

GENERATION OF HDR IMAGES IN NON-STATIC CONDITIONS BASED ON GRADIENT FUSION

Sira Ferradans¹, Marcelo Bertalmío², Edoardo Provenzi² and Vicent Caselles²

¹Centre De Recherche en Mathématiques de la Décision UMR CNRS 7534, Université de Paris Dauphine,
Place du Maréchal De Lattre De Tassigny, 75775, Paris Cedex 16, France

²Department of Information Technology, Universitat Pompeu Fabra, C/ Tànger 122-140, Barcelona, Spain

Keywords: High Dynamic Range Images, Gradient Fusion, Registration, Optical Flow.

Abstract: We present a new method for the generation of HDR images in non-static conditions, i.e. hand held camera and/or dynamic scenes, based on gradient fusion. Given a reference image selected from a set of LDR pictures of the same scene taken with multiple time exposure, our method improves the detail rendition of its radiance map by adding information suitably selected and interpolated from the companion images. The proposed technique is free from ghosting and bleeding, two typical artifacts of HDR images built through image fusion in non-static conditions. The advantages provided by the gradient fusion approach will be supported by the comparison between our results and those of the state of the art.

1 INTRODUCTION

Radiance in natural scenes can span several orders of magnitude, yet commercial cameras can capture only two orders of magnitude and reproduce Low Dynamic Range (LDR) photographs, often affected by saturated areas and loss of contrast and detail. During the last fifteen years, many techniques have been studied and proposed in order to expand the dynamic range of digital photographs to create the so-called High Dynamic Range (HDR) images, i.e. matrices whose entries are proportional to the actual radiance of the scene. For a thorough overview of the HDR imaging we refer the interested reader to the excellent book (Reinhard et al., 2005). HDR images have been widely used to capture scene illumination, but also for a realistic scene representation. That is to say, the wider dynamic range of HDR images allow to better capture details and color differences. This high precision is the reason why most photographers nowadays use cameras with a wider dynamic range than two orders of magnitude.

Two main problems remain to be solved in HDR imaging. First, while the technique proposed by Debevec and Malick (Debevec and Malik, 1997) can be considered as the ‘de facto standard’ for the creation of HDR images in *static conditions*, i.e. with perfectly still camera and without moving objects in the scene, no standard is available when these condition fail.

Second, HDR images cannot be entirely displayed or printed on the majority of commercial screens or printers, hence a so-called ‘Tone-Mapping’ transformation is needed to properly reduce their range without losing details and respecting as much as possible the original color sensation.

In this paper we deal only with the first problem, i.e. we study HDR image generation in non-static conditions, which are, by far, the most common. As we will explain in more detail later on, standard HDR image creation algorithms assume perfect alignment for the input images and no motion, the price to pay when this assumption fails is the generation of artifacts, in particular those called ‘ghosts’, i.e. objects with a translucent appearance induced by an incoherent image fusion, and ‘bleeding’, i.e. the diffusion of an artificial color over a flat image region, see Figure 1. To avoid this and other problems that we will discuss in more detail later, instead of synthesizing a new HDR image from the original sequence of LDR images, we will select just one of them to be the *reference* and then we will improve the associated radiance map by adding as many details as possible without introducing artifact. In doing this, a fundamental role will be played by gradient fusion, which has several advantages with respect to intensity fusion in this case. As we will show in Section 5, the results of the proposed technique compare well to the state of the art.



Figure 1: From top to bottom: an example of ghost and bleeding artifacts, respectively.

2 STATE OF THE ART IN HDR IMAGE CREATION

For the sake of readability, it is worthwhile to begin this section by introducing some concepts and notation that will be used throughout the paper. We define the *radiance* $E(x)$ to be the electromagnetic power per unit of area and solid angle that reaches the pixel x of a camera sensor with spatial support $\Omega \subset \mathbb{R}^2$. The *camera response function* f is the non-linear mapping $f : [0, +\infty) \rightarrow \{0, \dots, 255\}$ that transforms the radiance $E(x)$ acquired in the time Δt into $I(x)$, the *intensity* value of the pixel x , i.e. $f(E(x)\Delta t) = I(x)$. f is assumed to be semi-monotonically increasing, i.e. its derivative is supposed to be non-negative, since the digital values stored by the camera remain constant or increase as the radiance increases. In particular, the graph of f remains constant to 0 (black) for all the radiance values that do not overcome the sensibility threshold of the camera sensor and to 255 (white) for all those radiance values that exceed its saturation limit. For values that lie inside these two extremes, f is invertible and so we can use its inverse $f^{-1} : \{0, \dots, 255\} \rightarrow [0, +\infty)$ to pass from image intensity values $I(x)$ to the corresponding radiance values $E(x)$: $f^{-1}(I(x)) = E(x)\Delta t$. The camera response function can be built in controlled conditions by taking photos of a uniformly illuminated Macbeth chart with patches of known reflectance.

However, in many practical cases, the camera re-

sponse function is not known, thus various techniques for recovering f or its inverse without a Macbeth chart have been developed. These methods are called *chartless* and are all based on the same principle: since the camera sensor is limited by its sensibility threshold and saturation limit, in order to detect the whole scene radiance one has to take $N \geq 2$ shots with different time exposures Δt_j , $j = 1, \dots, N$. The next step consists in using the redundant information about the scene provided by the set of N digital images I_j to recover the inverse camera function f^{-1} . This is the only step that distinguishes the various algorithms proposed in the HDR imaging literature that we will briefly recall later. Let us for the moment suppose that we have indeed computed f^{-1} , then we can construct the j -th *partial radiance* $E_j(x)$ as follows: $E_j(x) \equiv f^{-1}(I_j(x)) / \Delta t_j$, $j = 1, \dots, N$. Notice that we use the adjective ‘partial’ when referring to the radiances E_j because they are built by applying f^{-1} to I_j , thus E_j cannot be considered a faithful representation of the entire radiance range, but only of the part corresponding to clearly visible details in the j -th image. Finally, in order to compute the set of final radiance values $\{E(x), x \in \Omega\}$ that will constitute the HDR image of the scene, one performs a weighted average of the partial radiances:

$$E(x) = \frac{\sum_{j=1}^N w(I_j(x)) E_j(x)}{\sum_{j=1}^N w(I_j(x))}, \quad (1)$$

where the weights w take their maximum in 127, the center of the LDR dynamic range, and decrease as the intensity values approach 0 or 255, when the information provided by I_j is imprecise. If one deals with *color images*, then this procedure must be repeated three times in order to recover three radiance functions relative to the red, green and blue spectral picks of the scene radiance.

As we have said, the algorithms for HDR image generation in static conditions are practically distinguished only for how they recover inverse camera response function f^{-1} . Mann and Picard (Mann and Picard, 1995; Mann, 2000) were the first to address this problem, they postulated a gamma-like analytical expression for f^{-1} and solved a curve fitting problem to determine the most suitable parameters. Mitsunaga and Nayar (Mitsunaga and Nayar, 1999) improved Mann and Picard’s work by assuming f^{-1} to be a polynomial and determining its coefficients through a regression model. Debevec and Malik (Debevec and Malik, 1997) worked with logarithmic data and did not impose any restrictive analytical form to f^{-1} , yet they required it to be smooth by adding a penalization term proportional to the second derivative of f^{-1} to the following optimization

problem: $\text{Min} \|\log f^{-1}(I_j(x)) - \log E(x) - \log \Delta t_j\|^2$. Their method is fast and gives reliable results, for this reason it has become the de facto standard in the field of HDR imaging with static conditions.

Let us now consider non static conditions, i.e. images taken with hand held camera and/or of dynamic scenes. In this case, a motion detection step must be added to register the set of input LDR images and a more careful fusion process must be considered to avoid artifacts.

Let us begin with registration: Ward (Ward, 2003) proposes to use multi-scale techniques to register a set of Median Threshold Maps (MTB), which are a binarization of the images with respect to their median value. Although this approach is independent of the exposure time, it depends on noise and histogram density around the median value. Grosch (Grosch, 2006) bases a local registration method on Ward's MTBs and Jacobs et al. (Jacobs et al., 2008) proposed an improvement of Ward's work by estimating motion with an entropy based descriptor. These methods give good results for the case of small movements generated by a hand held camera, but they tend to produce artifacts when dealing with large object motion, as e.g. people walking through the scene. Tomaszewska and Mantiuk (Tomaszewska and Mantiuk, 2007) and Heo et al. (Heo et al., 2011) propose to compute a global homography using RANSAC with SIFT descriptors which are based on gradients. Finally, Kang et al. (Kang et al., 2003) propose to use the camera function to boost the original images in order to facilitate the registration process.

Let us consider now the methods that focus mainly on a suitable improvement of the fusion step. As we have said before, if we assume a perfect correspondence among pixels in non-static conditions and perform a weighted average of radiances, then ghosts artifacts appear in areas where motion has occurred. Many methods in the state of the art, e.g. (Khan et al., 2006), (Heo et al., 2011), (Granados et al., 2010), (Grosch, 2006) and (Jacobs et al., 2008), focus on avoiding ghost formation by modifying eq. (1) in order to reduce the influence of pixels corresponding to moving objects in the process of intensity fusion to generate the HDR image. These methods may be robust to pixel saturation and small misalignments, but the areas that appear only in one image will be *copied* while the other image areas will be *averaged* from different images, thus new boundaries could be created in the process. In order to deal with these artifacts, Gallo et al. (Gallo et al., 2009) propose to create a vector field by copying patches of gradients from the best exposed areas that match a reference image. Then they blend the borders of neighboring patches

and integrate the vector field. The resulting images are ghost free but artifacts appear on the patch borders and flat regions. Finally, let us report that gradient-based methods can deal with the radiance differences among the image sequence but are sensitive to misalignments and can produce color bleeding as Eden et al. (Eden et al., 2006) have pointed out.

3 THE PROPOSED ALGORITHM

The method that we propose has several steps in cascade, before describing in detail all the steps it is worthwhile to give a brief summary of the whole algorithm. Firstly, as done in the static scenario, we apply the inverse camera function to every LDR image of the bracketing of pictures taken with different time exposures, obtaining the partial radiance map for each one of them. These radiance maps are radiometrically aligned using the camera function with the exposure time of the reference image. These modified images are used to compute a dense correspondence field, using the optical flow algorithm of Chambolle and Pock (Chambolle and Pock, 2011), for every image with respect to the reference one. These correspondences are subsequently filtered using a refinement step: for every pair of images, we compensate the motion between them and compute the absolute value image difference. The histogram of this image difference is modeled as a mixture of Gaussians, which allows us to distinguish between correct and erroneous correspondences. Finally, we use the corrected correspondences to obtain a gradient field that we integrate by solving a Poisson equation. To avoid color bleeding and other color artifacts we set a randomly selected group of points as Dirichlet boundary conditions. The intensity values of these points are computed through the Debevec-Malik intensity fusion of the aligned pixels.

The first step of the algorithm has already been described in the previous section, we present and discuss the following steps in separated subsections.

3.1 Radiance Registration

As previously declared, the goal of our method is to increase the detail rendition of the reference radiance map in the areas where the original LDR reference image was over/under exposed. The aim is to obtain these missing details from other images of the bracketing without creating artifacts produced by merging motion pixels. Thus, we need to distinguish motion pixels that we do not want to fuse from details that appear in other images (but not in the reference one)

that we actually want to fuse. To explain the problem let us consider the two HDR radiance maps shown in Figure 2.



Figure 2: Radiance maps of two pictures shown with the software HDRshop (available online at <http://www.hdrshop.com/>) useful to exemplify the radiance registration process (images courtesy of O.Gallo).

From the two images in Figure 2 we take the right one as reference. In this case we would like to introduce the detail of the texture on the wall present in the left image, avoiding the girls that cover it, given that we want to maintain the scene of the reference image. A registration algorithm based on intensity or geometry comparison cannot distinguish between wall details and people passing by, thus we need to modify the radiance maps before operating the registration, in such a way that the pixels corresponding to wall details are matched for fusion but not those corresponding to the girls. More generally, we need a *pre-registration step* that modifies the radiance maps so that the registration process will match the most precise details, i.e. those coming from dark areas of overexposed images and from bright areas of underexposed images.

Our tests have shown that the more efficient radiometric modification for this scope is given by the application of the camera response function with the (fixed) time exposure of the reference image, to each radiance map, as also observed by (Kang et al., 2003). Thus instead of I_j we work with the modified images $f(\Delta t_{\text{ref}} E_j(x))$. Notice that the reference image remains unchanged while for the others the transformation amounts to a normalization of the exposure time. Notice that, since the range of f is $\{0, \dots, 255\}$, after this process the radiance maps become a new set of LDR images. The advantage of this transformation is that the intensity levels become closer and that the image areas corresponding to the over/under-exposed values of the reference image also appear saturated in the other radiometrically aligned images.

After this pre-registration step, we can apply any optical flow method over the new set of LDR images. We used the method of Chambolle and Pock (Chambolle and Pock, 2011) which provides dense correspondences between pixels. Although the matchings are accurate, the method can produce errors on the

boundaries of moving objects and the areas where objects disappear due to motion. Thus a refinement step is needed to check whether the correspondences are correct or not.

3.1.1 Refinement Step

Let us assume that we have two radiometrically aligned images $I_{\text{ref}}, I_j, j = 1, \dots, N$, without motion. The pixel-wise distance between them defines the difference image $\|I_{\text{ref}}(x) - I_j(x)\|$ that depends on their noise. If we model this noise as additive and Gaussian, then the difference image histogram will be highly populated close to zero. As the motion (or mismatches) among the images increases, more modes appear in this histogram. We therefore propose modeling this difference histogram as a set of Gaussians using a Gaussian Mixture Model (GMM) where the most probable Gaussian is associated to the properly matched pixels and the other Gaussians represent matchings that are not correct. Let the most probable Gaussian be characterized by a median μ and a standard deviation σ . We define as a mismatch those correspondences with a distance from μ given by $\alpha\sigma$, where α is a parameter of the algorithm. After the refinement step, a pixel of the reference radiance map can have a set of correspondences in pixels of the other images of the sequence or no correspondence at all if the pixel belongs to an object that appears only in the reference image.

4 GRADIENT FUSION

After the steps previously discussed we know which pixels can be fused without generating artifacts in the final HDR image. Here we examine the process that will add details to the reference radiance map from the others in the set. Let us begin by observing that in general we do not have correspondences between all pixels of the reference image and pixels in all the other images of the bracketing, thus an intensity fusion can create artificial edges, as can be seen in Figure 3.

Note that the people are moving along the sequence, thus, their intensity values are being copied from the reference image to the final image while the other areas are computed by a weighted average among the corresponding pixels. This difference generates new edges around the subjects in motion. In order to avoid this problem, we propose using gradient fusion techniques in the Log-scale. Let us set for this scope $\tilde{I}_j = \log_{10}(I_j), j = 1, \dots, N$.



Figure 3: Detail of a result obtained by building radiance maps with an intensity fusion process, it can be seen that artificial edges appear (in particular around the moving people). Image (courtesy of O. Gallo) shown with HDRShop at a fixed time exposure.

Poisson Editing (Pérez et al., 2003) allows modifying the features of an image without creating new edges or artifacts in the LDR domain. The idea is to copy a target gradient field in a region Ω' of the image domain which is then integrated by solving a Poisson equation. Thus, the problem is to properly choose the target gradient field. To do that, recall that our aim is to improve the detail rendition of the final HDR image, so we must give more importance to gradients with largest norm. As discussed in (Piella, 2009), if we consider the weighted second moment matrix $G_s(x)$:

$$G_s(x) = \begin{pmatrix} \sum_{j=1}^N \left(s_j(x) \frac{\partial \tilde{I}}{\partial x_1} \right)^2 & \sum_{j=1}^N s_j^2(x) \frac{\partial \tilde{I}}{\partial x_1} \frac{\partial \tilde{I}}{\partial x_2} \\ \sum_{j=1}^N s_j^2(x) \frac{\partial \tilde{I}}{\partial x_2} \frac{\partial \tilde{I}}{\partial x_1} & \sum_{j=1}^N \left(s_j(x) \frac{\partial \tilde{I}}{\partial x_2} \right)^2 \end{pmatrix} \quad (2)$$

where the partial derivatives are computed at the point $x = (x_1, x_2)$ and $s_j(x) = \frac{\|\nabla \tilde{I}_j(x)\|}{\sqrt{\sum_{i=1}^N \|\nabla \tilde{I}_i(x)\|^2}}$, then the predominant gradient direction is given by the largest eigenvector of $G_s(x)$ and the corresponding eigenvalue gives its norm. In this process the sign of the gradient is lost. In our experiments, we have seen that the best results are achieved when we restore the sign to be that of the gradient with maximum modulus.

We can now define the target vector field for Poisson editing as follows: given a pixel x in the reference image, from all gradients of the bracketing of images at x , we select the gradient with largest modulus. Let us denote it by $M(x)$. Then the target vector V in x is defined as:

$$V(x) = \sqrt{\lambda} \varepsilon(x) \theta, \quad (3)$$

where λ is the largest eigenvalue of $G_s(x)$, θ its associated eigenvector, and

$$\varepsilon(x) = \begin{cases} -1 & \text{if } \langle \theta, M(x) \rangle \leq 0; \\ 1 & \text{otherwise.} \end{cases}$$

We stress that $V(x)$ is not necessarily a conservative vector field and, as remarked by Tao et al. (Tao et al., 2010) this can lead to bleeding effects along strong edges. In order to stop the bleeding effect, we have introduced Dirichlet boundary conditions at a set of randomly chosen points of Ω taken from the set of points for which there exist reliable correspondences. For them we have set the intensity as the radiance obtained with the Debevec-Malik intensity fusion.

5 RESULTS AND COMPARISONS

Now that the method has been described we can proceed to give some implementation detail. The optical flow method of Chambolle and Pock was applied to the luminance values of the radiometrically aligned images. For the refinement step we modeled the histogram of image differences as a GMM (Gaussian Mixture Model) with two components. A final step in the refinement process is introduced to avoid taking into account pixels on the boundaries: interpreting the mismatches as motion masks, we apply to these masks a dilation filter with a 6×6 structuring element (Gonzales and Woods, 2002). In the fusion process we randomly select the 10% of the total image area as possible Dirichlet points keeping only those with matches. Finally, we solve the Poisson equation using the Conjugate Gradient method for each color channel independently.

Let us now discuss the results of our proposal. Since HDR images cannot be represented on a LDR display, to show the results obtained from our algorithm we will show snapshots of the image provided by the free HDRShop software, available online at www.hdrshop.com, at a given exposure time. The comparison with the state of the art will be done following the procedure used in other papers, that is, by comparing the tone mapped versions of the HDR images.

First of all we would like to start by showing the improvement obtained using our method. In Figure 4 we can see two original partial radiance maps compared to our results. The enhancement is specially obvious on the building wall (first row), where our method recovers texture and color, and also in the trees, which show much more details (second row).

The parameter α controls the flexibility in the selection of matches in the registration process. Our



Figure 4: Left column: original reference radiance map at two different time exposures set in the HDRshop software. Right column: corresponding enhanced image with the method proposed in the paper. The parameter α was set to 1.

experiments have shown that a value of α that gives overall good performances is 0.5, for which we have found no artifacts in the final images. As we approach 1 we generate more details, but we might also create artifacts, as can be seen in Figure 5.



Figure 5: From left to right: results of our algorithm obtained by setting $\alpha = 0.5$ (no artifacts) and $\alpha = 1$ (artifacts appear), respectively.

The selection of the reference image can also influence the appearance of ghosts or artifacts. We have in fact observed that when a moving pixel with a color similar to the background is located in an area saturated by the camera function, then the refinement step may fail to distinguish background from foreground and color artifacts may appear. We can see an example in Figure 6: in the first row we present two images of the bracketing that were used as two different reference images. The corresponding results are shown in the second row. In both cases the parameter value is $\alpha = 1$. Note that the result shown on the left hand side has an artifact on the puppet's leg. The reason is that the reference image is saturated in that area, so the refinement step cannot distinguish between the puppet's leg and the ball.

Finally, we show the comparison between our re-

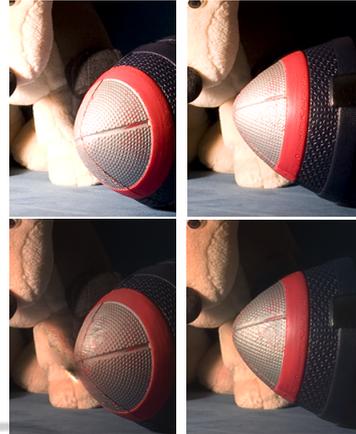


Figure 6: First row: two different original images taken as reference. Second row: output of our algorithm. It can be noticed an artifact in the left hand image due to the failure of the refinement step explained in the text.



Figure 7: First row, from left to right: tone mapped results of our HDR generation model and that of Gallo et al., respectively. Second row: detail magnification of the pictures above. The color difference is due to the different tone mapping operators used. We stress that the artifact visible in the image of Gallo et al's is not produced by the tone mapping operator, but by their HDR formation model.

sults and those of Gallo et al. (Gallo et al., 2009) who also use a gradient fusion technique. Since they present their results using tone mapped versions of their images, we also used a tone mapped version of ours. In Figure 7 we can see that our method avoids the generation of geometry artifacts as well as bleeding.

6 CONCLUSIONS

We have proposed a method for HDR image creation in the non-static scenario where photographs have been taken with a moving hand-held camera and/or the scene to be captured is dynamic. Our method is based on gradient fusion techniques that compares well to the state of the art. We have computed reliable correspondences by analyzing the differences of registered images using a GMM model. We have pointed out the importance of using gradient-based method instead of intensity-based methods. Finally, we proposed a technique in which a suitable selection of Dirichlet conditions permits to avoid or strongly reduce the bleeding artifacts that may appear after the integration of vector fields.

ACKNOWLEDGEMENTS

The authors would like to thank Dr. Orazio Gallo for kindly providing his images. They would also like to acknowledge partial support by PNPGC project, reference MTM2006-14836, and by GRC reference 2009 SGR 773 funded by the Generalitat de Catalunya. S. Ferradans was partially supported by ERC Grant n279593 SIGMA-Vision. E. Provenzi acknowledges the Ramoón y Cajal fellowship by Ministerio de Ciencia y Tecnología de España. V. Caselles also acknowledges 'ICREA Acadèmia' prize by the Generalitat de Catalunya.

REFERENCES

- Chambolle, A. and Pock, T. (2011). A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145.
- Debevec, P. and Malik, J. (1997). Recovering high dynamic range radiance maps from photographs. In *Proc. of the 24th annual conf. on Computer graphics*, pages 369–378.
- Eden, A., Uyttendaele, M., and Szeliski, R. (2006). Seamless image stitching of scenes with large motions and exposure differences. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, volume 2, pages 2498–2505.
- Gallo, O., Gelfand, N., Chen, W., Tico, M., and Pulli, K. (2009). Artifact-free high dynamic range imaging. *IEEE International Conference on Computational Photography (ICCP)*.
- Gonzales, R. and Woods, R. (2002). *Digital image processing*. Prentice Hall.
- Granados, M., Ajdin, B., Wand, M., Theobalt, C., Seidel, H.-P., and Lensch, H. (2010). Optimal hdr reconstruction with linear digital cameras. In *2010 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 215–222.
- Grosch, T. (2006). Fast and robust high dynamic range image generation with camera and object movement. In *Vision, Modeling and Visualization, RWTH Aachen*, pages 277–284.
- Heo, Y., Lee, K., Lee, S., Moon, Y., and Cha, J. (2011). Ghost-free high dynamic range imaging. In *Computer Vision - ACCV 2010*, volume 6495 of *Lecture Notes in Computer Science*, pages 486–500. Springer Berlin Heidelberg.
- Jacobs, K., Loscos, C., and Ward, G. (2008). Automatic high-dynamic range image generation for dynamic scenes. *IEEE Computer Graphics and Applications*, 28:84–93.
- Kang, S., Uyttendaele, M., Winder, S., and Szeliski, R. (2003). High dynamic range video. *ACM Trans. Graph.*, 22:319–325.
- Khan, E., Akyuz, A., and Reinhard, E. (2006). Ghost removal in high dynamic range images. In *IEEE International Conference on Image Processing*, pages 2005–2008.
- Mann, S. (2000). Comparametric equations with practical applications in quantigraphic image processing. *Image Processing, IEEE Transactions on*, 9(8):1389–1406.
- Mann, S. and Picard, R. W. (1995). On being undigital with digital cameras: Extending dynamic range by combining differently exposed pictures. In *Proceedings of IS&T*, pages 442–448.
- Mitsunaga, T. and Nayar, S. (1999). Radiometric self calibration. In *Computer Vision and Pattern Recognition, 1999. IEEE Computer Society Conference on*, volume 1, pages 374–380 Vol. 1.
- Pérez, P., Gangnet, M., and Blake, A. (2003). Poisson image editing. In *ACM SIGGRAPH 2003 Papers*, pages 313–318.
- Piella, G. (2009). Image fusion for enhanced visualization: A variational approach. *International Journal of Computer Vision*, 83:1–11.
- Reinhard, E., Ward, G., Pattanaik, S., and Debevec, P. (2005). *High Dynamic Range Imaging, Acquisition, Display, And Image-Based Lighting*. Morgan Kaufmann Ed.
- Tao, M., Johnson, M., and Paris, S. (2010). Error-tolerant image compositing. In *Proceedings of the 11th European conference on Computer vision: Part I, ECCV 2010*, pages 31–44.
- Tomaszewska, A. and Mantiuk, R. (2007). Image registration for multi-exposure high dynamic range image acquisition. In *Proc. Intl. Conf. Central Europe on Computer Graphics, Visualization, and Computer Vision (WSCG)*.
- Ward, G. (2003). Fast, robust image registration for compositing high dynamic range photographs from hand-held exposures. *Journal of graphic tools*, 8:17–30.