

# Improved Image Selection for Stack-Based HDR Imaging

Peter van Beek, Cheriton School of Computer Science, University of Waterloo, Canada

## Abstract

*Stack-based high dynamic range (HDR) imaging is a technique for achieving a larger dynamic range in an image by combining several low dynamic range images acquired at different exposures. Minimizing the set of images to combine, while ensuring that the resulting HDR image fully captures the scene's irradiance, is important to avoid long image acquisition and post-processing times. The problem of selecting the set of images has received much attention. However, existing methods either are not fully automatic, can be slow, or can fail to fully capture more challenging scenes. In this paper, we propose a fully automatic method for selecting the set of exposures to acquire that is both fast and more accurate. We show on an extensive set of benchmark scenes that our proposed method leads to improved HDR images as measured against ground truth using the mean squared error, a pixel-based metric, and a visible difference predictor and a quality score, both perception-based metrics.*

## Introduction

The sensor on a digital camera often cannot capture the full dynamic range of a natural scene, resulting in either dark, noisy regions or saturated regions in the image. Stack-based high dynamic range (HDR) imaging is a technique for achieving a larger dynamic range in an image by combining several low dynamic range (LDR) images acquired at different exposures [1, 2]. The resulting HDR image can then be displayed on an HDR monitor, tonemapped for display on a standard LDR monitor, used in computer vision algorithms for object recognition, or used as a source of illumination for computer-generated objects and scenes [3, 4].

An essential part of a stack-based approach is selecting the set of images to be combined. The set must be small, in order to decrease capture and post-processing times, but must accurately and fully capture the scene's irradiance. The problem has received much attention in the literature [5–15]. A low-cost method to select the set of images that is used in many cameras is to bracket around the best exposure at a fixed progression. However, the method is not fully automatic as the number of exposures  $N$  to acquire must be specified in advance, requiring information that is often not available to the photographer. This shortcoming is shared by several more elaborate methods that have been proposed (e.g., [5, 6, 11]). As well, restricting  $N$  to be some small, fixed value (e.g.,  $N \leq 3$ , [8, 14, 15]) will be insufficient in more challenging HDR scenes, and setting  $N$  to be a larger value would often acquire unnecessary images.

To accurately and fully capture a scene's irradiance, ideally a method must be adaptive to the irradiance distribution in a scene [9, 12, 15], and incorporate a model of camera noise [9, 10]. Barakat et al. [7] and Hasinoff et al. [10] propose methods that use an estimate of the extent of the dynamic range, and incorporate a noise model that ensures each pixel is properly exposed in each LDR image [7] or in the final HDR image [10]. Granados

et al. [9], Gallo et al. [12], and Seshadrinathan et al. [13] propose methods that use an estimate of the full irradiance distribution, the HDR histogram, in their selection of images and incorporate a noise model that ensures each pixel is properly exposed in the final HDR image. Pourreza-Shahri and Kehtarnavaz [15] propose a method for selecting images that is adaptive to the irradiance distribution in a scene by clustering the pixel values in a single well-exposed image.

As well, to be suitable in practice, a method must have low computational complexity. Unfortunately, if we consider just the state-of-the-art methods that are adaptive to the irradiance distribution in the scene, only Barakat et al. [7] and Pourreza-Shahri and Kehtarnavaz [15] propose methods whose computational complexity is polynomial. The remaining methods propose solving an optimization problem that is either explicitly NP-complete (e.g., formulated as an integer linear program [10]) or whose computational complexity is unknown but empirically known to be high [9, 12, 13].

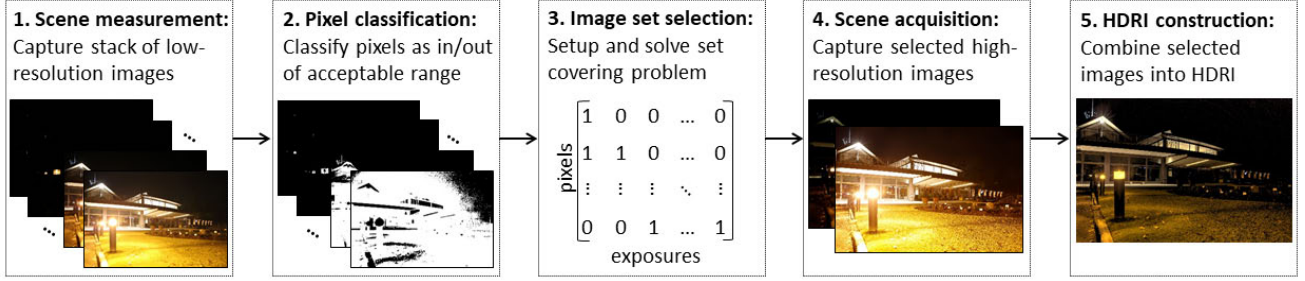
In this paper, we propose a method for selecting the set of exposures to acquire that is both fast and more accurate, particularly on challenging scenes. Following previous work, our proposed method is also adaptive to the irradiance distribution and incorporates a model of camera noise. However, our method formulates the selection problem as a *polynomially* solvable set covering problem. We show on a total of 110 benchmark scenes that overall, in addition to being fast, our proposed method leads to improved HDR images over the state-of-the-art methods as measured against ground truth using the mean squared error, a pixel-based metric, and a visible difference predictor and a quality score, both perception-based metrics. Our experimental evaluation is also the first to extensively evaluate existing state-of-the-art methods for image selection for stack-based HDR imaging.

## Background

In this section, we review some of the underlying concepts and routines used in methods for selecting the set of images.

Many cameras are able to record images in a proprietary RAW format in addition to the common JPEG format. The RAW images are linearly related to scene radiance [16, 17] and have a higher dynamic range (usually, 12–14 bits), while JPEG images are nonlinearly related to scene radiance and have a lower dynamic range (8 bits).

A camera radiometric response function  $f$  is the nonlinear mapping that determines how radiance in the scene becomes pixel values in a JPEG image through the imaging pipeline of the camera. In our experiments, a response function was estimated using multiple pairs of RAW and JPEG images, all taken of the same scene but each pair taken with a different exposure [2, 18–20]. Alternatively, a response function can be roughly modeled using a standard gamma correction curve. Once the mapping is known, JPEG pixel values can be inverted back to an estimate of the RAW




**Figure 1.** Pipeline for proposed end-to-end system for image set selection for construction of HDR image of a scene.

pixel value using the inverse function  $f^{-1}$ . Figure 2(a) shows an example radiometric response function for a Canon EOS 5D Mark III camera.

A camera noise level function  $\text{SNR}(\text{dB})$  maps a RAW pixel value to an estimate of the signal to noise ratio [21]. In our experiments, we followed Healey and Kondepudy’s [18] procedure for estimating camera noise from multiple images, although methods have also been proposed that estimate noise from a single image (see, e.g., [21]). The signal to noise ratio expressed in decibels is given by,  $\text{SNR}(\text{dB}) = 20 \cdot \log_{10}(\mu/\sigma)$ , where  $\mu$  is the RAW pixel value and  $\sigma$  is the estimated noise at that RAW pixel value. SNR increases monotonically (up to sensor saturation) with increased exposure time and decreases monotonically with increased ISO gain. The SNR at saturation is zero by definition. Figure 2(b) shows example SNR(dB) curves for a Canon EOS 5D Mark III camera. The empirical data can be fit with negligible residual error using the parametric noise model  $\sigma = \sqrt{\mu g + r^2 g^2 + c^2}$ , where  $g$  is the ISO gain,  $r$  is the read noise, and  $c$  is the noise component that does not depend on the signal or the gain (see, e.g., [10, 18]). For the Canon EOS 5D Mark III camera,  $c$  is zero in the best fit.

The best methods for selecting the set of images rely on either being given an estimate of the extent of the dynamic range of a scene, or on being given an estimate of the full irradiance distribution in the form of an HDR histogram. In our experiments, we followed Gallo et al.’s [12] procedure for estimating the HDR histogram by first capturing multiple low-resolution JPEG images from the live preview stream by sweeping the lens through the available shutter speeds and then combining the resulting LDR histograms into a HDR histogram in the cumulative distribution function domain.



**Figure 2.** For a Canon EOS 5D Mark III camera, (a) radiometric response function, (b) signal to noise ratio expressed in decibels as a function of  $\mu$ , the RAW pixel value.

## Selection as Set Covering

In this section, we present our method for selecting the set of images. Our approach is inspired by the focus stacking system proposed by Vaquero et al. [22]. For simplicity of presentation and because of their advantages, we assume that RAW images are to be selected and combined into an HDR image, although our approach can also be adapted to the selection of JPEG images. Our method proceeds as follows (see Figure 1).

*Step 1.* Capture a stack of low-resolution JPEG images from the live preview stream by sweeping the lens through the available shutter speeds. Note that these images can be acquired for “free” if one is already estimating the extent of the dynamic range or the HDR histogram (see last paragraph, “Background”).

*Step 2.* Classify each pixel in each low resolution JPEG image as to whether it is accurately captured. A pixel is accurately captured by an exposure if its grayscale conversion falls within the interval  $[I_{\min}, I_{\max}]$ , where  $I_{\min}$  is the darkest pixel  $p$  such that  $\text{SNR}(\text{dB})$  at  $f^{-1}(p)$  is above a given threshold and  $I_{\max}$  is the brightest pixel value such that two or more component channels are rarely saturated.

*Step 3.* The key step is that we formulate the selection of images for HDR imaging as a set covering problem.

**Definition 1** (Set covering). Let  $A = [a_{ij}]$  be an  $m \times n$  (0-1)-matrix with a cost  $w_j$  associated with each column. A row  $i$  of  $A$  is covered by a column  $j$  if  $a_{ij}$  is equal to one. The set covering problem is to find a subset of the columns  $C \subseteq \{1, \dots, n\}$  that minimizes the total cost  $\sum_{j \in C} w_j$  such that every row is covered; i.e., for every  $i \in \{1, \dots, m\}$  there exists a  $j \in C$  such that  $a_{ij} = 1$ .

Let  $\{t_1, \dots, t_n\}$  be the ordered set of available shutter speeds on a camera. In the set covering instance, a row represents a pixel across the stack of low-resolution JPEG images (Step 1), a column represents a possible exposure setting  $t_j$  for that pixel, and an entry  $a_{ij}$  is 1 if and only if the pixel has been accurately captured by that exposure (Step 2). If the goal is to minimize the number of images selected,  $w_j = 1$ , and if the goal is to minimize the total capture time,  $w_j = t_j + t_{\text{over}}$ , where  $t_{\text{over}}$  is the overhead between image acquisitions. In general, solving set covering is NP-complete [23]. However, here the set covering instance can be constructed to have the consecutive ones property, allowing the selection of images to be computed in polynomial time [24, 25].

**Definition 2** (Consecutive ones property). A set covering problem is said to have the consecutive ones property if the ones in each row of the matrix  $A$  appear consecutively.

If all costs  $w_j$  are one, the following simple reduction rules alone solve an instance. Let  $M_j = \{i : a_{ij} = 1\}$  and  $N_i = \{j : a_{ij} = 1\}$ , where  $i \in \{1, \dots, m\}$  and  $j \in \{1, \dots, n\}$ .

- R1. If  $N_{i_1} \subseteq N_{i_2}$ , row  $i_2$  can be removed.
- R2. If  $M_{j_1} \subseteq M_{j_2}$  and  $w_{j_1} \geq w_{j_2}$ , column  $j_1$  can be removed.

If the costs are the shutter speeds, a polynomial algorithm can be used to find a solution after reduction [24, 25].

Steps 4 & 5. As the final steps in our proposed method (see Figure 1), acquire high-resolution RAW images at the exposures specified by the solution to the set covering instance, and combine them into a single HDR image using existing techniques (e.g., [2, 26, 27]).

In the form stated above, our proposed method chooses an exposure by setting the shutter speed, keeping the aperture and ISO gain fixed. Keeping the aperture fixed at the camera’s native ISO is desirable in stack-based imaging to reduce image noise and increase dynamic range (for the Canon EOS 5D Mark III camera, the native ISO is 100; see Figure 2(b) for the effect of ISO on image noise and dynamic range). However, increasing the ISO gain can be useful when the camera is not mounted on a tripod and a minimum shutter speed may be required to reduce the impact of camera shake, or to handle dynamic scenes [8, 10, 13]. Our proposed set covering method can seamlessly handle a fixed but higher ISO gain setting to ensure a minimum shutter speed: the low resolution stack will be acquired at the higher ISO (Step 1), the set covering will take into account the higher ISO when setting the interval  $[I_{min}, I_{max}]$  for determining whether a pixel has been accurately captured (Step 2 & 3), and the high resolution images will also be acquired at the higher ISO (Step 4).

## Experimental Evaluation

In this section, we experimentally evaluate the effectiveness of our proposed method for image selection.

We compare our proposed set covering method against four representative state-of-the-art methods: (i) Barakat et al. [7], (ii) Hasinoff et al. [10], (iii) Pourreza-Shahri and Kehtarnavaz [15], and (iv) Seshadrinathan et al. [13]<sup>1</sup>. All five methods rely, to varying degrees, on knowing the camera response function, the noise level function, and the HDR histogram of the scene. The same routines were used across all methods (see “Background”). As well, when needed a threshold of 20 dB was used for acceptable SNR, and a range of [20, 230] was used for determining whether a pixel has been accurately captured. The grayscale pixel value of 20 corresponds to an SNR of 20 dB for our Canon EOS 5D Mark III camera and the value of 230 was empirically determined to be the grayscale threshold where two or more component channels are rarely saturated. Pourreza-Shahri and Kehtarnavaz’s [15] method requires the specification of a parameter  $w$ , where  $w$  is used in clustering the dark and bright regions of the well-exposed image. We set  $w = 8$ , as in Pourreza-Shahri and Kehtarnavaz’s [15] experiments. Similarly, Seshadrinathan et al.’s [13] method requires the specification of  $N$ , where  $N$  is an upper bound on how many exposures to consider. Because of efficiency considerations, Seshadrinathan et al.’s [13] set  $N$  to be

three in their experiments. In our experiments we set  $N$  to be five; larger values are impracticable.

We test the methods on the following benchmarks. In each benchmark, only the shutter speed was varied and the aperture and ISO gain were kept fixed.

- The HDR Photographic Survey [29] suite consists of 105 benchmark image sets<sup>2</sup>. Each image set consists of nine images (with one exception) of a scene captured with a Nikon D2X camera. In each scene, the exposure step was set to one stop and a  $4288 \times 2848$  high resolution RAW image was acquired at each of the nine shutter speeds using the camera’s continuous auto-bracketing function. We converted each RAW image to JPEG using Nikon’s ViewNX2 software and downsampled to give a  $964 \times 640$  low resolution JPEG image that simulates the image that would have been acquired from the live preview stream at that shutter speed.
- We acquired five benchmark image sets using a camera remote control application we implemented. A Canon EOS 5D Mark III camera mounted on a tripod was tethered to a computer via a USB cable and controlled by software that makes use of the Canon SDK (Version 2.11). In each scene, the exposure step was set to  $1/3$  of a stop and a  $5760 \times 3840$  high resolution RAW image was acquired at each of the 55 possible shutter speeds. A  $960 \times 640$  low resolution JPEG image was also acquired from the live preview stream at each shutter speed.

A ground truth HDR image was constructed using all of the RAW high-resolution images in a benchmark (except if the image would only add noise such as being fully saturated) and compared against the HDR image constructed using only the RAW images selected by each method. HDR images were constructed using Photomatrix Pro<sup>3</sup>. We compare a method’s HDR image, referred to as the *test* HDR image, against a *ground truth* HDR image using the following performance measures.

- *Quality correlate*. The HDR-VDP-2.2.1 image quality metric, which quantifies the visual distortion of the test HDR image from the ground truth HDR image with a single quality score [30, 31].
- *Visual difference prediction*. The HDR-VDP-2.2.1 visual difference prediction metric, which estimates the probability at each pixel that an average human observer will detect a visually significant difference between the test HDR image and the ground truth HDR image [30, 31].
- *Mean squared error*. The mean squared error between the test HDR image and the ground truth HDR image.

Our experimental evaluation is the first to extensively evaluate existing state-of-the-art methods for image selection for stack-based HDR imaging, with significantly more methods, benchmarks, and performance measures being used compared to the limited numbers used in previous evaluations.

Figure 3(a) summarizes the *quality* of the test HDR images against the ground truth HDR images, as measured by the quality correlate, for the 105 HDR Photographic Survey benchmarks.

<sup>1</sup>The MATLAB implementations and benchmark images are available at: <https://anonymous>.

<sup>2</sup><http://www.rit-mcsl.org/fairchild/HDR.html>

<sup>3</sup><https://www.hdrsoft.com/>

Figure 3(b)&(c) summarizes the *errors* for the 105 HDR Photographic Survey benchmarks. To measure errors, for each benchmark we summarize the probability at each pixel that an average human observer will detect a visually significant difference from ground truth with the percentage of pixels that are greater than or equal to 0.75; i.e., the percentage of pixels where a difference is very likely to be detected (see Figure 3(a)). As well, to measure errors, for each benchmark we summarize the mean squared error between the test HDR images and the ground truth HDR images (see Figure 3(b)). It can be seen that on these benchmarks, even with the limited range of choice in each benchmark—only nine images are available for selection, each a full stop apart—our set covering method achieves improvements across all three performance measures.

Figure 4 summarizes the results for the five Canon benchmarks<sup>4</sup>. The (non-tonemapped) images encode the probability  $p$  that an average human observer would detect a significant difference from ground truth: blue,  $p = 0.0$ ; cyan,  $p = 0.25$ ; green,  $p = 0.50$ ; yellow,  $p = 0.75$ ; and red,  $p = 1.0$ . It can be seen that on these benchmarks, with a more extensive range of choice in each benchmark—55 images are available for selection, each 1/3 of a stop apart—our set covering method consistently achieves excellent results as measured by all three performance metrics. On these benchmarks, the existing state-of-the-art methods fail to fully capture the more challenging scenes, whereas our set covering method successfully captures the scenes.

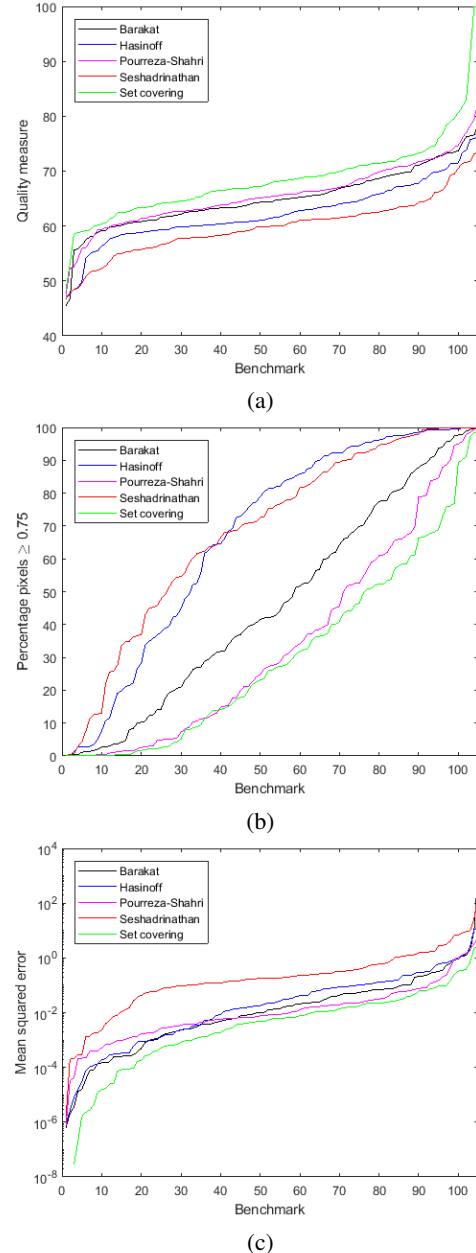
The methods can also be compared using speed as the performance measure. The methods can be clustered into fast (Barakat et al. [7], Pourreza-Shahri and Kehtarnavaz [15], and our set covering), and slow (Hasinoff et al. [10], Seshadrinathan et al. [13]). A MATLAB implementation of our set covering method took approximately 2/3 sec. for pixel classification for 55 images (Step 2) and 1/3 sec. for image set selection (Step 3). In comparison, Hasinoff et al.’s [10] method took an average of 82.4 sec. The key reason for the speed and computational complexity differences between the fast and slow methods is that the fast methods place a threshold on acceptable SNR for each LDR image whereas the slow methods place the threshold on the final HDR image. The former allows fast greedy methods whereas the latter raises the computational complexity significantly (either known to be NP-complete or empirically shown to be high).

Finally, the methods can also be compared using the median and 75th percentile number of images selected across all 110 benchmarks: Barakat et al. [7], 4 and 4 images; Hasinoff et al. [10], 3 and 3 images; Pourreza-Shahri and Kehtarnavaz [15], 3 and 4 images; Seshadrinathan et al. [13], 1 and 2 images; and our proposed set covering method, 3 and 4 images, respectively. Thus, the improvement in accuracy of our proposed method is not at the expense of increasing the capture time for most scenes.

## Conclusion

We proposed a method for selecting the set of images to combine in stack-based HDR imaging that is fast and offers improved accuracy. Our technique minimizes the set of images to combine,

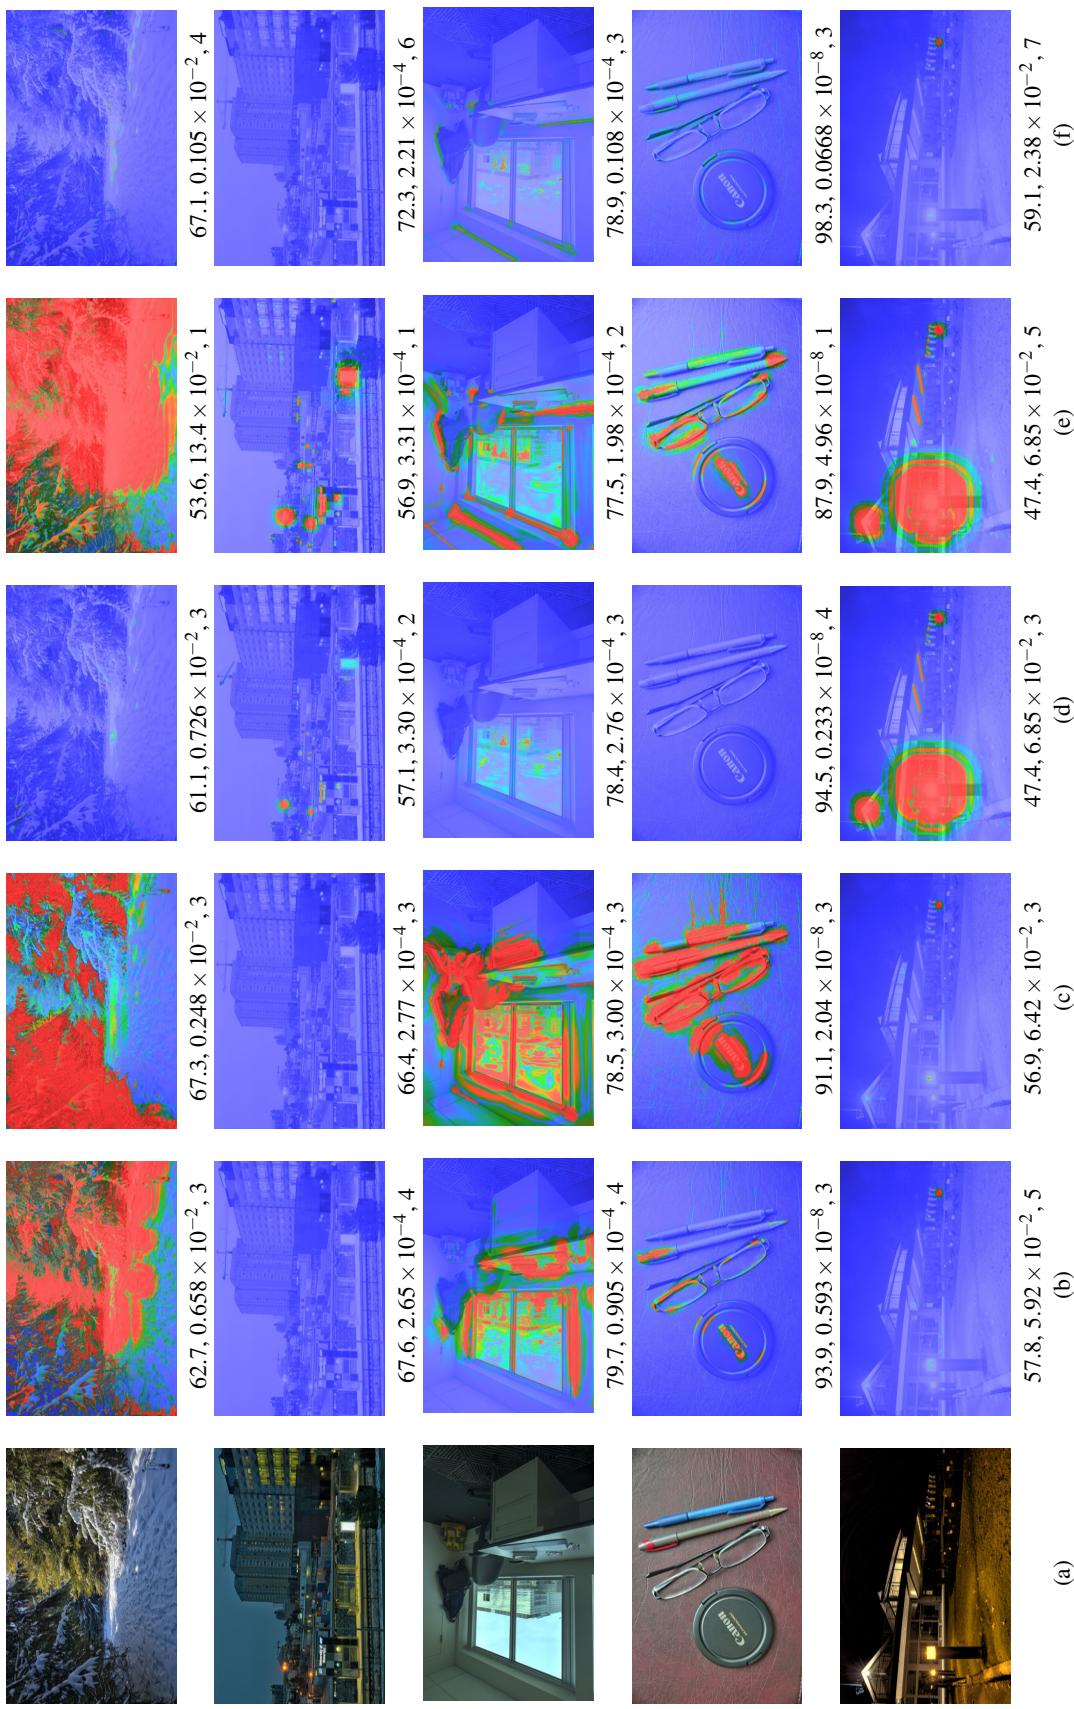
<sup>4</sup>Two of the scenes have dark regions (second and last row). In both of these scenes, Hasinoff et al.’s [10] method selected exposures that included 76 multiple captures of the longest shutter speed (32 sec.). In these cases, we restricted the exposure set to single captures; thus the results shown for these scenes may be somewhat worse than they would be.



**Figure 3.** For each method of selecting images and each benchmark in the HDR Photographic Survey: (a) quality correlate; (b) percentage of pixels with a probability  $\geq 0.75$  of a visually significant difference; and (c) mean squared error between test HDR images and ground truth HDR images. Within a method, a performance measure is sorted from smallest value to largest. Best viewed in color; our proposed set covering method has the highest quality and the lowest error.

while ensuring that the resulting HDR image faithfully captures the scene’s irradiance. On 110 benchmark scenes, our proposed method gave improved HDR images as measured against ground truth using a pixel-based metric and two perception-based metrics. As well, our experimental evaluation was the first to extensively evaluate existing state-of-the-art methods for image selection for stack-based HDR imaging.





**Figure 4.** For various benchmark scenes, (a) ground truth HDR image (tonemapped), (b)–(f) visual representation of prediction of visually significant differences from ground truth [31], mean squared error, and number of images selected for various exposure selection methods: (b) Barakat et al. [7]; (c) Pourreza-Shahri and Kehtarnavaz [15]; (d) Hasinoff et al. [10]; (e) Seshadriathan et al. [13]; and (f) our proposed set covering method. Best viewed in color; zoom in for additional detail.

## References

- [1] Steve Mann and Rosalind W. Picard, "On being "undigital" with digital cameras: Extending dynamic range by combining different exposed pictures," in *Proceedings of IS&T Annual Meeting*, 1993, pp. 442–448.
- [2] Paul E. Debevec and Jitendra Malik, "Recovering high dynamic range radiance maps from photographs," in *Proceedings of ACM SIGGRAPH*, 1997, pp. 369–378.
- [3] Erik Reinhard, Greg Ward, Sumanta Pattanaik, Paul Debevec, Wolfgang Heidrich, and Karol Myszkowski, *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*, Morgan Kaufmann, 2nd edition, 2010.
- [4] Christian Bloch, *The HDRI Handbook 2.0: High Dynamic Range Imaging for Photographers and CG Artists*, Rocky Nook, 2012.
- [5] Ting Chen and Abbas El Gamal, "Optimal scheduling of capture times in a multiple capture imaging system," in *Proceedings of Sensors and Camera Systems for Scientific, Industrial, and Digital Photography Applications III*, SPIE Vol. 4669, 2002.
- [6] Michael D. Grossberg and Shree K. Nayar, "High dynamic range from multiple images: Which exposures to combine," in *Proceedings of the ICCV Workshop on Color and Photometric Methods in Computer Vision*, 2003, pp. 1–8.
- [7] Neil Barakat, A. Nicholas Hone, and Thomas E. Darcie, "Minimal-bracketing sets for high-dynamic-range image capture," *IEEE Transactions on Image Processing*, vol. 17, pp. 1864–1875, 2008.
- [8] Natasha Gelfand, Andrew Adams, Sung Hee Park, and Kari Pulli, "Multi-exposure imaging on mobile devices," in *Proceedings of the ACM International Conference on Multimedia*, 2010, pp. 823–826.
- [9] Miguel Granados, Boris Ajdin, Michael Wand, Christian Theobalt, Hans-Peter Seidel, and Hendrik P. A. Lensch, "Optimal HDR reconstruction with linear digital cameras," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2010, pp. 215–222.
- [10] Samuel W. Hasinoff, Frédo Durand, and William T. Freeman, "Noise-optimal capture for high dynamic range photography," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2010, pp. 553–560.
- [11] Keigo Hirakawa and Patrick J. Wolfe, "Optimal exposure control for high dynamic range imaging," in *Proceedings of the IEEE International Conference on Image Processing*, 2010, pp. 3137–3140.
- [12] Orazio Gallo, Marius Tico, Natasha Gelfand, and Kari Pulli, "Metering for exposure stacks," *Eurographics*, vol. 31, pp. 479–488, 2012.
- [13] Kalpana Seshadrinathan, Sung Hee Park, and Oscar Nestares, "Noise and dynamic range optimal computational imaging," in *Proceedings of the IEEE International Conference on Image Processing*, 2012, pp. 2785–2788.
- [14] Kun-Fang Huang and Jui-Chiu Chiang, "Intelligent exposure determination for high quality HDR image generation," in *Proceedings of the IEEE International Conference on Image Processing*, 2013, pp. 3201–3205.
- [15] Reza Pourreza-Shahri and Nasser Kehtarnavaz, "Exposure bracketing via automatic exposure selection," in *Proceedings of the IEEE International Conference on Image Processing*, 2015, pp. 320–323.
- [16] Ayan Chakrabarti, Daniel Scharstein, and Todd Zickler, "An empirical camera model for internet color vision," in *Proceedings of the British Machine Vision Conference*, 2009.
- [17] Richard Szeliski, *Computer Vision: Algorithms and Applications*, Springer, 2010.
- [18] Glenn E. Healey and Raghava Kondepudy, "Radiometric CCD camera calibration and noise estimation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 16, pp. 267–276, 1994.
- [19] Tomoo Mitsunaga and Shree K. Nayar, "Radiometric self calibration," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2005, pp. 374–380.
- [20] Seon Joo Kim, Hai Ting Lin, Zheng Lu, Sabine Süsstrunk, Stephen Lin, and Michael S. Brown, "A new in-camera imaging model for color computer vision and its application," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, pp. 2289–2302, 2012.
- [21] Ce Liu, Richard Szeliski, Sing Bing Kang, C. Lawrence Zitnick, and William T. Freeman, "Automatic estimation and removal of noise from a single image," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 30, pp. 299–314, 2008.
- [22] Daniel Vaquero, Natasha Gelfand, Marius Tico, Kari Pulli, and Matthew Turk, "Generalized autofocus," in *Proceedings of the IEEE Workshop on Applications of Computer Vision*, 2011.
- [23] Michael R. Garey and David S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman, 1979.
- [24] George L. Nemhauser and Laurence A. Wolsey, *Integer and Combinatorial Optimization*, Wiley, 1988.
- [25] Anita Schöbel, "Set covering problems with consecutive ones property," Technical report, Universität Kaiserslautern, 2004.
- [26] Kristian Kirk and Hans Jørgen Andersen, "Noise characterization of weighting schemes for combination of multiple exposures," in *Proceedings of the British Machine Vision Conference*, 2006.
- [27] Mark A. Robertson, Sean Borman, and Robert L. Stevenson, "Estimation-theoretic approach to dynamic range enhancement using multiple exposures," *Journal of Electronic Imaging*, vol. 12, pp. 219–228, 2003.
- [28] Steffen Mecke and Dorothea Wagner, "Solving geometric covering problems by data reduction," in *Proc. of European Symposia on Algorithms*, 2004, pp. 760–771.
- [29] Mark D. Fairchild, "The HDR photographic survey," in *Color and Imaging Conference*, 2007, pp. 233–238.
- [30] Rafat Mantiuk, Kil Joong Kim, Allan G. Rempel, and Wolfgang Heidrich, "HDR-VDP-2: A calibrated visual metric for visibility and quality predictions in all luminance conditions," *ACM Transactions on Graphics*, vol. 30, no. 4, pp. 40:1–40:14, 2011.
- [31] Manish Narwaria, Rafal K. Mantiuk, Mattheiu Perreira Da Silva, and Patrick Le Callet, "HDR-VDP-2.2: A calibrated method for objective quality prediction of high dynamic range and standard images," *Journal of Electronic Imaging*, vol. 24, 2015.

## Author Biography

Peter van Beek received a BSc from the University of British Columbia (1984) and an MMath and PhD in Computer Science from the University of Waterloo (1986, 1990). He has been a faculty member in Computer Science at the University of Waterloo since 2000. In 2008, he was named a Fellow of the Association for Artificial Intelligence. His research interests span the field of AI with a focus on optimization, constraint programming, and applied machine learning.