



Tesis Fin de Máster
Máster en Ingeniería de Sistemas e Informática
Curso 2009-2010

Reverse Tone Mapping for Suboptimal Exposure Conditions

ANEXOS

Belén Masiá Corcoy

Septiembre de 2010

Director: Diego Gutiérrez Pérez

Departamento de Informática e Ingeniería de Sistemas
Centro Politécnico Superior
Universidad de Zaragoza

Contents

1	Prologue	1
2	Introduction	1
2.1	High Dynamic Range Imaging: background	1
2.2	Reverse Tone Mapping: the problem	2
2.3	Approach and contributions of this work	3
3	Previous Work	5
3.1	Reverse tone mapping	5
3.2	User studies	6
4	Experiment One: rTMO Evaluation	7
4.1	Introduction	7
4.2	Stimuli and subjects	7
4.3	Procedure	9
4.4	Results and Discussion	9
5	Experiment Two: HDR vs. LDR Monitor	12
5.1	Introduction	12
5.2	Results and Discussion	12
6	Expanding Over-exposed Content	14
6.1	Introduction	14
6.2	Gamma expansion	14
6.3	Validation	16
7	Selective Reverse Tone Mapping	18
7.1	Introduction	18
7.2	Using the Zone System for rTM	19
7.3	Content-aware rTM	21
7.3.1	Detecting salient features	21
7.3.2	Expanding the dynamic range	22
7.4	Results and Discussion	22
8	Conclusions and future work	24
8.1	Conclusions	24
8.2	Future work	25

References	27
A Complete set of stimuli	31
B Objective validation: complete results	33
C Subjective validation for the dark images series	36
D Saliency Detection	37
D.1 Learning-based saliency detection	37
D.2 Saliency Cuts	38
E Publications	40

Appendix A. Complete set of stimuli

This appendix includes the complete set of stimuli used in our tests for rTMO evaluation. The set consists of 9 scenes, and four exposures per scene. For each of the 36 images here depicted, 3 HDR renditions were obtained using a representative subset of reverse tone mapping operators: Banterle’s operator [13], LDR2HDR [17], and linear contrast scaling [8]. The resulting quadruples were the stimuli presented to the subjects of the experiment.

A. Complete set of stimuli




Figure A.1: Complete set of stimuli used in our experiments.

Appendix B. Objective validation: complete results

Results of the objective image quality metric developed by Aydin et al. [24] comparing the original LDR images (reference images) with the corresponding outputs after reverse tone mapping them with the different operators. Each row shows the results of the linear expansion, LDR2HDR, Banterle’s operator and our gamma expansion, respectively. Green, blue and red identify loss of visible contrast, amplification of invisible contrast and contrast reversal respectively. Our γ expansion does not lose any contrast, while minimizing gradient reversals. Additionally, it reveals more detail in significant areas of the images. The objective evaluation was run only for the scenes of the bright series, for which our γ expansion was devised.

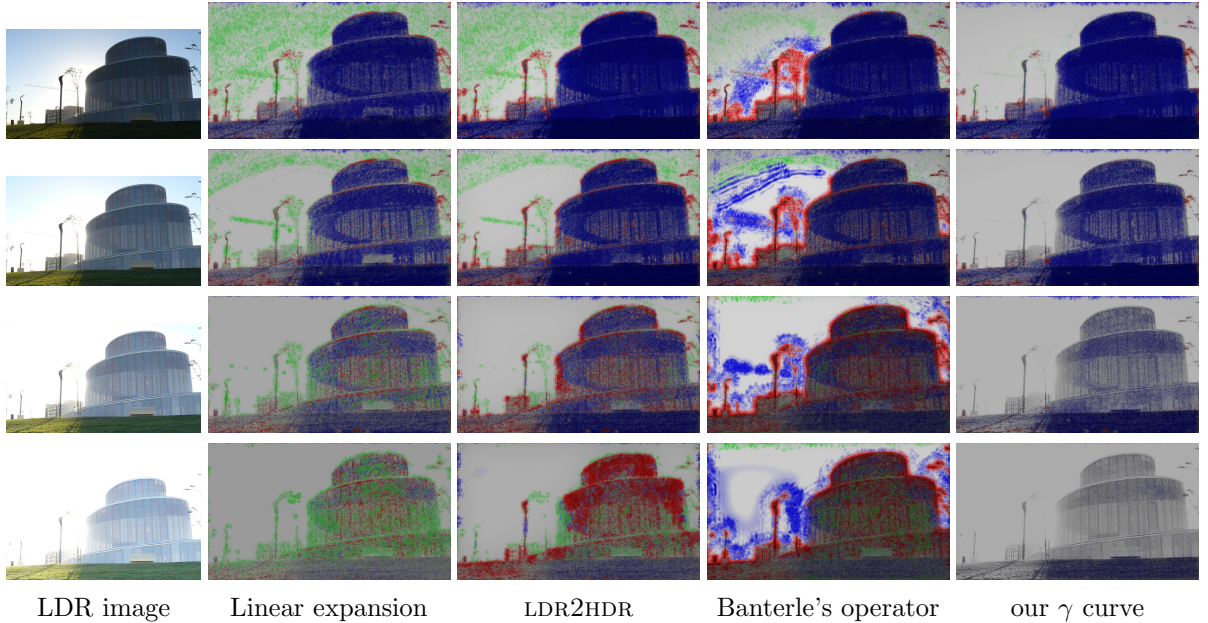


Figure B.1: Objective validation: *Building* scene. Green, blue and red identify loss of visible contrast, amplification of invisible contrast and contrast reversal respectively.

B. Objective validation: complete results




Figure B.2: Objective validation: *Sunset* scene. Green, blue and red identify loss of visible contrast, amplification of invisible contrast and contrast reversal respectively.

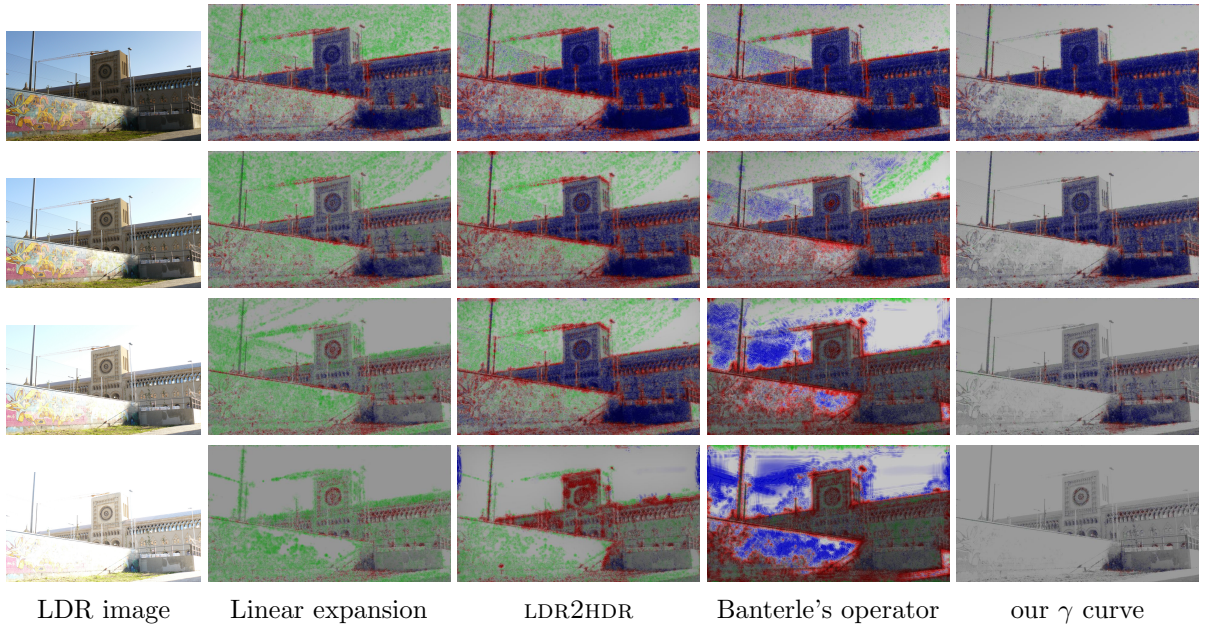


Figure B.3: Objective validation: *Graffiti* scene. Green, blue and red identify loss of visible contrast, amplification of invisible contrast and contrast reversal respectively.

B. Objective validation: complete results




Figure B.4: Objective validation: *Strawberries* scene. Green, blue and red identify loss of visible contrast, amplification of invisible contrast and contrast reversal respectively.

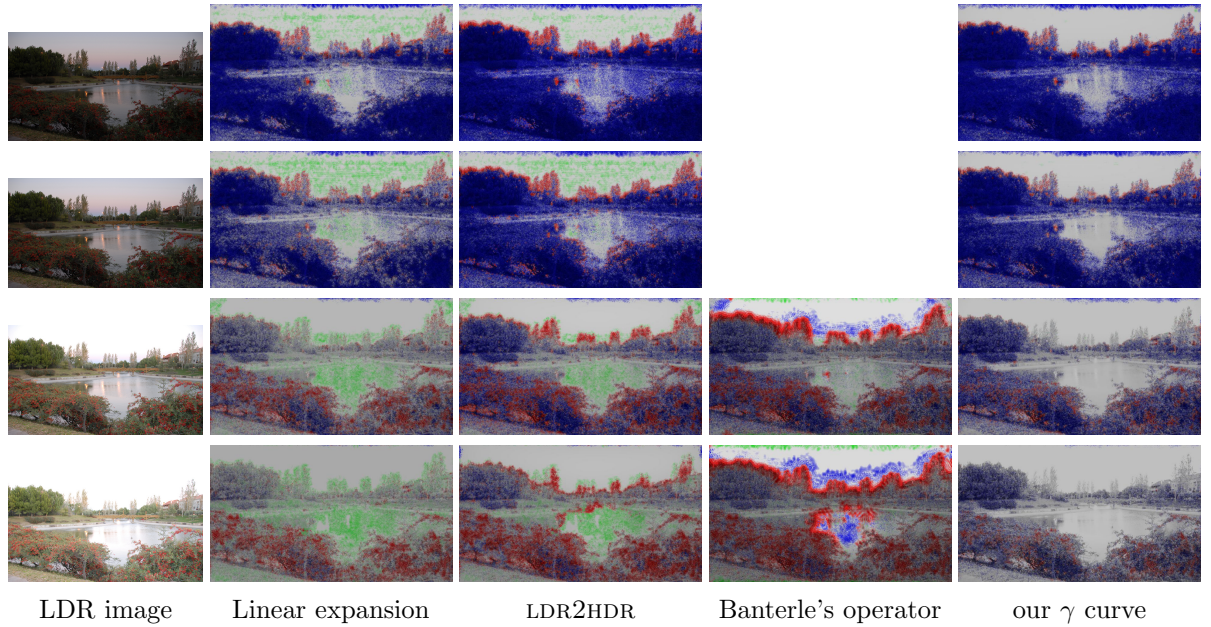


Figure B.5: Objective validation: *Lake* scene. Green, blue and red identify loss of visible contrast, amplification of invisible contrast and contrast reversal respectively. The missing images are due to some error in the application which made it impossible to obtain the corresponding comparison.

Appendix C. Subjective validation for the dark images series

This appendix includes the results obtained in the subjective evaluation of our gamma expansion for the scenes belonging to the dark images series (namely Car, Flowers, Crayons, Pencils). As expected, we found no significant improvements over the tested rTMOs, given that our expansion is designed for bright images.




Figure C.1: Results for the dark images series. The blue bars represent the mean ratings across subjects for the four rTMOs (LDR2HDR, Banterle's, Linear and LDR) with increasing exposure levels (see Chapter 4 for details). Error bars in the last chart represent standard errors on the mean. The red line in the first four charts represents the mean ratings for our proposed γ -curve expansion (see Chapter 6 for details). It can be seen that for dark images our method rates similar to the other four rTMOs. As mentioned in Chapter 8, the absence of significant improvement derives from the fact that our expansion method is designed for bright images.

Appendix D. Saliency Detection

This appendix includes a brief description of the saliency detection methods used in this work. For a more comprehensive explanation we refer the reader to the original papers [38, 39].

D.1 Learning-based saliency detection

This method, introduced by Liu and colleagues [38], delivers, for each input image I , a binary saliency map $A = a_i$, where a_i takes values 1 or 0 depending on whether each pixel belongs or not to the salient object, respectively. In essence, they formulate the problem as a Conditional Random Field (CRF) in which $P(A|I)$ is inferred using a combination of salient features. Learning using a large training database is used to determine the optimal linear combination of the computed salient features.

Given an image I , whose saliency is to be computed, the objective is to obtain a binary saliency mask A . To do this $P(A|I)$ is computed as:

$$P(A|I) = \frac{1}{Z} \exp(-E(A|I)), \quad (\text{D.1})$$

where Z is the partition function (equivalent to a normalizing factor) and the energy $E(A|I)$ is defined as:

$$E(A|I) = \sum_i \sum_{k=1}^K \lambda_k F_k(a_i, I) + \sum_{i,i'} S(a_i, a_{i'}, I). \quad (\text{D.2})$$

The first term of Equation D.2 corresponds to the linear combination of saliency features, so that λ_k are the coefficients which are calculated by learning and $F_k(a_i, I)$ are the K feature functions employed. As for the second term, i and i' denote two adjacent pixels, and $S(a_i, a_{i'}, I)$ is intended so that the pixels in the homogeneous inner part of the salient object are included as salient ones. The function S is thus designed so that the likelihood that two adjacent pixels are assigned different labels decreases the more similar in color the pixels are.

The feature functions F_k follow the expression:

$$F_k(a_i, I) = \begin{cases} f_k(i, I), & a_i = 0 \\ 1 - f_k(i, I), & a_i = 1 \end{cases}, \quad (\text{D.3})$$

with $f_k(i, I)$ being a different function depending on the feature computed but always taking values within the $[0, 1]$ interval. In their work, Liu et al. choose to use three different feature functions at different levels: multi-scale contrast at local level, center-surround histogram at regional level, and color spatial distribution at a global level. The experiments performed in their work show how the combination of the three yields an optimal result. Figure 7.5 (bottom row) shows an example of these feature functions used and the final saliency mask obtained from them.

Training has to be performed for the CRF in order to obtain the coefficients λ_k which determine the influence of each feature function. To do this, the training set is a large database of images (ca. 21,000 images) where the salient object has been manually labeled. The obtention of λ_k is posed as a maximization problem where the objective function is the sum of the log-likelihood (details on how to solve the optimization problem can be found in their paper):

$$\boldsymbol{\lambda} = \begin{pmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \end{pmatrix} = \arg \max_{\boldsymbol{\lambda}} \sum_n \log P(A_j | I_j; \boldsymbol{\lambda}), \quad (\text{D.4})$$

where $I_j, j = 1..N$, refer to the images in the training set and A_j to their corresponding saliency binary masks.

D.2 Saliency Cuts

This method, presented by Fu et al. [39] is essentially a combination of two techniques: the use of graph cuts for object segmentation [40, 41] and the spectral approach to saliency detection of Hou et al. [42].

Interactive graph cuts methods pose object segmentation as a minimal graph cuts problem. The nodes of the graph are formed by image pixels, and the two terminal nodes $\{s, t\}$ correspond to object and background, respectively. These marking of pixels as object or background by the user constitute the hard constraints on the problem, while soft constraints which take into account boundary and region information are also incorporated. The problem of finding minimal cuts in the graph is then solved via a max-flow algorithm [45].

As for the saliency detection, following the spectral residual approach to the problem the saliency map is computed as:

$$S(x) = g(x) * \mathfrak{F}^{-1}[\exp(R(f) + P(f))]^2, \quad (\text{D.5})$$

where $R(f)$ is the spectral residual, obtained as $L(f) - A(f)$, $L(f)$ being the log spectrum of the input image (after downsampling it) and $A(f)$ being the general shape of log spectra. $g(x)$ is a

Gaussian filter used to smooth the final saliency map, \mathfrak{F}^{-1} denotes the Inverse Fourier Transform, and $P(f)$ represents the phase spectrum of the image (the reader can refer to the original paper for a comprehensive description). In the Saliency Cuts implementation this map $S(x)$ is then binarized to obtain an object saliency map $S^o(x)$. This binary saliency map, together with the auto-labeling used for the background and the salient object when performing the segmentation, can be seen in Figure 7.5 (top right).

The idea behind the Saliency Cuts framework is that even though interactive graph cuts can yield very accurate segmentations when proper priors are used, it usually requires a skillful user to select the appropriate regions to feed the algorithm. However, saliency regions detected by the algorithm by Hou and colleagues can serve as seeds to the graph cuts segmentation process, thus obtaining an automatic and accurate separation between the salient object and the background. The limitations of the method lie in the fact that they can only detect a single object and in their assumption that the salient object is at the center of the image, while the sides are always assumed to belong to the background.

Appendix E. Publications

This appendix includes a copy of the publications which have stemmed from the present work.

The first of them, Evaluation of Reverse Tone Mapping Through Varying Exposure Conditions, which covers most of the contents of this master thesis, was presented at SIGGRAPH Asia 2009 in Yokohama, Japan. Articles accepted for this conference are also published in the journal Transactions on Graphics, whose JCR 2009 index is 2/93 in the category of Computer Science, Software Engineering (5-year impact factor 5.012, highest of its category). Transactions on Graphics is regarded as the top journal in the field of Computer Graphics.

The second publication, Selective Reverse Tone Mapping, a smaller project which covers the work described in Chapter 7, has been accepted for publication in Congreso Español de Informática Gráfica 2010.

Evaluation of Reverse Tone Mapping Through Varying Exposure Conditions

Belen Masia¹

Sandra Agustin¹

Roland W. Fleming²

Olga Sorkine³

Diego Gutierrez^{1,4}

¹Universidad de Zaragoza

²Max Planck Institute for Biological Cybernetics

³New York University

⁴Instituto de Investigación en Ingeniería de Aragón (I3A)

Abstract

Most existing image content has low dynamic range (LDR), which necessitates effective methods to display such legacy content on high dynamic range (HDR) devices. Reverse tone mapping operators (rTMOs) aim to take LDR content as input and adjust the contrast intelligently to yield output that recreates the HDR experience. In this paper we show that current rTMO approaches fall short when the input image is not exposed properly. More specifically, we report a series of perceptual experiments using a Brightside HDR display and show that, while existing rTMOs perform well for under-exposed input data, the perceived quality degrades substantially with over-exposure, to the extent that in some cases subjects prefer the LDR originals to images that have been treated with rTMOs. We show that, in these cases, a simple rTMO based on gamma expansion avoids the errors introduced by other methods, and propose a method to automatically set a suitable gamma value for each image, based on the image key and empirical data. We validate the results both by means of perceptual experiments and using a recent image quality metric, and show that this approach enhances visible details without causing artifacts in incorrectly-exposed regions. Additionally, we perform another set of experiments which suggest that spatial artifacts introduced by rTMOs are more disturbing than inaccuracies in the expanded intensities. Together, these findings suggest that when the quality of the input data is unknown, reverse tone mapping should be handled with simple, non-aggressive methods to achieve the desired effect.

CR Categories: I.3.7 [Computing Methodologies]: Computer Graphics—; I.4.10 [Computing Methodologies]: Image Processing and Computer Vision—Image Representation

Keywords: Tone management, high dynamic range imaging, image processing, perception, psychophysics, human visual system

1 Introduction

High dynamic range display devices are becoming increasingly common [Seetzen et al. 2004], yet very large amount of existing low dynamic range legacy content and prevalence of 8-bit photography persist. This presents us with the problem of reverse tone mapping. The aim of reverse tone mapping operators (rTMOs) is to endow low dynamic range (LDR) imagery with the appearance of a higher dynamic range without introducing objectionable artifacts. Ideally, an rTMO should take a standard LDR image as input

and reconstruct as accurately as possible the true luminance values of the original scene. As depicted in Figure 1, this is an ill-posed problem. For most scenes and imaging devices, the image data is irreversibly distorted by unknown nonlinearities, sensor noise, lens flare, blooming, and perhaps most importantly, sensor saturation, which clips high intensities to a constant value. Reverse tone mappers must somehow reconstruct the missing data, or boost the contrast in a way that does not cause the clipped regions to appear visually unpleasant.

Existing rTMOs tackle this ill-posed problem in different ways, leading them to succeed and fail in different conditions. For example, some reverse tone mapping strategies may handle small clipped highlights well, but cause large saturated regions to appear unnatural. Conversely, other rTMOs may avoid introducing artifacts in over-exposed conditions, but fail to enhance under-exposed images sufficiently. The key is to understand which strategies produce the best possible visual experience, for which a number of user studies have recently been conducted [Yoshida et al. 2006; Seetzen et al. 2006; Akyüz et al. 2007; Banterle et al. 2009b]. These experiments have yielded many valuable insights which may guide future rTMO and even HDR display design. However, they have been applied only to subjectively *correctly exposed* images, usually with knowledge of the dynamic range of the original, real-world scene. A key challenge in rTMO design is how to handle non-optimal LDR content, particularly images that are incorrectly exposed.

Our research is dedicated to finding non-intrusive ways to take advantage of the higher dynamic range of the display medium, irrespective of the dynamic range of the original image. Reverse tone mapping also sheds light on a general problem in signal processing: taking partial, distorted or corrupted data and reconstructing the original as faithfully as possible. Here our quality criterion is perceptual faithfulness rather than physical accuracy.

The vast amount of LDR legacy content spans a large range of exposures. Under- or over-exposure may be due to different reasons, including bad choices by the photographer or pure artistic intentions. Legacy professional material may have been shot to make the most appropriate use of the dynamic range available at the time, very different from what is currently available. Additionally, the information about the dynamic range of the real scene is typically not recorded. It is therefore crucial to extend previous studies by taking into consideration varying exposure conditions for a set of images without additional information.

We have performed a series of psychophysical studies assessing how rTMOs handle images across a wide range of exposure levels. We have found that, while existing rTMOs perform sufficiently well for dimmer (under-exposed) images, their performance systematically decreases for brighter (over-exposed) input images. This suggests that there is a need for an rTM method that effectively deals with over-exposed content. We show that simply boosting the dynamic range by means of an adaptive γ curve achieves good results that outperform the current rTMOs, and propose a simple method to obtain a suitable value of γ for each image.

We additionally observe that artifacts produced by some rTMOs are also visible in low dynamic range renditions of the images. This is

because many artifacts are not simply due to inappropriate intensity levels, but also have a spatial component. We perform a second user study to shed light on which type of inaccuracies introduced by reverse tone mapping most hamper our perception of the final image. This information can further help future rTMO design.

2 Previous Work

Reverse tone mapping. Dynamic range expansion, along with related subsequent problems such as contour artifacts, has been initially addressed by bit-depth extension techniques [Daly and Feng 2003] and decontouring methods [Daly and Feng 2004]. However, these techniques are designed for extension to bit-depths much lower than that of HDR displays. More recently, a few works have looked at the problem of reverse tone mapping for the display of LDR images and videos on HDR displays. The general approach of these reverse tone mapping techniques has been to identify the bright areas within the image, and in particular areas that have been clamped due to sensor saturation, such as light sources. Those areas are typically significantly expanded, while the rest is left unchanged or mildly expanded, to prevent noise amplification. We offer here a brief discussion on reverse tone mapping techniques, and refer the reader to the work by Banterle and colleagues [2009a] for a comprehensive review on the topic.

Banterle et al. [2006; 2007] apply the inverse of Reinhard’s tone mapping operator [Reinhard et al. 2002] to the LDR image and detect areas of high luminance in the resultant HDR image. They then produce a so-called *expand-map* by density estimation of the bright areas, and use this map to interpolate between the LDR image and the initial inverse tone mapped HDR image, thus modulating the expansion range. This framework has been extended to video by designing a temporally-coherent version of the expand-map [Banterle et al. 2008]. The LDR2HDR framework of Rempel et al. [2007] is similar in spirit, but their expand-map (which they term *brightness enhancement function*) can be computed in real time using the GPU. The image intensity is first linearized, and a binary mask is computed by thresholding the saturated pixels; the brightness enhancement map is computed as a blurred version of the binary mask, combined with an edge stopping function to retain contrast of prominent edges. The contrast of the LDR image is then scaled according to the enhancement map. Note that the expansion is affected by the size of the bright objects: larger objects may receive more brightness boost. Recently, Kovaleski and Oliveira [2009] presented a reverse tone mapping technique which is also based on real-time computation of a brightness enhancement function, but substitutes a bilateral filter for the combination of a Gaussian blur and an edge stopping function used by Rempel et al. [2007].

Meylan et al. [2006; 2007] explicitly focus on specular highlight detection and apply a steep linear tone mapping curve to the presumably clamped areas, whereas the rest of the image is expanded by a mild linear curve. A more sophisticated segmentation and classification of bright areas in the image is done in the work of Didyk and colleagues [2008]: they segment the bright image areas and label them as diffuse surfaces, light sources, specular highlights and reflections using a trained classifier. Different expansion functions are designed for each class to reproduce the dynamic range more accurately (in particular, the luminance of light sources and highlights is expanded more than that of reflections, while bright diffuse surfaces are not expanded). The method is suitable for high-quality video enhancement thanks to the temporal coherence of the segmentation and the expansion function. Finally, Wang et al. [2007] propose to fill in the texture information of the clamped bright areas by transferring texture from other (well exposed) areas, although the method may not be viable if a suitable region for transferring detail is not found elsewhere. Both methods [Didyk et al. 2008; Wang

et al. 2007] rely on user assistance to guide the process, whereas we are interested in more automatic approaches.

User studies. It is now generally accepted that HDR displays provide a richer visual experience than their LDR counterparts. However, different parameters such as luminance, contrast or spatial resolution influence our visual experience, which makes it difficult to come up with an ideal combination. Additionally, image content probably also affects our preferences. In computer graphics, several researchers have performed a series of user studies, the findings of which may even influence future hardware development.

Yoshida et al. [2006] judged subjective preference (without a reference image) and fidelity (by comparing to a real world scene) for a series of tone mapped images. Users could adjust brightness, contrast and saturation for each individual image. Although their work was geared towards the design of a forward tone mapping operator, their conclusions are also useful for rTMO development: they found that, in general, brighter images were preferred over dimmer ones. Interestingly, however, in certain cases users would break this tendency and keep a significant portion of the image dark, reducing overall brightness and giving more importance to contrast.

Seetzen et al. [2006] analyzed the influence of luminance, contrast and amplitude resolution of HDR displays, to guide future display designs. Their studies show that the preferred luminance and contrast levels are related: for a given contrast, perceived image quality increases with peak luminance, reaches a maximum and then slowly decreases.

Akyüz and colleagues [2007] performed a series of psychophysical studies which revealed that a linear range expansion of the LDR image could surpass the appearance of a true HDR image, suggesting that simple solutions may suffice for reverse tone mapping. Recently, Banterle et al. [2009b] have presented a psychophysical evaluation of existing reverse tone mapping techniques, the results of which indicate that nonlinear contrast enhancement may yield better results overall.

These previous studies provide useful insight into the desirable behavior of tone mapping operators. A key difference with our work is that they were performed on correctly exposed images, whereas we are interested in analyzing reverse tone mapping across varying exposure conditions. In this work, we define over-exposed pixels as those with values ≥ 254 , and under-exposed pixels as those with null values [Rempel et al. 2007; Martin et al. 2008].

3 Experiment One: rTMO Evaluation

To assess the overall performance of an rTMO, it is important to evaluate it across a range of different imaging conditions. To this end, we have performed a user study in which subjects directly compared the output of three reverse tone mapping schemes (plus standard LDR visualization) across a range of exposures, from clearly under-exposed to clearly over-exposed images. We asked subjects to rate the appearance of the reverse tone mapped images on a calibrated Brightside DR37-P monitor (32.26" wide and 18.15" high), with a black level of 0.015 cd/m^2 and a peak luminance of over 3000 cd/m^2 . Calibration of the Brightside monitor was performed to confirm linearity and stable performance during the experiment and to enable comparison to specific intensities in cd/m^2 should the need have arisen in the analysis, as per standard practice in psychophysics. Temperature compensation was turned off to avoid changes in intensity (this was possible thanks to the air conditioning in the room). The LDR versions of the images were displayed by approximately matching the contrast to a typical desktop TFT (Dell).

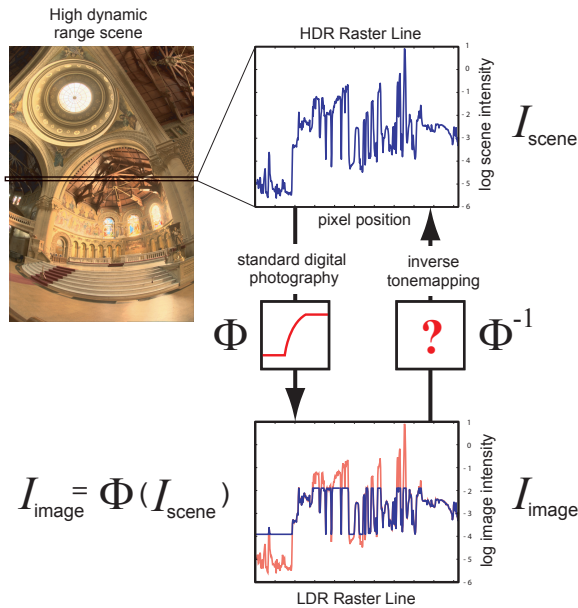


Figure 1: The reverse tone mapping problem. Standard imaging loses data by transforming the raw scene intensities I_{scene} through some unknown function Φ , which clips and distorts the original scene values to create the I_{image} , shown in the bottom panel (values clipped from the original are shown in red). The goal of an rTMO is to invert Φ to reconstruct the original scene data, or to convincingly “fake” it.

Ambient luminance was kept at about 20 cd/m^2 , and the participants were seated approximately one meter away from the monitor. Based on the subjects’ ratings, we can infer which rTMOs are most effective at recreating the experience of an HDR scene without visually objectionable side-effects. As opposed to other studies, we do not provide a ground truth HDR image for direct comparison, since it is almost always unavailable in the case of legacy content.

Stimuli: The stimuli consist of photographs of nine scenes with different lighting conditions, captured with a Nikon D200 at an original resolution of 3872 by 2592 (down-sampled for visualization purposes on the Brightside monitor, which has a 1920 by 1080 pixel resolution). Each scene was captured with four different exposure times. Five scenes were made up of bright images (from approximately correct exposure to clearly over-exposed), and the remaining four were made up of dark images (from clearly under-exposed to approximately correct). Figure 2 shows a representative image of each scene, while Figure 3 shows the four exposures for two example scenes. The stimuli (please refer to the supplementary material for the complete series of all the scenes) have been obtained from a previous study on exposure perception [Martin et al. 2008], where the authors analyze basic image data to try to obtain a correlation between image statistics and the perception of under- and over-exposure.

From each exposure in the bracketed sequence, we obtained three candidate renditions for display on the HDR monitor using a representative subset of reverse tone mapping algorithms: LDR2HDR [Rempel et al. 2007], Banterle’s operator [Banterle et al. 2006] and linear contrast scaling [Akyüz et al. 2007]. Except for the straightforward linear scaling (in Yxy color space, and thus performed on linearized values) we obtained the images from the authors of the original algorithms, in order to ensure accuracy in the implementation. For the LDR2HDR algorithm the parameters used were 150 pixels for the standard deviation of the large Gaussian blur

applied to the mask, a brightness amplification factor $\alpha = 4$ and a gradient image baseline width for divided differences of 5 pixels, plus a 9×9 -pixel kernel for the antialiasing blur and a 4-pixel radius for the *open* operator used to clean up the final edge stopping function (please refer to the original paper for a detailed explanation of these parameters). In the case of Banterle’s operator, when generating the expand-map, the parameters of the density estimation were a radius ranging from 16 to 42 pixels (smaller radius for lower exposures) and a threshold of 1 to 4 light sources (lower threshold for higher exposures), being 2048 the number of generated light sources for Median Cut sampling. In both cases, Banterle’s operator and LDR2HDR, images were linearized using gamma correction ($\gamma = 2.2$). We also added a fourth LDR rendition in which the original images are presented within a luminance range matched to a typical desktop TFT monitor. The goal of this fourth image is to study whether the established assumption that visual preference is given to HDR holds over a range of exposures.

Subjects: A gender-balanced set of twelve subjects with normal or corrected-to-normal acuity and normal color vision were recruited to participate in the experiment. All subjects were unaware of the purpose of the study, and were unfamiliar with HDR imaging.

Procedure: Participants viewed the stimuli on the Brightside HDR display in a dark room. On each trial, subjects were presented with all four renditions of a given exposure of a given scene in a 2×2 array (a *stimulus quadruple*). The positions of the four renditions within the array were random across trials, and the order of the trials was random with the constraint that consecutive trials did not present the same scene. The subjects’ task was to rate the quality of the four renditions on a scale from 1 to 7, according to how accurately the images depicted how the scene would appear to the subject if they were actually present in the scene. Thus the key criterion for comparison was the subjective fidelity of the renditions. Subjects were given unlimited time for each trial and could modify their rating of any of the renditions on a given trial before proceeding to the next trial. Additionally, they were free to assign the same values to all four renditions on a given trial, although they were instructed to try to use as much of the 1-7 scale as possible within the experiment as a whole. To aid them in setting their scale, and to accustom them to the experimental procedure, the subjects were presented with a number of practice trials before the start of the experiment.

Results: Several conclusions can be drawn from this test. First, for our images, there was a clear difference in perceived quality between the bright and the dark series: subjects clearly preferred the reverse tone mapped depictions of darker images over brighter ones. This can be seen in Figure 4: not only is the overall mean value significantly higher in the former case, but it is relatively stable across exposure as well. In contrast, for the bright images, there is a general downward trend in ratings across the four exposure levels.

Note that this gradual decrease in performance does not correlate with the subjective perception of quality of the original LDR image: in a previous pilot study, users picked different exposures for each series as the subjective best, not necessarily the same as the objective best (defined as the one with the smallest proportion of under- and over-exposed pixels [Akyüz et al. 2007]). The trend instead correlates with the proportion of over-exposed pixels and the mean luminance, which do increase with exposure.

Secondly, we can observe *systematic* differences between the rTMOs. On average, subjects rated the LDR2HDR and the Linear rTMOs best (the difference between the two failed to reach sta-



Figure 2: Representative samples of the stimuli used in our tests. Top: bright images (Building, Lake, Graffiti, Strawberries, Sunset), each showing a certain degree of over-exposure. Bottom: dark images (Car, Flowers, Crayons, Pencils), with varying degrees of under-exposure.



Figure 3: The complete bracketed sequence for the Building and Flowers scenes.

(i-j)	$p_b(i, j)$	$p_d(i, j)$
LDR2HDR - Banterle's	2.0532e-21	2.8633e-7
LDR2HDR - Linear	0.5734	0.0283
LDR2HDR - LDR	1.7762e-6	1.4976e-11
Banterle's - Linear	1.1739e-22	0.0013
Banterle's - LDR	4.4489e-11	0.1938
Linear - LDR	1.4697e-7	2.0538e-6

Table 1: Results of the Wilcoxon rank sum tests for the bright and dark series (denoted by subindices b and d respectively). Values of $p < 0.05$ are considered to indicate statistically significant differences between rTMOs. Thus, all differences were significant except for LDR2HDR vs. Linear in the bright series and Banterle vs. LDR in the dark series.

tistical significance), followed by the LDR images, and finally the output of the Banterle's rTMO (see Figure 4). Pairwise Wilcoxon rank sum tests (similar to a non-parametric version of the t-test) reveal that these differences were significant to $p < 0.05$, except for LDR2HDR vs. Linear in the bright series and Banterle's operator vs. the LDR depiction in the dark series (see Table 1 for the complete results).

It is important to note, however, that this ordering does not hold for all conditions. For instance, the LDR depiction was systematically ranked lower than two of the rTMOs, suggesting that indeed HDR visualization is still preferred over LDR, even for under- and over-exposed images. Surprisingly, though, it ranked higher on average than Banterle's rTMO for bright images. The poor overall performance of Banterle's rTMO with this data set is probably due to the fact that it often exaggerates the errors in poorly exposed images, resulting in intrusive artifacts. This becomes clear when we measure the extent to which each rTMO yields outlier rating values for each image. We calculate the median rating for each image across rTMOs. We then obtain the *outlier index* as the difference in rating

for each rTMO relative to this median value. When an rTMO is *neutral*, simply reflecting the overall quality of the exposure of the image, then the outlier index tends to be close to zero. However, when an rTMO stands out relative to the others (for example due to the introduction of artifacts), then the outlier index tends to deviate from zero. In Figure 5, we plot the histogram of the outlier index values for the three rTMOs and the LDR depiction. It is notable that for LDR2HDR, Linear and LDR, the distribution tends to be relatively tightly tuned, while for Banterle's the spread is much broader. This means that on the one hand, when it performs well, it tends to equal or exceed the others. However, it sometimes introduces substantial artifacts that cause the images to look worse than if they were not reverse tone mapped at all.

Although this seems to contradict a recent study where Banterle's operator actually outperformed other rTMOs [Banterle et al. 2009b], it is important to note that the experiments carried out in both cases differ significantly: first of all, in the work by Banterle et al. [2009b] the LDR source images were again well exposed, which is the regime within which Banterle's rTMO performs well, as we also found. However, when the source material is less flattering, we found that the algorithm sometimes produces clearly visible artifacts, which leads to lower ratings. Second, in [Banterle et al. 2009b] the authors used a 2AFC paradigm with direct ground truth comparison, whereas we propose a rating approach, which allows users to report their relative subjective preferences. Both tasks are valid ways of assessing fidelity. However, ours has the advantage that it is closer to the real usage scenario: in general the ground truth is unknown and is not presented for comparison.

4 Experiment Two: HDR vs. LDR Monitor

We notice that artifacts produced by LDR2HDR and Banterle's rTMOs are typically visible in low dynamic range renditions of the images. This is because they generally have a spatial component:




Figure 4: Top: bright images series. The blue bars represent the mean ratings across subjects for the four rTMOs (LDR2HDR, Banterle’s, Linear and LDR) with increasing exposure levels (see Section 3). The last chart clearly shows the downward trend in perceived image quality. Error bars represent standard errors on the mean. The red line in the first four charts represents the mean ratings for our proposed γ -curve expansion (see Section 5). It can be seen that it rates generally higher and is more stable. Bottom: same information for the dark images series, showing higher overall means and a more stable perceived quality across exposures.




Figure 5: Distribution of outlier indices for all four rTMOs. Top: bright series. Bottom: dark series.

they are not simply due to inappropriate intensity levels for certain features, but they also include fringes, visibly boosted noise and other artifacts. To analyze this, we performed a second experiment with seven new subjects, which was identical to the first experiment, except that on each trial, the 2×2 stimulus array was tone mapped using histogram adjustment [Ward et al. 1997]¹. The array was then presented on a standard TFT monitor (note that this means that the LDR control condition now appears much darker than on a normal TFT).

In Figure 6, we plot the average ratings for each image in the LDR control condition against the average ratings in the HDR condition. As can be seen from the scatter plot, the ratings in the LDR control condition correlated extremely strongly with the ratings in the original experiment on the HDR monitor ($r^2 = 0.9018$). We found no significant difference between bright and dark images.

This result does not imply that the images look *the same* in LDR as in HDR: the subjects were not asked to compare these conditions directly, and previous studies have confirmed that HDR depictions are preferred over LDR [Akyüz et al. 2007]. Indeed, none of the




Figure 6: Scatter plot showing a strong correlation between ratings on an HDR monitor and ratings when the images were tone mapped back down to LDR and presented on a standard TFT monitor.

subjects saw both renditions. However, it does demonstrate that the pattern of preferences is extremely well conserved. In other words, the images that were less preferred on the HDR monitor were also less preferred when tone mapped back down to LDR. This has two important implications. First, the strong correlation found suggests that a reasonably predictive evaluation of a rTMO could be made without directly testing on an HDR monitor. Second, as noted, the subjective ratings of HDR images that have been generated from LDR images seem to depend more on the presence or absence of disturbing spatial artifacts than on the exact intensities of different features. A similar observation (confirmed by our test) was made by Aydin et al. [2008]: they noted that the key issue in image reproduction is to accurately maintain the important features while preserving overall structure, whereas achieving an optical match becomes relatively less important. This becomes even more salient given that the dark-adaptation state of the observer is typically unknown, making absolute intensities meaningless to the user.

The design philosophy that emerges from these considerations is

¹We have used the *pcond* program in *Radiance* to tone map the stimuli.

that it is generally better to apply simpler, less-aggressive rTMO schemes if the original image is imperfect. Failing to fully recreate the HDR experience is less disturbing to users than unintended artifacts that can occur when poorly-exposed images are adjusted too aggressively. In the following section we present a simple and robust approach to boosting the dynamic range of over-exposed images, and show that it is less prone to artifacts than other rTMOs.

5 Expanding over-exposed content

Our experiments have shown that the danger with computationally sophisticated reverse tone mapping schemes is the potential to make the image appear worse than before processing, through the introduction of objectionable artifacts. However, the goal of a rTMO is to make the image content look better in general and avoid, under any circumstances, making it look worse. Simple global reverse tone mappers, such as linear scaling and gamma boosting, never cause polarity reversals, ringing artifacts or spuriously boost regions well beyond their context. Our first experiment clearly indicates that there is room for improvement in devising an rTMO for bright input images with large saturated areas, whilst darker images turn out much better. We thus focus on the former in this section.

Examining the bright sequence in Figure 3 we observe that as exposure increases, more detail is lost as pixel values become saturated, and colors fade to white. It thus seems reasonable to attempt to depict the image in a way that the remaining details become more prominent, as opposed to boosting saturated areas as existing rTMOs do. Note that we do not aim to recover information lost to over-exposure, for which existing hallucination techniques may work [Wang et al. 2007], but rather to increase perceived quality.

We make the following key observations, which have been confirmed by previous studies on reverse tone mapping: on the one hand, darker HDR depictions are usually preferred for bright input LDR images [Meylan et al. 2006]; on the other hand, in many cases contrast enhancements improve perceived image quality [Rempel et al. 2007]. These suggest expansion of the linearized luminance values following a simple γ curve, which has the desired effect of darkening the overall appearance of the images while increasing contrast. Linearization of the luminance values prior to the dynamic range expansion was done with a gamma curve ($\gamma = 2.2$), following the findings by Rempel et al. [2007] which note that simple gamma correction can be used for linearization instead of the inverse of the camera response without producing visible artifacts. To avoid amplifying noise, a bilateral filter [Tomasi and Manduchi 1998] can be used prior to expansion [Rempel et al. 2007]. Gamma expansion may potentially boost noise; however, over-exposed images tend to be significantly less noisy than under-exposed ones. Our psychophysical tests confirmed that noise amplification did not affect the final perceived quality.

Obviously, the problem with the proposed expansion lies in automatically obtaining an image-dependent suitable γ value, to avoid the cumbersome manual readjustment of the display settings for each individual image to be shown. For this, we first obtain a mea-

sure of image brightness, for which we compute its key value; this key acts as an indicator of whether the scene is subjectively dark or light. Since overall brightness can be approximated with log-luminance [Tumblin and Rushmeier 1993; Reinhard et al. 2002], we estimate the key of an image as [Akyüz and Reinhard 2006]:

$$k = \frac{\log L_{avg} - \log L_m}{\log L_M - \log L_m} \quad (1)$$

where $\log L_{avg} = (\sum_{x,y} \log(L(x,y) + \delta)) / n$. L_m and L_M are the minimum and maximum image luminances respectively, n is the number of pixels and $L(x,y)$ is the pixel luminance. The small offset δ prevents singularities when $L(x,y) = 0$. We exclude 1% of the highest and lowest pixel values following the suggestion in [Akyüz and Reinhard 2006], to make the estimation less sensitive to outliers. We asked users in a pilot study to manually adjust the value of γ in a set of images, and fitted empirical data with a linear regression $\gamma = a \cdot k + b$ (with $a = 10.44$ and $b = -6.282$), which relates γ as a function of the image key ($r^2 = 0.82$). We have used this expression in this work to compute the reverse tone mapped results in this paper. Table 2 shows the key and γ values used for all the stimuli.

To provide a subjective evaluation of the performance of this strategy, we repeated Experiment One (Section 3), substituting the LDR depiction with our γ -expanded versions in order to maintain the 2×2 stimulus array. The red line in Figure 4 shows the results.

Experiment One provides useful information about the *subjective* perception of image quality. However, we are also interested in evaluating our approach from an *objective* point of view. The problem is the fact that the intended comparison needs to be performed between an LDR and an HDR image. Recently, Aydin and colleagues [2008] have presented a novel image quality metric which identifies visible distortions between two images, independently of their respective dynamic ranges. The metric uses a model of the human visual system, and classifies visible changes between a reference and a test image. The authors identify three types of structural changes: loss of visible contrast (when contrast visible in the reference image becomes invisible in the second one), amplification of invisible contrast (when invisible contrast in the reference image becomes visible in the second one), and reversal of visible contrast (when contrast polarity is reversed in the second image with respect to the reference). It is important to remember that, as Rempel and colleagues noted [2007], contrast enhancement tends to increase perceived quality, and therefore is a desired outcome of the rTMO.

Figure 7 shows the results of this metric² comparing two of the original LDR images (reference images) with the corresponding outputs using linear expansion, LDR2HDR, Banterle’s operator and our proposed γ curve. Our method reveals more detail, shows no loss of contrast and minimizes gradient reversals. Note that while our approach may fail to utilize the dynamic range to its full extent in some cases, it has the important and experimentally validated advantage of avoiding objectionable and unpredictable artifacts.

6 Discussion and Conclusions

Previous works on the perception of HDR images and rTM design have assumed that the input images were, in general, correctly exposed. While these provide valuable knowledge that could guide the development of both HDR display hardware and reverse tone mapping algorithms, existing LDR legacy content actually covers a wide range of exposures, including material that suffers from bad exposure. As currently designed, existing rTMOs tend to boost

	1	2	3	4
Building	0.697 / 1.22	0.762 / 1.5	0.816 / 1.75	0.845 / 2.6
Lake	0.7714 / 1.1	0.7453 / 1.2	0.7487 / 1.5	0.7830 / 2.25
Graffiti	0.7666 / 1.2	0.8193 / 1.35	0.8738 / 1.5	0.9184 / 1.75
Strawberries	0.6696 / 1.22	0.7218 / 1.35	0.7218 / 1.55	0.8479 / 1.9
Sunset	0.7022 / 1.1	0.8103 / 1.35	0.8016 / 1.4	0.8713 / 1.75

Table 2: Key and γ values for the five scenes and the four exposure levels.

²We have used the online implementation provided by the original authors of the paper: <http://drim.mpi-inf.mpg.de/generator.php>

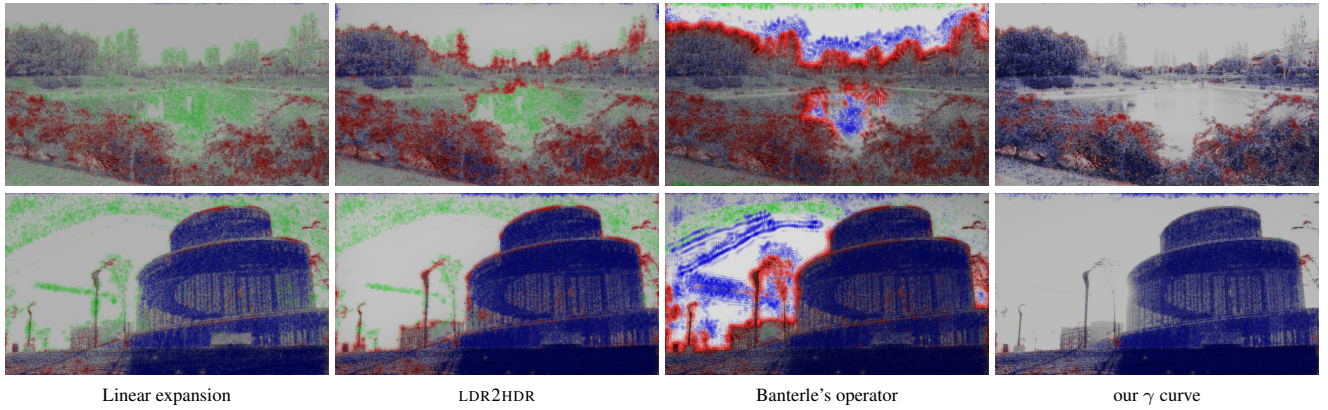


Figure 7: Comparing the results of several rTMOs with the image quality metric from Aydin et al.[2008]. The reference LDR images are Lake (top) and Building (bottom) as depicted in Figure 2 (which correspond to the third and second exposure levels in the series. Please refer to the supplementary material for all the exposures in all the scenes). Green, blue and red identify loss of visible contrast, amplification of invisible contrast and contrast reversal respectively. Our γ expansion does not lose any contrast, while minimizing gradient reversals. More importantly, it reveals more detail in the most significant areas of the images (trees, grass, bushes and buildings in the images shown).

over-exposed areas more than the rest of the image. The strategy works well for small areas such as light sources or highlights if the rest of the image is correctly exposed, but no performance evaluation on generally over-exposed imagery had been performed.

Experiment One shows that performance of rTMOs decreases for input images containing a large number of over-exposed pixels, while they seem to perform significantly better for darker images. This suggests that for bright images the consensual approach of boosting bright areas could be improved. We have shown that a simple rTMO based on γ expansion, without the need for explicitly detecting saturated areas, outperforms existing rTMOs in these cases, and propose an empirical expression to automatically find a suitable γ as a function of the image's key, without user interaction. This rTMO has the desired properties of boosting contrast and detail in non-saturated areas of the image, visually compensating for the lack of information in the saturated ones.

We have performed two validation studies, both subjective and objective. The first one has confirmed that our approach increases the perceived image quality for these kind of images. Pairwise Wilcoxon rank sum tests revealed that the differences in rating were statistically significant with respect to all other rTMOs tested. Given that it produces darker overall images with increased contrast, this result is in accordance with previous suggestions [Meylan et al. 2006; Rempel et al. 2007]. The second evaluation uses a recently published image quality metric which operates with arbitrary dynamic ranges [Aydin et al. 2008]. The metric concludes that our method reveals more detail in non-saturated areas, does not reduce contrast and shows less gradient reversals than the other rTMOs tested. Thus, the artists' original intentions are better preserved.

In both experiments we used typical numbers of subjects for a within-subject design in psychophysics, and the results were highly coherent across subjects. In Experiment One the reported results are statistically significant to the $p < 0.05$ level, meaning that the chances that the outcome of the pairwise comparisons would change after running more subjects from the same population is less than 5%. Indeed, for many of the results, the probability is many orders of magnitude lower than this, which implies that the qualitative pattern of the results is well conserved across subjects. Likewise, data from Experiment Two exhibit a correlation coefficient of 0.9018, notably conclusive in statistical terms.

Our findings seem to indicate that superior rTMOs should take into

account global statistics about the image, and not just individual pixel values. We have derived a simple strategy based on the key value of the images, but more sophisticated strategies could also be devised, possibly including high-level semantics.

We also ran the same expansion on the images from the dark series: as expected, we found no significant improvements over the tested rTMOs, given that our expansion is designed for bright images (please refer to the supplementary material for the complete data).

The results from our second experiment confirm that spatial artifacts are more disturbing than inaccuracy in reproduced intensity levels [Aydin et al. 2008]. We found a very strong correlation in the pattern of preferences when viewing images on HDR and LDR displays. This does not mean that the images looked the same, but it does suggest that the artifacts that emerge with poorly-exposed input images are spatial in nature and severe enough that HDR evaluation is not necessary: they can also be clearly seen in LDR.

We do not aim to create new depictions of LDR material, which would potentially interfere with the original intentions and artistic vision. Our goal is much like that of an audio mastering engineer: we wish to increase the illusion of power, presence and fidelity in the final display medium, while preserving the author's original vision of the content. Our results complement those in the work by Akyüz et al. [2007], where the authors show that, for *correctly exposed* imagery, a simple linear expansion works well and suggest that sophisticated treatment of LDR data may not be necessary. In fact, our work is consistent with that of Akyüz et al. [2007] in the sense that our proposed γ curves approach linear scaling when the image is approximately correctly exposed. Together, both studies suggest that potentially complex operators might not be needed.

The conclusions drawn aim to be valuable for further development of HDR display technology, HDR imaging in general and the development of future LDR expansion algorithms in particular. However, further tests on LDR expansion are desirable. As the community investigates this issue further, this and similar studies will surely be extended and updated. Future reverse tone mapping strategies could allow the user to control dynamic range expansion based on her own preferences or intended goal. Aspects such as tonal balance or mood could potentially be independently controlled by the user, in a similar manner to existing tools for LDR or HDR images [Bae et al. 2006; Lischinski et al. 2006; Farbmán et al.

2008]. Similarly, reverse tone mapping for video content is a key challenge in this field. In order to develop operators that gracefully handle changes in exposure over time, it is crucial to first understand how they fail in the static case, for which we hope this work stimulates future research.

7 Acknowledgements

We would like to express our gratitude to the anonymous reviewers for their valuable comments. We would also like to thank Miguel Martin for contributing the original images, the authors of the other rTMOs, especially Francesco Banterle, Allan Rempel and Wolfgang Heidrich, for running their algorithms on our images, Aydin et al. for making their quality metric publicly available, and the participants in the experiments for their time. This research has been partially funded by the Spanish Ministry of Science and Technology (TIN2007-63025) and the Gobierno de Aragón (project OTRI 2009/0411). Roland Fleming was funded by DFG FL624/1-2, by BW-FIT grant “Information at your Fingertips” and by the Max Planck Society. Olga Sorkine’s research is supported in part by an NYU URCF grant.

References

- AKYÜZ, A. O., AND REINHARD, E. 2006. Color appearance in high dynamic range imaging. *SPIE Journal of Electronic Imaging* 15, 3.
- AKYÜZ, A. O., FLEMING, R., RIECKE, B. E., REINHARD, E., AND BÜLTHOFF, H. H. 2007. Do HDR displays support LDR content?: a psychophysical evaluation. *ACM Trans. Graph.* 26, 3, 38.
- AYDIN, T. O., MANTIUK, R., MYSZKOWSKI, K., AND SEIDEL, H.-P. 2008. Dynamic range independent image quality assessment. *ACM Trans. Graph.* 27, 3, 69.
- BAE, S., PARIS, S., AND DURAND, F. 2006. Two-scale tone management for photographic look. *ACM Trans. Graph.* 25, 3, 637–645.
- BANTERLE, F., LEDDA, P., DEBATTISTA, K., AND CHALMERS, A. 2006. Inverse tone mapping. In *GRAPHITE '06*, ACM, New York, NY, USA, 349–356.
- BANTERLE, F., LEDDA, P., DEBATTISTA, K., CHALMERS, A., AND BLOJ, M. 2007. A framework for inverse tone mapping. *Vis. Comput.* 23, 7, 467–478.
- BANTERLE, F., LEDDA, P., DEBATTISTA, K., AND CHALMERS, A. 2008. Expanding low dynamic range videos for high dynamic range applications. In *Proc. of the Spring Conference on Computer Graphics*, ACM, New York, NY, USA.
- BANTERLE, F., DEBATTISTA, K., ARTUSI, A., PATTANAIK, S., MYSZKOWSKI, K., LEDDA, P., BLOJ, M., AND CHALMERS, A., 2009. High dynamic range imaging and LDR expansion for generating HDR content. *Annex Eurographics 2009*, April.
- BANTERLE, F., LEDDA, P., DEBATTISTA, K., BLOJ, M., ARTUSI, A., AND CHALMERS, A. 2009. A psychophysical evaluation of inverse tone mapping techniques. *Computer Graphics Forum* 28, 1, 13–25.
- DALY, S., AND FENG, X. 2003. Bit-depth extension using spatiotemporal microdither based on models of the equivalent input noise of the visual system. In *Proc. of Color Imaging VIII: Processing, Hardcopy, and Applications*, SPIE, vol. 5008, 455.
- DALY, S., AND FENG, X. 2004. Decontouring: prevention and removal of false contour artifacts. In *Proc. of Human Vision and Electronic Imaging IX*, SPIE, vol. 5292, 130.
- DIDYK, P., MANTIUK, R., HEIN, M., AND SEIDEL, H.-P. 2008. Enhancement of bright video features for HDR displays. *Computer Graphics Forum* 27, 4, 1265–1274.
- FARBMAN, Z., FATTAL, R., LISCHINSKI, D., AND SZELISKI, R. 2008. Edge-preserving decompositions for multi-scale tone and detail manipulation. *ACM Trans. Graph.* 27, 3, 67.
- KOVALESKI, R. P., AND OLIVEIRA, M. M. 2009. High-quality brightness enhancement functions for real-time reverse tone mapping. *The Visual Computer* 25, 5-7 (April), 539–547.
- LISCHINSKI, D., FARBMAN, Z., UYTENDAELE, M., AND SZELISKI, R. 2006. Interactive local adjustment of tonal values. *ACM Trans. Graph.* 25, 3, 646–653.
- MARTIN, M., FLEMING, R., SORKINE, O., AND GUTIERREZ, D. 2008. Understanding exposure for reverse tone mapping. In *Congreso Español de Informática Gráfica*, 189–198.
- MEYLAN, L., DALY, S., AND SÜSSTRUNK, S. 2006. The reproduction of specular highlights on high dynamic range displays. In *IS&T/SID 14th Color Imaging Conference*.
- MEYLAN, L., DALY, S., AND SÜSSTRUNK, S. 2007. Tone mapping for high dynamic range displays. In *Proc. IS&T/SPIE Electronic Imaging: Human Vision and Electronic Imaging XII*, vol. 6492.
- REINHARD, E., STARK, M., SHIRLEY, P., AND FERWERDA, J. 2002. Photographic tone reproduction for digital images. *ACM Trans. Graph.* 21, 3, 267–276.
- REMPEL, A. G., TRENTACOSTE, M., SEETZEN, H., YOUNG, H. D., HEIDRICH, W., WHITEHEAD, L., AND WARD, G. 2007. Ldr2Hdr: on-the-fly reverse tone mapping of legacy video and photographs. *ACM Trans. Graph.* 26, 3, 39.
- SEETZEN, H., HEIDRICH, W., STUERZLINGER, W., WARD, G., WHITEHEAD, L., TRENTACOSTE, M., GHOSH, A., AND VOROZCOVS, A. 2004. High dynamic range display systems. *ACM Trans. Graph.* 23, 3, 760–768.
- SEETZEN, H., LI, H., YE, L., WARD, G., WHITEHEAD, L., AND HEIDRICH, W. 2006. Guidelines for contrast, brightness, and amplitude resolution of displays. In *Society for Information Display (SID) Digest*, 1229–1233.
- TOMASI, C., AND MANDUCHI, R. 1998. Bilateral filtering for gray and color images. In *ICCV*, 839–846.
- TUMBLIN, J., AND RUSHMEIER, H. E. 1993. Tone reproduction for realistic images. *IEEE Computer Graphics and Applications* 13, 6, 42–48.
- WANG, L., WEI, L.-Y., ZHOU, K., GUO, B., AND SHUM, H.-Y. 2007. High dynamic range image hallucination. In *Eurographics Symposium on Rendering*, 321–326.
- WARD, G., RUSHMEIER, H., AND PIATKO, C. 1997. A visibility matching tone reproduction operator for high dynamic range scenes. *IEEE Trans. on Visualization and Computer Graphics* 3, 4, 291–306.
- YOSHIDA, A., MANTIUK, R., MYSZKOWSKI, K., AND SEIDEL, H.-P. 2006. Analysis of reproducing real-world appearance on displays of varying dynamic range. *Computer Graphics Forum* 25, 3, 415–426.

Selective Reverse Tone Mapping

Belen Masia

Universidad de Zaragoza

belenmasia@gmail.com

Roland Fleming

Max Planck Institute for Biological Cybernetics

roland.fleming@tuebingen.mpg.de

Olga Sorkine

New York University

sorkine@cs.nyu.edu

Diego Gutierrez

Universidad de Zaragoza,

Instituto de Investigacion en Ingenieria de Aragon

diegog@unizar.es

Abstract

High dynamic range (HDR) displays are becoming more common, which has given rise to a number of reverse tone mapping techniques. The goal of these techniques is to expand the dynamic range of all the existing low dynamic range content to fit that of these displays. Most of the reverse tone mapping operators, however, fail to offer a good solution in cases where the input images contain large saturated areas. In this paper we present an interactive higher-level approach to reverse tone mapping. Inspired by the Zone System used in photography, it can also be used as an artistic tool where both the tonal balance and the mood of the final depiction can be adjusted by the user.

1 Introduction

The increasing dynamic range of modern monitors enables them to display images with much greater contrast, similar to some real-world scenes. High dynamic range (HDR) content visualized on these monitors therefore recreates the real world more faithfully than low dynamic range (LDR) images viewed on conventional displays. However, there exists a great amount of legacy content which has been recorded and stored in low dynamic range formats, which now needs to be dealt with for a

correct visualization on HDR displays.

Reverse tone mapping (rTM) refers to the process of expanding the range of an LDR input image or video to create an HDR depiction which matches real-world luminance values as faithfully as possible. Obviously, accurate reconstruction of these real-world luminances is an impossible task, since some information is inevitably lost due to data quantification, sensor saturation and nonlinearities in the camera response when capturing and encoding the LDR image. Most existing reverse tone mapping operators (rTMOs) do not aim at recovering that lost information, but they try to convey a convincing HDR depiction instead. A notable exception to this is the HDR hallucination technique [WWZ*07], which presents a user-guided approach to fill in missing information of clamped areas by transferring texture detail from other parts of the image.

To produce a pleasant HDR image from LDR input, existing rTMOs work under the general assumption that highly saturated pixels need to be expanded much more than the rest. As a result, bright image areas representing features like highlights, or the sun in the sky, are largely boosted, thus counter-parting the clamping of information in the LDR image and better representing the real-world experience.

Even though these techniques can produce appealing results for a wide range of LDR con-



Figure 1: Examples of images containing large saturated areas.

tent, there are some cases in which the general approach of boosting bright areas may not be the best way to proceed, as shown recently by Masia and colleagues [MAF*09]. These cases include images -such as those shown in Figure 1- which contain large saturated areas, either because of artistic purposes or due to a bad exposure.

In this paper we show how a tailored approach to dynamic range expansion can be a good alternative in those cases which are unfavorable for existing rTMOs. We present two different techniques, one based in Ansel Adams' Zone System [Ada83], and another based on detection of salient features, which allow the user to control dynamic range expansion based on her own preferences or intended goal. The techniques can also be used in combination with each other. This provides a new method for reverse tone mapping and an artistic tool where tonal balance and mood of the final HDR image can be adjusted by the user (in a similar manner to existing tools for LDR or HDR images [BA83, LLW04, BPD06, LFUS06, FFLS08]).

2 Previous Work

In the last few years, several new reverse tone mapping techniques have been proposed (we refer the reader to the existing literature for a comprehensive review on the topic [BDA*09]). Daly and Feng were among the first to address the problem of dynamic range expansion [DF03, DF04], although the specifics of HDR displays were not taken into account, given that the first prototypes had not been available yet. Banterle and col-

leagues [BLDC06, BLD*07] introduce the term *inverse* tone mapping since their approach consists of literally inverting Reinhard's tone mapping operator [RSSF02]. Subsequent density estimation of the bright areas of the input LDR image yields an *expand-map*, which guides the non-linear range expansion. In almost concurrent work, Meylan et al. [MDS06, MDS07] simplify the problem to a simple highlight detection by luminance thresholding, applying a piece-wise linear reverse tone mapping function that allocates more range to those highlights.

To extend the techniques to video, Rempel et al. [RTS*07] combine a Gaussian-blurred *brightness enhancement function* with an edge-stopping function in real time. A bilateral filter is used instead by Kovaleski and Oliveira [KO09], while Banterle and colleagues add temporal coherence to their original expand map [BLDC08]. Didyk et al. [DMHS08] label bright areas in the image as diffuse surfaces, light sources, specular highlights and reflections using a trained classifier and relying on additional user intervention. A different expansion function is then proposed for each type.

All these techniques share the common insight that bright (or even clamped) areas in the image need to be identified and boosted significantly more than the rest of the pixels. A different approach was proposed by Akyüz and colleagues [AFR*07]: by means of a series of psychophysical studies, they show that a simple linear expansion can yield results that are preferred even over a true HDR image. Recently, Masia et al. [MAF*09] show that although existing rTMOs work well for under-exposed input data, their performance

drops with over-exposed content. They propose a simple gamma expansion for those cases, whose value they derive from the key of the input image.

In contrast with previous works, we propose here a higher-level approach to reverse tone mapping where, in addition to pixel intensity data, saliency information at object level is leveraged as well. Coupled with some user input, this allows us to assign more dynamic range to the predicted regions of interest.

3 Using the Zone System for rTM

The so-called Zone System was introduced by Ansel Adams as a guide to produce good photographs with correct tonal values [Ada83]. Exposure is the main factor which determines the way in which the luminance values of the scene are finally mapped to the limited range of values which can be reproduced by the photograph; choosing the right exposure is therefore one of the most important concerns of a photographer. Common exposure meters are designed to aid in this task by measuring luminance values of the scene (or object of interest) and suggesting the lens aperture and shutter speed values. However, irrespective of the scene -its lighting or content-, the values provided by an exposure meter are always such that the object of interest will appear as middle gray in the final image, which in many cases will not be the adequate election. A simple example which illustrates this problem is that of photographing a black and white checkerboard and a scene which is all black except for a white square: the same exposure settings should be used in both of them, yet the reading of an exposure meter would give very different exposure settings for each one. Ansel Adams' Zone System provides a simple way of, using this middle gray reading of exposure meters, choosing the best exposure settings. This system is not only a tool for photographers still widely in use today [Joh99, Gra97], but also a formalization of sensitometry principles which provides deep insight into how mapping of tonal values works. Reinhard et al. [RSSF02] already rely on it as a basis for



Figure 2: Division of luminance in zones according to Ansel Adams' System.

their well-known tone mapper, and posterior works on interactive tone management have also built on this system [LFUS06]. Following Adams' technique the luminance values in a scene can be divided into ten different luminance zones (0 through IX, see Figure 2) according to the equation given by Koren [Kor]:

$$p = ((\exp(v \sin(\pi \frac{zone-1}{16})) - 1) / (\exp(v) - 1))^\psi, (1)$$

where p represents the zone limits in normalized pixel luminances and ψ is the encoding function responsible for non-linearities in the LDR values (usually the inverse of a γ function). The value $v = 5.25$ is set so that zone V on a properly calibrated monitor appears as middle gray [Ada83], defined as 21% of the maximum screen brightness level (this is similar to 18% reflectance referenced to 90% white, which is pure white on good photographic paper). Equation 1 is designed so that input values of zero and one map to zones 0 and IX respectively, while the sine function is responsible for the compression required at high pixel levels. By working in XYZ or Yxy space, color information remains unchanged [AFR*07]. Once the luminance range




Figure 3: *Left*: Input LDR image. *Right*: The result of luminance decomposition for zone-based reverse tone mapping.

of the LDR input image is divided in zones according to Equation 1 (see Figure 3) the reverse tone mapping process is done by assigning different expansion functions to the different zones. Although in theory these functions




Figure 4: Zone-based reverse tone mapping. *Left*: HDR image obtained by linearly expanding luminance values, and corresponding expansion function. *Top center*: Original LDR image. *Right*: HDR image obtained with a piece-wise linear expansion function based on the Zone System, and corresponding graph showing this expansion function.

could be as complex as desired, we choose to use linear functions for each zone, as they offer a good balance between simplicity and control over the expansion. Thus, the resulting rTM function is piece-wise linear. The darkest and the brightest zones (0 and IX, respectively) of the LDR image are mapped to the lowest and the highest luminance values of the HDR display. A second constraint is that the rTM function must be monotonically increasing, as otherwise gradient reversals may appear that spoil the final depiction. Adjusting the slopes of each of the zones may seem like an involved process; however, in the end it somehow resembles what photographers constantly do, as it translates to assigning ranges of the HDR image luminance to each zone of the LDR input image. Besides, the calculation of the resulting HDR image is almost immediate, thus allowing the user to try different curves before choosing the final one. As an example, Figure 4, *right* shows an HDR image obtained by using a piece-wise linear curve on which only three values were specified: Zone IV being assigned 10% of the HDR image luminance range, Zone VI 40% of that range, and Zone VII 60% of it, which translates to adjusting

three points of the LDR-HDR curve shown. We can also appreciate how this simple tuning of the rTM function yields a more appealing depiction than the linear scaling (shown to be on par in subject preference with the HDR image itself by Akyüz and colleagues [AFR*07]). Additionally, this zone-based expansion can also be used as part of a bigger rTM framework, as examples in Section 5 show.

4 Content-aware rTM

As noted before, the general approach in rTM is to allocate most of the additional dynamic range that an HDR display offers to saturated areas in the scene. However, this may not always be the optimal choice. To our knowledge, none of the previous techniques have taken into consideration the *semantics* of the scene. In an image where a large region of it is saturated, such as the leaf in the snow in Figure 1, treating in a different way the object of interest (in this case the leaf) and the saturated background (the snow) can lead to more visually appealing results than boosting the saturated area while leaving the leaf nearly untouched. The same reasoning applies to the

rest of the images in Figure 1, and in general to images which, either as a result of the artist's choice, or because of wrong exposure, contain large saturated areas. Moreover, when dealing with these type of images, linearly expanding the dynamic range (which in general terms is the other rTM alternative offered by the literature) would result in a significant loss of visible contrast, which is a crucial characteristic of these type of images.

We therefore propose to use a higher-level approach in these cases, taking into account the content of the scene and detecting the object of interest in order to use different reverse tone mapping functions for it and for the background. To separate the region of interest from the background a saliency detector can be used.

4.1 Detecting salient features

Saliency detection techniques pursue the objective of detecting those regions where the viewer's attention concentrates when looking at the image. Even though it is an active field where research continues to offer new and improved methods, a series of detectors exist which are able to offer convincing results in a wide variety of images. In general, saliency detection is performed by developing more or less complex models of the human visual system and using them in combination with image metrics. Most models of attention are based on the fact that at the first stage of visual attention low-level visual features (i.e. edges, intensities, orientations) are extracted. Following this, many existing methods obtain low level features on a first stage, and on a second stage they compute saliency based on these features, as does the well-known work of Itti et al. [IKN98]. However, for many purposes it is necessary to perform a third stage in which object segmentation is applied to extract salient *objects* instead of just a map of salient locations. In our case the need for this third stage in the saliency detection is obvious, as we look for an accurate separation between the object of interest and the background. From within the saliency detection techniques developed in the last years, we found two of them to meet

our needs and applied them to our content-aware reverse tone mapping framework. They are both briefly summarized below.

Learning-based saliency detection.

This method, introduced by Liu and colleagues [LSZ*07], delivers, for each input image I , a binary saliency map $A = a_i$, where a_i takes values 1 or 0 depending on whether each pixel belongs or not to the salient object, respectively. In essence, they formulate the problem as a Conditional Random Field (CRF) in which $P(A|I)$ is inferred using a combination of salient features. Learning using a large training database is used to determine the optimal linear combination of the computed salient features.

Given an image I , whose saliency is to be computed, the objective is to obtain a binary saliency mask A . To do this $P(A|I)$ is computed as:

$$P(A|I) = \frac{1}{Z} \exp(-E(A|I)), \quad (2)$$

where Z is the partition function (equivalent to a normalizing factor) and the energy $E(A|I)$ is defined as:

$$E(A|I) = \sum_i \sum_{k=1}^K \lambda_k F_k(a_i, I) + \sum_{i,i'} S(a_i, a_{i'}, I). \quad (3)$$

The first term of Equation 3 corresponds to the linear combination of saliency features, so that λ_k are the coefficients which are calculated by learning and $F_k(a_i, I)$ are the K feature functions employed. As for the second term, i and i' denote two adjacent pixels, and $S(a_i, a_{i'}, I)$ is intended so that the pixels in the homogeneous inner part of the salient object are included as salient ones. The function S is thus designed so that the likelihood that two adjacent pixels are assigned different labels decreases the more similar in color the pixels are.

The feature functions F_k follow the expression:

$$F_k(a_i, I) = \begin{cases} f_k(i, I) & a_i = 0 \\ 1 - f_k(i, I) & a_i = 1 \end{cases}, \quad (4)$$

with $f_k(i, I)$ being a different function depending on the feature computed but always taking values within the $[0, 1]$ interval. In their work, Liu et al. choose to use three different feature functions at different levels: multi-scale contrast at local level, center-surround histogram at regional level, and color spatial distribution at a global level. The experiments performed in their work show how the combination of the three yields an optimal result. Figure 5 (bottom row) shows an example of these feature functions used and the final saliency mask obtained from them. Training has

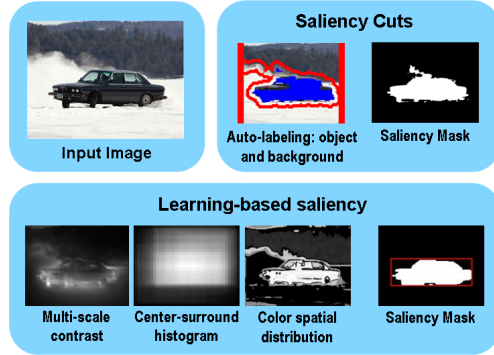


Figure 5: Saliency detection with the different methods. *Top left*: Input image. *Top right*: Saliency detection using the Saliency Cuts algorithm. *Bottom row*: Saliency detection using the learning-based saliency detection approach (images from the saliency database publicly available at <http://research.microsoft.com/~jiansun/>). Further details can be found in the text.

to be performed for the CRF in order to obtain the coefficients λ_k which determine the influence of each feature function. To do this, the training set is a large database of images (ca. 21,000 images) where the salient object has been manually labeled. The obtention of λ_k is posed as a maximization problem where the objective function is the sum of the log-likelihood (details on how to solve the optimization problem can be found in their paper):

$$\lambda = \begin{pmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \end{pmatrix} = \arg \max_{\lambda} \sum_n \log P(A_j | I_j; \lambda), \quad (5)$$

where I_j , $j = 1..N$, refer to the images in the training set and A_j to their corresponding saliency binary masks.

Saliency Cuts. This method, presented by Fu et al. [FCLL08] is essentially a combination of two techniques: the use of graph cuts for object segmentation [BJ01, RKB04] and the spectral approach to saliency detection of Hou et al. [HZ07].

Interactive graph cuts methods pose object segmentation as a minimal graph cuts problem. The nodes of the graph are formed by image pixels, and the two terminal nodes $\{s, t\}$ correspond to object and background, respectively. These marking of pixels as object or background by the user constitute the hard constraints on the problem, while soft constraints which take into account boundary and region information are also incorporated. The problem of finding minimal cuts in the graph is then solved via a max-flow algorithm [BK04].

As for the saliency detection, following the spectral residual approach to the problem the saliency map is computed as:

$$S(x) = g(x) * \mathfrak{F}^{-1}[\exp(R(f) + P(f))]^2, \quad (6)$$

where $R(f)$ is the spectral residual, obtained as $L(f) - A(f)$, $L(f)$ being the log spectrum of the input image (after downsampling it) and $A(f)$ being the general shape of log spectra. $g(x)$ is a Gaussian filter used to smooth the final saliency map, \mathfrak{F}^{-1} denotes the Inverse Fourier Transform, and $P(f)$ represents the phase spectrum of the image (the reader can refer to the original paper for a comprehensive description). In the Saliency Cuts implementation this map $S(x)$ is then binarized to obtain an object saliency map $S^o(x)$. This binary saliency map, together with the auto-labeling used for the background and the salient object when performing the segmentation, can be seen in Figure 5 (top right).

The idea behind the Saliency Cuts framework is that even though interactive graph cuts can yield very accurate segmentations when proper priors are used, it usually requires a skillful user to select the appropriate regions to feed the algorithm. However, saliency re-

gions detected by the algorithm by Hou and colleagues can serve as seeds to the graph cuts segmentation process, thus obtaining an automatic and accurate separation between the salient object and the background. The limitations of the method lie in the fact that they can only detect a single object and in their assumption that the salient object is at the center of the image, while the sides are always assumed to belong to the background.

4.2 Expanding the dynamic range

Once the division in object of interest and background has been performed, different expansion functions can be used for each. These expansion functions can be of any type. Given that we are focusing on an interactive approach where the user guides the reverse tone mapping process, we choose again to use piece-wise linear functions after a separation in luminance zones as explained in Section 3. Resulting HDR images obtained with this rTM framework and the corresponding saliencies and expansion functions are shown in the Results section.

5 Results and Discussion

Figure 6 shows an example of a complete pipeline using our rTM approach, combining the two techniques described in the previous sections. The original image is segmented yielding a binary mask containing the salient object, and a division in luminance zones of the input image is performed. Next, the user can adjust the range of luminance in the HDR final image that will be assigned to each zone, both for the seals and for the background independently. This allows the user to easily manipulate the tonal balance of the image to get the best depiction. In this case a non-linear curve (shown in blue in the graph) has been applied to the seals, thus increasing their contrast and making them more salient; the snow has been just linearly expanded. Segmentation has been performed using Saliency Cuts (seeds used for the foreground and background are shown in blue and red, respectively). Even

though both of the saliency detection methods presented produce segmentations accurate enough for our purposes given input images which are not excessively complex, for increasing complexity (either morphological or related to luminance values) manual segmentation may be necessary. The presence of more than two salient objects in the image also requires a manual segmentation, as the methods discussed cannot segment more than one object. In the results presented in Figure 7 the object of interest was segmented manually and, again, different zone-based piece-wise linear expansion functions were used for the salient object and for the background.

The interactive nature of the approach presented implies that the functions for reverse tone mapping, which determine how the high dynamic range image will look, are adjusted and tuned with low dynamic range renditions of the images as feedback. This is reasonable due to the fact that recent psychophysical experiments have demonstrated that the subjective quality of HDR images that have been generated from LDR images depends more on the presence of absence of spatial artifacts than on the exact luminance values, and thus a reasonably predictive evaluation of an HDR image can be done with an LDR depiction of it [MAF*09].

6 Conclusions and Future Work

In this work we have presented an interactive approach to reverse tone mapping which can be useful for a wide variety of images, especially those containing large saturated areas. The basis of our method is inspired by photographer Ansel Adams' well-known Zone System, which allows us to divide the luminance range of the image into zones. With the aid of this division in zones, and in an interactive process, a piece-wise linear function to expand the LDR image can be provided by the user. Furthermore, our technique includes the possibility of using higher-level information as a guide for the expansion, segmenting the image in the object of interest and the background and using different expansion functions




Figure 6: Complete pipeline using our rTM approach. *From left to right*: Input LDR image, auto-labeling of salient object (blue) and background (red) and binary saliency mask, expansion functions for the salient object (blue) and the background (red), and final HDR image. Original image copyright of National Geographic.




Figure 7: Reverse tone mapping using different zone-based expansion functions for the salient object and the background. *From left to right*: Input LDR image, manually obtained saliency mask, expansion functions for the salient object (blue) and the background (red) and final HDR image. Original image courtesy of Leandro Fessia, all rights reserved.

for each. This interactive approach offers a tool to expand the dynamic range of a scene with significant yet intuitive control over the final result. Besides, being able to freely adjust the luminance ranges of the zones makes it possible to obtain very different HDR depictions of the same input image, potentially providing an artistic tool for photographers and artists in general.

Regarding future work, adding a fitting step of the piece-wise linear rTM functions proposed to smoother ones would be desirable. In the same sense, when dealing with content-aware rTM, taking care of the luminance transitions in the boundary between the objects of interest and the backgrounds would be necessary, either by somehow smoothing the binary mask or by placing constraints to the relationship between both -the object's and the background's- expansion functions. Besides, thorough comparison between the proposed

rTM technique and existing reverse tone mapping operators by means of psychophysical experiments would certainly be interesting for the field. Additionally, salient object detection is an open field of research, and our approach would definitely benefit from future advances in this field. Other lines of future research could involve the design of a contrast-based rTMO, following the findings of the work by Mantiuk et al. [MMS06], which shows promising results in the field of contrast processing of HDR images, working in visual response space.

7 Acknowledgements

The authors would like to thank Sandra Agustin, Miguel Martin and Leandro Fessia. We would also like to thank the authors of the Saliency Cuts paper for running their algorithm on our images, and Tie Liu and colleagues for making their results pub-

licly available. Besides, we would like to thank the Flickr users (Gavatron, revdode, greckor, MikeWebkist) who have shared their photographs. This research has been partially funded by the Spanish Ministry of Science and Technology (TIN2007-63025) and the Gobierno de Aragon projects OTRI 2009/0411 and CTPP05/09. Belen Masia was funded by a grant from the Gobierno de Aragon. Roland Fleming was funded by DFG FL624/1-2, by BW-FIT grant Information at your Fingertips and by the Max Planck Society. Olga Sorkines research is supported in part by an NYU URCF grant.

References

- [Ada83] ADAMS A.: *The Print*. The Ansel Adams Photography series. Little, Brown and Company, 1983.
- [AFR*07] AKYÜZ A. O., FLEMING R., RIECKE B. E., REINHARD E., BÜLTHOFF H. H.: Do HDR displays support LDR content?: a psychophysical evaluation. *ACM Trans. Graph.* 26, 3 (2007), 38.
- [BA83] BURT P. J., ADELSON E. H.: The Laplacian pyramid as a compact image code. *IEEE Transactions on Communications* 31, 4 (1983), 532–540.
- [BDA*09] BANTERLE F., DEBATTISTA K., ARTUSI A., PATTANAIK S., MYSZKOWSKI K., LEDDA P., BLOJ M., CHALMERS A.: High dynamic range imaging and LDR expansion for generating HDR content. Annex Eurographics 2009, April 2009.
- [BJ01] BOYKOV Y. Y., JOLLY M. P.: Interactive graph cuts for optimal boundary and region segmentation of objects in N-D images. In *ICCV* (2001), pp. I: 105–112.
- [BK04] BOYKOV Y. Y., KOLMOGOROV V.: An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26, 9 (2004), 1124–1137.
- [BLD*07] BANTERLE F., LEDDA P., DEBATTISTA K., CHALMERS A., BLOJ M.: A framework for inverse tone mapping. *Vis. Comput.* 23, 7 (2007), 467–478.
- [BLDC06] BANTERLE F., LEDDA P., DEBATTISTA K., CHALMERS A.: Inverse tone mapping. In *GRAPHITE '06* (New York, NY, USA, 2006), ACM, pp. 349–356.
- [BLDC08] BANTERLE F., LEDDA P., DEBATTISTA K., CHALMERS A.: Expanding low dynamic range videos for high dynamic range applications. In *Proceedings of the Spring Conference on Computer Graphics* (New York, NY, USA, 2008), ACM.
- [BPD06] BAE S., PARIS S., DURAND F.: Two-scale tone management for photographic look. *ACM Trans. Graph.* 25, 3 (2006), 637–645.
- [DF03] DALY S., FENG X.: Bit-depth extension using spatiotemporal microdither based on models of the equivalent input noise of the visual system. In *Proceedings of Color Imaging VIII: Processing, Hardcopy, and Applications* (June 2003), vol. 5008, 455, SPIE.
- [DF04] DALY S., FENG X.: Decontouring: prevention and removal of false contour artifacts. In *Proceedings of Human Vision and Electronic Imaging IX* (June 2004), vol. 5292, 130, SPIE.
- [DMHS08] DIDYK P., MANTIUK R., HEIN M., SEIDEL H.-P.: Enhancement of bright video features for HDR displays. *Computer Graphics Forum* 27, 4 (2008), 1265–1274.
- [FCLL08] FU Y., CHENG J., LI Z., LU H.: Saliency cuts: An automatic approach to object segmentation. In *International Conference on Pattern Recognition (ICPR)* (2008).
- [FFLS08] FARBMAN Z., FATTAL R., LISCHINSKI D., SZELISKI R.: Edge-preserving decompositions for multi-scale tone and detail manipulation. *ACM Trans. Graph.* 27, 3 (2008), 67.
- [FLW02] FATTAL R., LISCHINSKI D., WERMAN M.: Gradient domain high dynamic range compression. *ACM Trans. Graph.* 21, 3 (July 2002), 249–256.
- [Gra97] GRAVES C.: *The zone system for 35mm photographers*. Focal Press, 1997.
- [HZ07] HOU X. D., ZHANG L. Q.: Saliency detection: A spectral residual approach. In *Computer Vision and Pattern Recognition* (2007).
- [IKN98] ITTI L., KOCH C., NIEBUR E.: A model of saliency-based visual attention for rapid scene analysis. *IEEE Trans. Pattern Anal. Mach. Intell.* 20, 11 (1998), 1254–1259.
- [Joh99] JOHNSON C.: *The practical zone system*. Focal Press, 1999.
- [KO09] KOVALESKI R. P., OLIVEIRA M. M.: High-quality brightness enhancement functions for real-time reverse tone mapping. *The Visual Computer* 25, 5-7 (April 2009), 539–547.

- [Kor] KOREN N.: A simplified zone system. www.normankoren.com/zonesystem.html.
- [LFUS06] LISCHINSKI D., FARBMAN Z., UYTENDAELE M., SZELISKI R.: Interactive local adjustment of tonal values. *ACM Trans. Graph.* 25, 3 (2006), 646–653.
- [LLW04] LEVIN A., LISCHINSKI D., WEISS Y.: Colorization using optimization. *ACM Trans. Graph.* 23, 3 (Aug. 2004), 689–694.
- [LSZ*07] LIU T., SUN J., ZHENG N.-N., TANG X., SHUM H.-Y.: Learning to detect a salient object. In *IEEE Computer Vision and Pattern Recognition* (2007).
- [MAF*09] MASIA B., AGUSTIN S., FLEMING R. W., SORKINE O., GUTIERREZ D.: Evaluation of reverse tone mapping through varying exposure conditions. *ACM Transactions on Graphics (SIGGRAPH Asia)* 28, 5 (2009), 1–8.
- [MDS06] MEYLAN L., DALY S., SÜSSTRUNK S.: The reproduction of specular highlights on high dynamic range displays. In *IS&T/SID 14th Color Imaging Conference* (2006).
- [MDS07] MEYLAN L., DALY S., SÜSSTRUNK S.: Tone mapping for high dynamic range displays. In *Proc. IS&T/SPIE Electronic Imaging: Human Vision and Electronic Imaging XII* (2007), vol. 6492.
- [MMS06] MANTIUK R., MYSZKOWSKI K., SEIDEL H.-P.: A perceptual framework for contrast processing of high dynamic range images. *ACM Trans. Graph.* 3, 3 (2006), 286–308.
- [RKB04] ROTHER C., KOLMOGOROV V., BLAKE A.: "grabcut": interactive foreground extraction using iterated graph cuts. *ACM Trans. Graph.* 23, 3 (2004), 309–314.
- [RSSF02] REINHARD E., STARK M., SHIRLEY P., FERWERDA J.: Photographic tone reproduction for digital images. *ACM Trans. Graph.* 21, 3 (2002), 267–276.
- [RTS*07] REMPEL A. G., TRENTACOSTE M., SEETZEN H., YOUNG H. D., HEIDRICH W., WHITEHEAD L., WARD G.: Ldr2Hdr: on-the-fly reverse tone mapping of legacy video and photographs. *ACM Trans. Graph.* 26, 3 (2007), 39.
- [WWZ*07] WANG L., WEI L.-Y., ZHOU K., GUO B., SHUM H.-Y.: High dynamic range image hallucination. In *Eurographics Symposium on Rendering* (2007), pp. 321–326.