

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. Once the luminance range of the




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

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 could be as complex as desired, we choose to use linear




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.

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. Expansion is performed on the luminance channel, and the RGB channels are then recovered. Saturation can be tuned when recovering chromaticities in order to obtain the best 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 be-

ing 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 expres-

sion:

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

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 regional 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

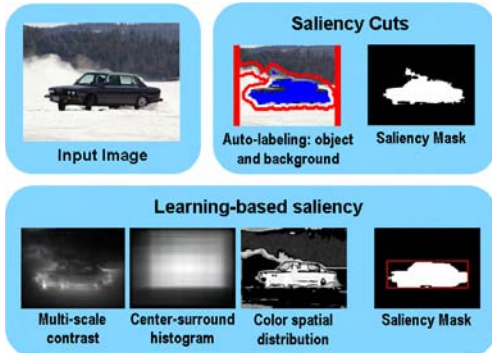


Figure 5: Saliency detection with the different methods. *Top left*: Input image. *Top right*: Saliency detection using the Saliency Cuts algorithm. *Bottom row*: Saliency detection using the learning-based saliency detection approach (images from the saliency database publicly available at <http://research.microsoft.com/~jiansun/>).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

mization 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 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.

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

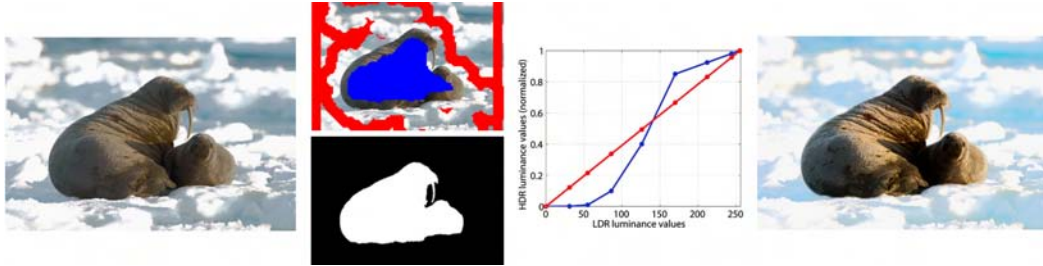


Figure 6: Complete pipeline using our rTM approach. *From left to right:* Input LDR image, auto-labeling of salient object (blue) and background (red) and binary saliency mask, expansion functions for the salient object (blue) and the background (red), and final HDR image. Original image copyright of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

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 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. It would also be interesting to work in Yxy color space instead of RGB to automatically keep ratios between color channels constant. 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. Finally, 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 Fesia. 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 publicly 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., PATTANAİK S., MYZKOWSKI 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.