in RAW uncompressed format. We consider these RAW images as HDR images, since pixel values are proportional to scene luminance and stored at high bit depth.

Fig. 4, compares results for the same input image using different tone mappers [2, 3, 13, 14] and parameters. For each TMO, we empirically selected one or two of the most relevant parameters of each algorithm[2]. We selected the best and worst values for each set of parameters according to the contrast misuse score, and show the images corresponding to the best, worst, and default parameters of each TMO. Fig. 4-d, localizes the best (green), default (blue), and worst (red) parameter sets according to the contrast misuse score in a 2D waste/loss plot (see sec. 4.1).

The runtime of the parameter selection depends on the speed of the actual tone mappers. To construct each waste/loss map, we sampled a 9×9 grid on the plane defined by the two selected parameters for each TMO. For each grid point, a tone mapped image and its contrast misuse score are computed. The scores for intermediate parameters are bilinearly interpolated. On average, the construction of a waste/loss map takes around two minutes using unoptimized MATLAB code for a 1449×968 image.

We draw two conclusions from the results shown in fig. 4. First, the perceived quality of the results is empirically correlated with the contrast score in all TMOs since the best result contains less visible noise than the worst result, without incurring detail loss. It could be argued that the worst result of Durand and Dorsey [2] can be perceptually preferable, despite clearly visible noise, since darker parts become brighter at this parameter setting. This perceptible preference can be accounted for in our algorithm by optimizing for contrast loss only. Experimentally, if contrast waste is ignored (i.e. setting $\beta = 1$), the former worst parameter setting now obtains the best score on this scene and TMO. Therefore, our algorithm has the flexibility to express the user intent through selection of the β parameter. Second, a TMO's default parameters can significantly differ from the optimal parameters, and our algorithm provides a systematic way to select suitable values for a given input image. Fig. 5 presents additional comparisons, and the supplementary material shows a systematic comparison of our method on four TMOs, on all of our 11 test images.

4.2. Application II: noise-aware tone curves

In this section, we propose another application of our novel HDR noise metrics. That is, we present a simple, yet effective strategy to generate a noise-optimal tone curve that can be adapted to a given image. Our approach explicitly shapes the tone curve to avoid, as much as possible, the

[2]Since the method of Durand and Dorsey does not have parameters that cause high variation in the results, we used the two first components of the PCA model of camera response curves (see sec. 4.2) as parameters to control the shape of the tone curve used to compress the base layer.

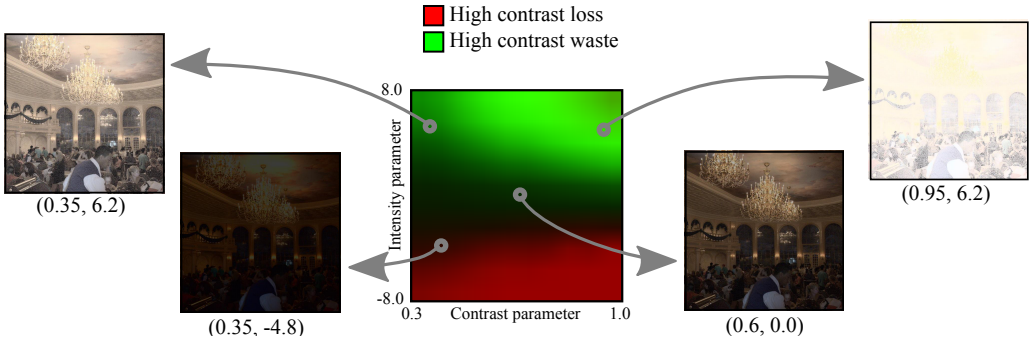

Figure 3. Contrast scores for different configurations of the TMO in [13]. The central *waste/loss plot* shows a color map of the contrast waste and loss scores obtained for combinations of two parameters of the TMO, contrast and intensity. Contrast waste is depicted in green, contrast loss in red. Low scores correspond to values close to black. Noise is more apparent in images with high contrast waste (top left), high contrast loss makes images look washed out (bottom left). Images with high contrast waste and loss scores look noisy and flat (top right). When values for both scores are low, resulting images use the available display contrast optimally (best seen in PDF).

conditions of contrast waste and contrast loss in the result. Fig. 6 presents the approach. First, we use the PCA model of camera response curves from [6] to define the space of possible tone curves. Then, we sample the first two components of the PCA model, and select the component weights $p = (\cdot, \cdot)$ that produce the minimum contrast misuse $\bar{S}$. With an interface similar to the one presented in sec. 4.1, the user can quickly browse the space of tone-curve parameters, and control the trade-off between contrast waste and contrast loss by adjusting the parameter β.

5. Conclusion and future work

In conclusion, this paper proposed two metrics—contrast waste and contrast loss—that measure the efficiency of existing TMOs in allocating the available display contrast. The metrics are based on camera noise and contrast perception models. We further applied these models to propose a principled way to 1) improve the robustness of TMOs in low light conditions by allowing a user to intuitively navigate the space of TMO parameters; and 2) create noise-aware tone curves. Through an empirical validation, we showed that the robustness of existing tone mapping operators can be improved automatically by including these models in the selection of adequate parameters. Our method enables users to obtain feedback about the expected quality of existing tone mappers, and to apply them reliably in automatic settings, even for images in low light conditions.

Currently, contrast is only evaluated on adjacent pixels, so we model only its high frequency content. Capturing lower contrast frequencies would be possible with the use of an image pyramid, which we plan to explore next. In addition, the proposed visualization scheme in sec. 4.1 is only practical for a pair of TMO parameters at a time, a higher dimensional space could only be seen one 2-D slice at a time. Similarly, the automatic selection of the optimal set of parameters for a given image would become exponentially slower to compute for higher dimensional spaces.

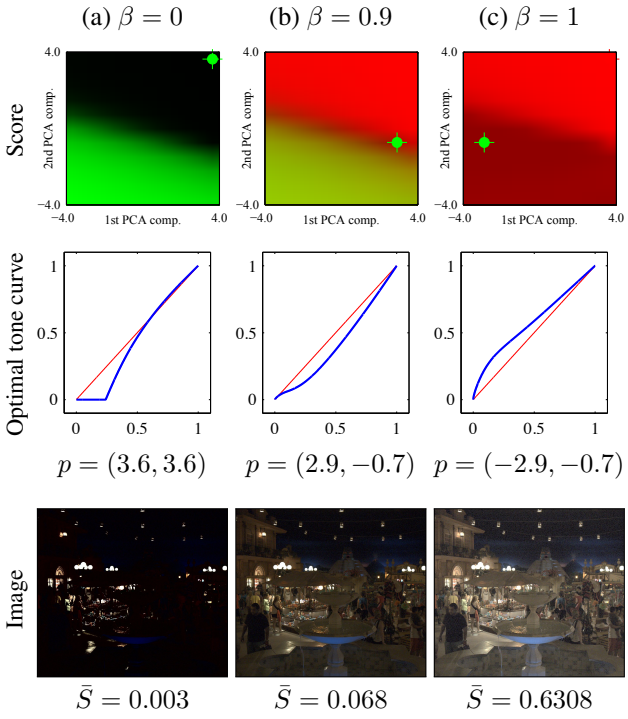

Figure 6. Tone curves with low contrast misuse sampled from a PCA model of camera responses. Top row: *waste/loss plot* for PCA parameters (red: high loss, green: high waste). Middle row: tone curve (blue) with minimum contrast misuse $\bar{S}$. Bottom: Image tone-mapped with the optimal tone curve. The optimal tone curve depends on the relative importance of contrast waste and loss set by the user: (a) Ignoring contrast loss ($\beta = 0$), results in a dark noiseless image. (b) Increasing the weight of contrast loss ($\beta = 0.9$), reduces detail loss while suppressing noise. (c) Ignoring contrast waste ($\beta = 1$) results in visible noise artifacts.

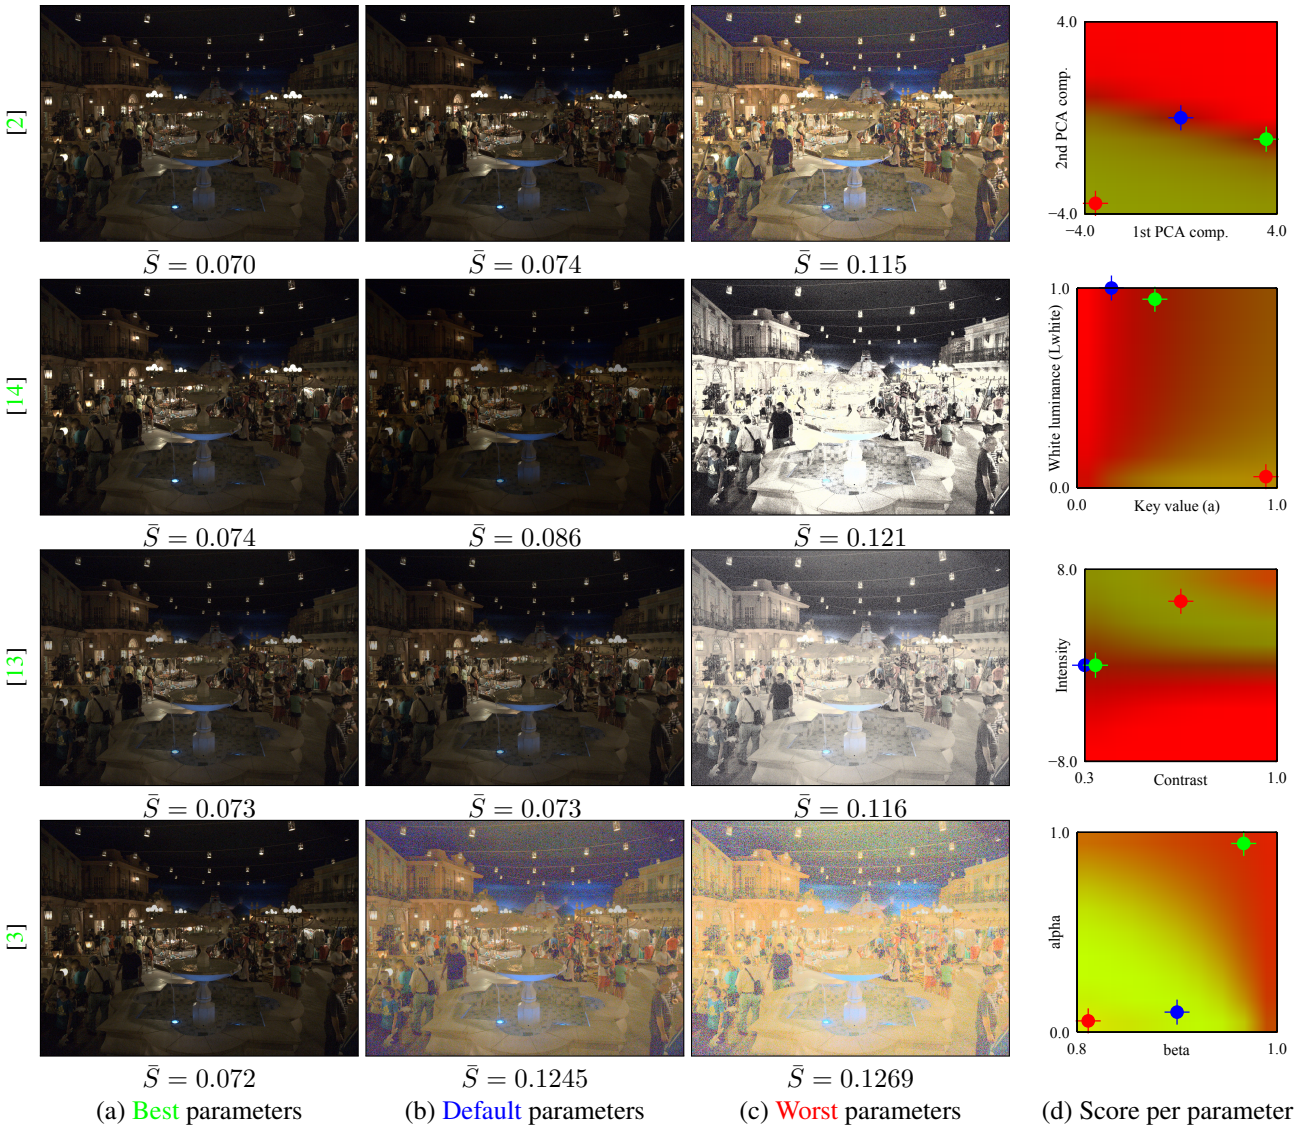

(a) Best parameters (b) Default parameters (c) Worst parameters (d) Score per parameter

Figure 4. TMO parameter optimization: best (a), default (b), and worst (c) parameters for different existing TMOs according to the contrast misuse metric ($\beta = 0.9$). The waste/loss plot (d) illustrates the contrast score for each parameter combination (red: high contrast loss, green: high contrast waste). The best (green), default (blue), and worst (red) parameters are marked. Several of the default parameters of TMOs lie close the optimal but they still can be improved if minimizing the contrast misuse score for the particular scene (best seen by zooming into PDF). **Please see the supplementary material for more examples**.

A more efficient (perhaps parallel) computation of our contrast metrics would be an interesting area of future work.

References

[1] T. O. Aydin, R. Mantiuk, K. Myszkowski, and H.-P. Seidel. Dynamic range independent image quality assessment. *ACM TOG*, 27(3):69:1–69:10, Aug. 2008.

[2] F. Durand and J. Dorsey. Fast bilateral filtering for the display of high-dynamic-range images. *ACM TOG*, 21(3):257–266, July 2002.

[3] R. Fattal, D. Lischinski, and M. Werman. Gradient domain high dynamic range compression. *ACM TOG*, 21(3):249–256, July 2002.

[4] J. A. Ferwerda, S. N. Pattanaik, P. Shirley, and D. P. Greenberg. A model of visual adaptation for realistic image synthesis. In *Proc. SIGGRAPH*, pages 249–258, 1996.

[5] M. Granados, B. Ajdin, M. Wand, C. Theobalt, H.-P. Seidel, and H. P. A. Lensch. Optimal HDR reconstruction with linear digital cameras. In *Proc. CVPR*, 2010.

[6] M. D. Grossberg and S. K. Nayar. Modeling the space of camera response functions. *IEEE Trans. PAMI*, 26(10):1272–1282, 2004.

[7] P. Irawan, J. A. Ferwerda, and S. R. Marschner. Perceptually based tone mapping of high dynamic range image streams. In *Proc. EGSR*, pages 231–242, 2005.

[8] J. Janesick. CCD characterization using the photon transfer technique. In K. Prettyjohns and E. Derenlak, editors, *Proc. Solid State Imaging Arrays*, volume 570, pages 7–19. SPIE, 1985.

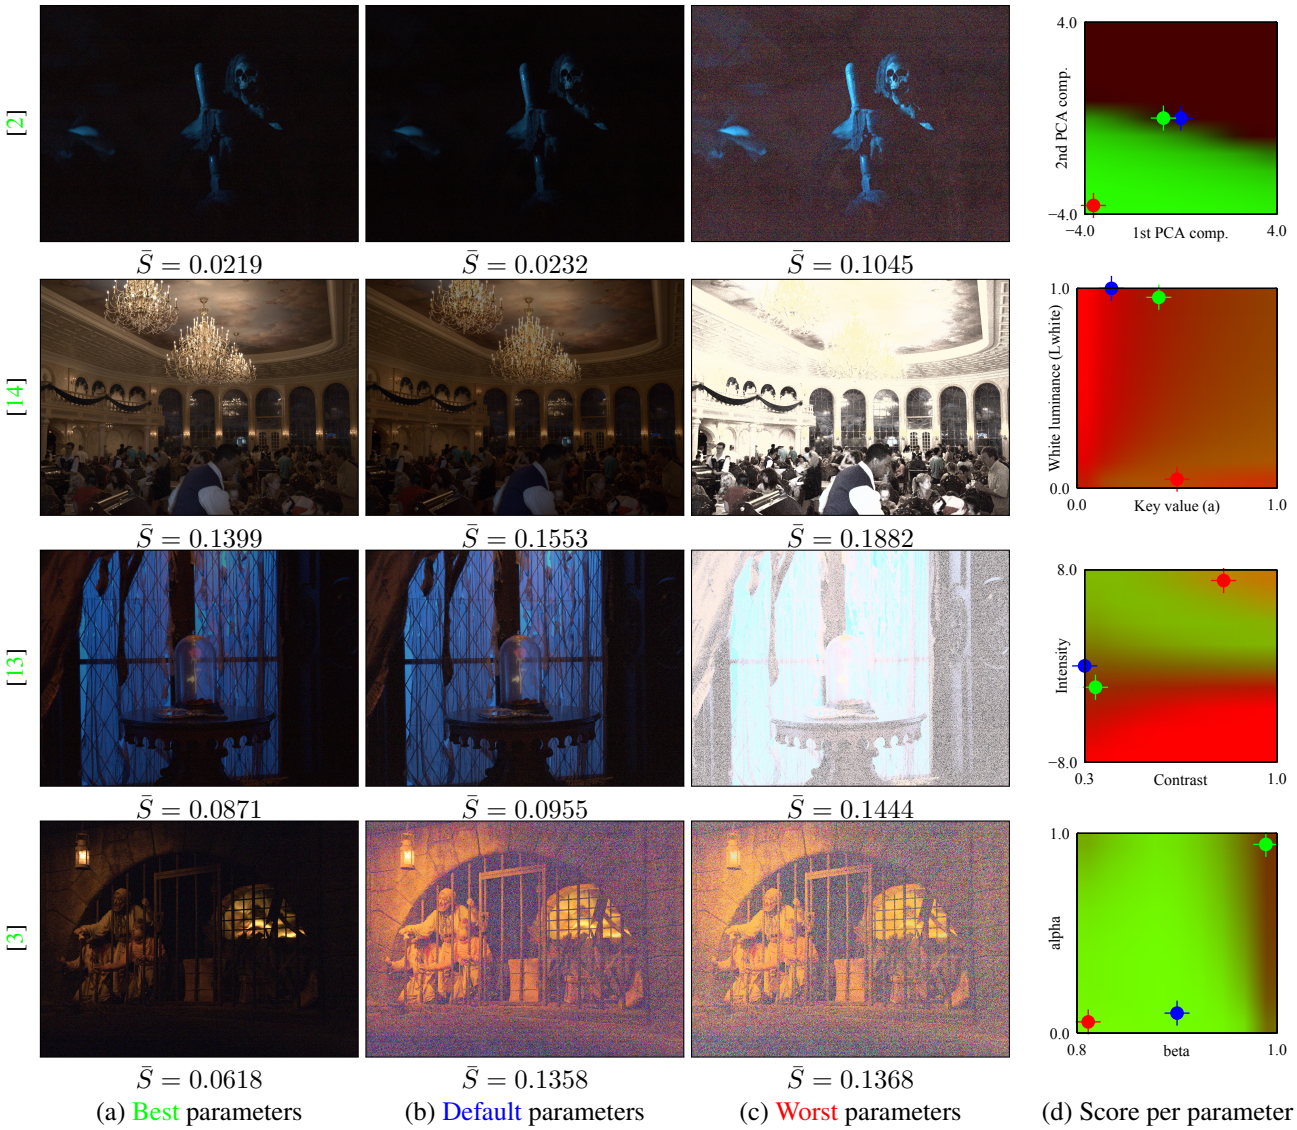

(a) Best parameters (b) Default parameters (c) Worst parameters (d) Score per parameter

Figure 5. TMO parameter optimization: best (a), default (b), and worst (c) parameters for different existing TMOs according to the contrast misuse metric ($\beta = 0.9$). The score plot (d) illustrates the contrast score for each parameter combination (red: high contrast loss, green: high contrast waste). The best (green), default (blue), and worst (red) parameters are marked. Several of the default parameters of TMO lie close the optimal but they still can be improved if minimizing the contrast misuse score for the particular scene (best seen by zooming into PDF). **Please see the supplementary material for more examples**.

[9] J. Janesick. *Scientific charge-coupled devices*. SPIE Press, 2001.

[10] A. G. Kirk and J. F. O'Brien. Perceptually based tone mapping for low-light conditions. *ACM TOG*, 30(4):42:1–42:10, July 2011.

[11] C. Liu, R. Szeliski, S. B. Kang, C. L. Zitnick, and W. T. Freeman. Automatic estimation and removal of noise from a single image. *IEEE PAMI*, 30(2):299–314, 2008.

[12] S. N. Pattanaik, J. Tumblin, H. Yee, and D. P. Greenberg. Time-dependent visual adaptation for fast realistic image display. In *Proc. SIGGRAPH*, pages 47–54, 2000.

[13] E. Reinhard and K. Devlin. Dynamic range reduction inspired by photoreceptor physiology. *IEEE TVCG*, 11(1):13–24, 2005.

[14] E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda. Photographic tone reproduction for digital images. *ACM TOG*, 21(3):267–276, July 2002.

[15] E. Reinhard, G. Ward, S. Pattanaik, P. Debevec, W. Heidrich, and K. Myszkowski. *HDR Imaging - Acquisition, Display, and Image-Based Lighting, Second Edition*. Morgan Kaufmann, 2010.

[16] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.*, 13(4):600–612, 2004.