Published in final edited form as:

Proc IEEE Workshop Appl Comput Vis. 2012 January 9; 2012: 225–232. doi:10.1109/WACV.2012.6163015.

Image Alignment for Multiple Camera High Dynamic Range Microscopy

Brian S. Eastwood and Elisabeth C. Childs
Colby College, Waterville, Maine

Abstract

This paper investigates the problem of image alignment for multiple camera high dynamic range (HDR) imaging. HDR imaging combines information from images taken with different exposure settings. Combining information from multiple cameras requires an alignment process that is robust to the intensity differences in the images. HDR applications that use a limited number of component images require an alignment technique that is robust to large exposure differences. We evaluate the suitability for HDR alignment of three exposure-robust techniques. We conclude that image alignment based on matching feature descriptors extracted from radiant power images from calibrated cameras yields the most accurate and robust solution. We demonstrate the use of this alignment technique in a high dynamic range video microscope that enables live specimen imaging with a greater level of detail than can be captured with a single camera.

1. Introduction

Cameras suffer from a fundamental limitation of dynamic range. A wide variety of common scenes exhibit a large variation in light intensity—*e.g.* indoor spaces with windows, specular surfaces, back-lit objects, and dark shadows. Camera image sensors, however, can only record a finite range of image intensities. The photographer selects exposure settings that determine what range of intensities are properly recorded in the acquired image. Image regions with intensity levels outside the range captured by the sensor suffer from lack of detail, appearing either underexposed or overexposed.

High dynamic range (HDR) imaging addresses this limitation by combining information from multiple images acquired with different exposure settings. Using a larger range of exposure settings yields a higher potential dynamic range in the composite image. HDR imaging has become increasingly popular in recent years, with tools appearing in consumer digital cameras and post-processing software for professional and amateur photographers.

HDR imaging provides benefits to microscopy as well. Microscopy specimens can exhibit a wide range of optical densities, and consequently a microscope scene can contain more intensity variation than can be captured with a single image. As microscopy image analysis becomes increasingly computer-based, the need to acquire accurate and detailed digital images increases. Improper image exposure poses problems to many image processing techniques, and HDR imaging can eliminate this problem. Current HDR microscopy systems, however, are limited to observing stationary specimens—no known systems have adequately addressed HDR imaging of moving specimens.

This paper presents two novel results. First, we address the challenge of aligning HDR component images acquired by multiple cameras at different exposure settings. Because a

greater exposure separation enables a greater dynamic range, we evaluate three alignment techniques that are robust to exposure differences. Second, we address the problem of HDR microscopy of moving specimens by simultaneously capturing multiple images using a pair of cameras. The system described enables observation of moving specimens with a greater level of detail than is possible with existing single camera systems.

The remainder of this paper is organized as follows. Section 2 provides an overview of HDR imaging and presents the challenges in acquiring HDR images of moving objects. Section 3 discusses the camera calibration process including procedures adapted for microscopy. Section 4 evaluates image alignment techniques for multiple cameras operating at different exposure settings. Section 5 presents the implementation of a high dynamic range video microscopy system. Section 6 discusses the broader implications of this work and suggests avenues of future research.

2. HDR Imaging

Dynamic range is a measure of the range of intensity values present in a scene or accurately recorded in an image. Dynamic range is conveniently measured on the exposure value (EV) scale:

$$R = \log_2 \frac{E_{max}}{E_{min}}, \quad (1)$$

where E is a linear measure of intensity. The EV scale is used by photographers to measure exposure control, where 1 EV corresponds to a doubling of light sensed through increasing the shutter speed, lens aperture, or sensor gain.

HDR imaging extends the dynamic range of an image sensor by combining multiple images taken with different exposure settings. The imaging technique involves two steps: (1) calibrating a camera response function and (2) combining the linear measures of light intensity obtained from multiple exposures of a single scene.

2.1. Camera Response Function

A camera exposure samples the image irradiance function formed by the lens on the image sensor. Digital image sensors accumulate charge at each photosite proportional to the image irradiance, integrated over the area of the photosite and scaled by the spectral response and quantum efficiency of the sensor [8]. The quantity integrated at each photosite is radiant power, ϕ . Integrated over the exposure time, t , each photosite records a measure of radiant energy as a voltage. The camera response function transforms this radiant energy measurement into a pixel intensity, z , recorded as a digital number [3, 7].

$$z = f(\Phi t) \quad (2)$$

While most image sensors have a linear response to radiant energy over much of their range, the pixel intensities recorded in a digital image cannot be taken as linear measures of light intensity. Specifically, there are discontinuities in the camera response function at the noise floor and saturation points of the device. Additionally, some camera manufacturers intentionally employ nonlinear response functions [5].

Several techniques exist to estimate a camera's response function or inverse response function. The inverse response function is a map of pixel intensity value to the radiant energy incident on the image sensor,

$$f^{-1}(z) = \Phi t. \quad (3)$$

Debevec and Malik recover a discrete inverse response function from a series of images of a static scene taken at different shutter speeds [5]. Bell *et al.* directly measure the camera response to a known radiant power provided by a reference illuminant [2]. Armed with an inverse camera response function, the radiant power incident on an image sensor can be recovered from an image acquired at a known shutter speed,

$$\Phi = \frac{f^{-1}(z)}{t} \quad (4)$$

2.2. Image Composition

HDR image composition combines the linear measures of radiant power recovered from multiple images taken with different shutter speeds. Each image contains different properly exposed regions corresponding to the radiant power that is properly sampled by the sensor. Combining images with a weighted average yields a HDR radiant power map,

$$\Phi_{hdr}(x, y) = \frac{\sum_i w(I_i(x, y)) \Phi_i(x, y)}{\sum_i w(I_i(x, y))}, \quad (5)$$

where I_i and Φ_i are the i^{th} pixel intensity and radiant power images in the series. The weighting function $w(\cdot)$ serves to discount unreliable measurements from the shoulder regions of the response function [3, 5].

If the component radiant power images span the entire range of intensities represented in the scene, the composite image has all regions properly exposed. The result is that detail is recorded in all regions of the image. Additionally, because the weighted average of Equation 5 combines multiple measurements at each pixel, HDR images exhibit a noise reduction in addition to the extended dynamic range.

2.3. Moving Objects

HDR imaging is broadly applicable to macroscopic and microscopic applications. The use of HDR imaging in microscopy becomes important for accurate computational analysis of specimens. Without it, images may contain underexposed or overexposed regions where the signal from the specimen is completely lost. Analysis algorithms get no information in these regions. For example, motion detection algorithms cannot recover motion within constant-valued regions; segmentation and classification algorithms have no information by which to discriminate objects.

HDR imaging is straightforward to apply to fixed specimens, and several existing computer-integrated microscopy systems include such a tool [2]. When observing moving specimens, however, one needs to account for the motion between frames acquired at different exposures. No known systems have addressed this problem adequately to date.

Several macroscopic imaging techniques address the issue of moving subjects in HDR imaging and compositing. Uyttendaele *et al.* detect moving objects in panorama compositing, and select a single instance of each object for the final image [11]. This approach is impractical in microscopy because it is often the moving objects that exhibit HDR content, so information needs to be combined from multiple images of these objects.

Kang *et al.* use rigid registration to account for global camera motion and scene motion in disjoint regions of the frame [10]. This approach is not practical in microscopy because the motion of biological specimens is rarely rigid. Nonrigid motion estimation in images with different exposure values is difficult because intensity constancy cannot be assured [1].

Wilburn *et al.* use a large array of cameras operating at different shutter speeds to capture detail in all regions of a scene [13]. Of these approaches, this is most similar to ours. However, we explore the constraint of having a very limited number of cameras with which to capture the entire range of light intensities present in a scene. We also provide additional details of system implementation related to microscopy imaging.

We resolve the motion issue by using two (or more) cameras to acquire images simultaneously. The cameras observe approximately the same field of view, but use different shutter speeds to sample a broader range of radiant power.

This approach presents several issues. First, the camera response functions need to be adjusted to the same radiant power scale. Second, the images acquired from each camera need to be aligned before assembling the HDR composite. For maximum flexibility in a live experiment, the alignment procedure should be robust to differences in exposure value.

This approach places the emphasis of image composition on increasing dynamic range, not noise reduction. With only two images combined in each composite, any noise reduction is minor. Moreover, the dynamic range increase comes from pushing the two exposure values as far apart as possible. This reduces the number of pixels that are well-exposed in both images and averaged in the composite.

3. Camera Calibration

Calibration for HDR microscopy imaging involves several stages. The camera response function enables recovering light intensity measurements in a linear space. Color balancing corrects for differences in the spectral emission of the microscope lamp and the response of different color pixels. Flat-fielding accounts for non-uniform illumination of the lamp. In non-microscopy applications, flat-field correction is similar to vignetting correction.

3.1. Camera Response and Color Correction

The inverse camera response function (Equation 3) needs to be recovered for each camera. This relates the digital number recorded at each pixel in the image to a radiant energy accumulated at the corresponding photosite. For our purposes, we select the approach proposed by Debevec and Malik [5]. The advantage of this method is that it requires no special calibration target; the disadvantage is that each response function is recovered in an arbitrary radiant power scale. A separate response function is recovered for each color channel on each camera.

Color calibration rescales the response functions to the same linear radiant power scale. That is, the inverse camera response functions are related by constant scale factors, w_c

$$f_c^{-1}(z) = w_c \Phi t \quad (6)$$

Color calibration involves selecting a small image patch, Ω , of a gray object. The average radiant power measured in each color channel should be equivalent within this patch.

$$\mu_c = \frac{1}{n} \sum_{(x,y) \in \Omega} \Phi_c(x,y), \quad (7)$$

where n is the number of pixels in the patch and Φ_c is the radiant power measured in channel c . The radiant power for each channel can therefore be adjusted to the same scale, Φ' , in relation to a reference radiant power, μ_0 :

$$\Phi'(x,y) = \Phi_c(x,y) * \frac{\mu_0}{\mu_c}. \quad (8)$$

To correlate the response functions from multiple cameras, each camera uses an image patch of the same gray object and a single reference average radiant power, μ_0 , is selected for all cameras.

In practice, a gray image patch can be acquired for microscope calibration by moving the stage to an empty region of the slide. An image acquired here records only the light from the illuminating lamp and color correction serves to neutralize the illumination spectrum. The region selected for color correction should not contain underexposed or overexposed pixels. In non-microscopy applications, color balancing is often done using a gray card target.

3.2. Flat-field Correction

Flat-field correction adjusts for non-uniform illumination and accounts for fixed-pattern noise in the image sensor [9, Ch. 12]. A flat image is found by selecting a field of view that is known to be uniform. The measured radiant power should be consistent across this field. A mean flat image, Φ_f , is the temporal average of a series of m radiant power images, Φ_i , of the flat scene.

$$\Phi_f(x,y) = \frac{1}{m} \sum_i \Phi_i(x,y) \quad (9)$$

The mean value of the flat image is the expected radiant power at all pixels,

$$\mu_f = \frac{1}{n} \sum_{(x,y)} \Phi_f(x,y), \quad (10)$$

where n is the number of pixels in the image and the summation is performed over the entire frame. A correction that rescales radiant power to be uniform everywhere in the field is then:

$$\Phi'(x,y) = \Phi(x,y) \frac{\mu_f}{\Phi_f(x,y)}. \quad (11)$$

Both color and flat-field correction apply linear scale factors to each radiant power measure, and are therefore very similar. Color correction provides a per-channel scaling while flat-field correction provides a per-pixel scaling. In practice, a single flat-field image can also be used to perform color correction.

It is important to emphasize that both correction factors operate with the linear radiant power measurements. It is common to apply these techniques using pixel value images, but such an approach suffers from the nonlinearity of pixel intensities previously mentioned.

It is also important to note the assumptions of invariant measurements in the calibration procedures. Recovering the camera response function involves taking repeated measurements at individual pixels; what matters is that the radiant power incident at each pixel remains constant over the sequence of images. Color correction involves comparing measurements within a small image patch; what matters is that the mean radiant power within the patch is equivalent for each color channel. Flat-field correction involves comparing measurements across the entire image frame; what matters is that the radiant power is equivalent at all pixels.

4. Image Alignment

Images need to be aligned before compositing an HDR image. Placing the images in the same reference frame aligns the multiple radiant power measurements taken at each image location. The weighted average of radiant power measurements (Equation 5) can then be used to assemble a composite HDR image.

The image analysis literature provides a broad array of techniques for aligning images, including global image registration, phase correlation, and feature-based alignment [1]. Choosing an appropriate technique involves selecting an approach that is robust to the variation present in the images and selecting a motion model that describes the motion between images.

In HDR applications, the images to be aligned are acquired at different exposure settings. An image alignment procedure can be completed on images with the same exposure settings if the alignment is guaranteed to remain fixed. Some applications, however, require alignment to be carried out with different exposure settings. For example, an image series taken with a hand-held camera exhibits small camera motion between frames [12]. Multiple camera systems that use moving cameras also need to be aligned on the fly.

In multiple camera microscopy systems, changing objective lenses or adjusting focus can alter the image alignment. Most microscopes use telecentric optics to provide orthographic projection, eliminating alignment differences from focusing, but this effect varies from system to system.

Under these conditions, the alignment technique needs to be robust to exposure differences. We discuss three such approaches below.

4.1. Alignment Methods

Ward proposes a method to align images acquired with different shutter speeds based on global alignment of median thresholded images [12]. This technique leverages the fact that the median intensity level present in a scene is not dependent on exposure. Correlation of median threshold bit masks provides a more reliable alignment metric than correlation of the images themselves.

Median threshold alignment is performed using a multiscale, coarse-to-fine optimization on the correlation of median threshold bit masks. At each pyramid level, a moving image is shifted ± 1 pixel along each axis with respect to the fixed image. The translation that produces the best correlation metric is propagated to the next finer level of the pyramid. Upsampling the images before alignment provides subpixel accuracy—we use two additional pyramid levels to obtain alignment within 0.25 pixel. While the technique can be adapted to recover any rigid registration, we do not investigate this here.

We propose a minor modification to median threshold alignment. Ward notes that, depending on the scene composition, an intensity value other than the median—*e.g.* the 25th

percentile—may be a better threshold choice. We therefore use three threshold bit mask images simultaneously, obtained at the 25th, 50th, and 75th percentiles.

Another class of exposure-robust alignment methods is based on matching image feature descriptors. Brown *et al.* describe one such approach [4]. The locations of salient features in each image are determined using an interest point detector, such as the Harris corner detector. Image descriptors are extracted from the neighborhood around each interest point—we use axis-aligned image intensity patches within a 41×41 pixel region downsampled to a 7×7 descriptor. Initial feature correspondences between the two images are determined using an exhaustive pair-wise search for minimum Euclidean distance between descriptors. Correspondences are eliminated for which the ratio of scores from the nearest match to the second nearest match is too large. Random sample consensus (RANSAC) is used to find the best collection of feature correspondences that provides a consistent transform between images. The transform model recovered by RANSAC is flexible—we use either translation or affine transforms.

Within this framework, the choice of feature descriptors accounts for the difference in exposure between images. One simple approach is to use normalized descriptors extracted from the pixel values obtained directly from the camera. Here, the normalization of each descriptor provides robustness to exposure value differences.

A second approach is to use normalized descriptors extracted from the radiant power images recovered from the camera response function. Here, descriptors extracted from well-exposed regions of both images are measurements taken in the same linear space, and should therefore be invariant to exposure differences.

Note that all of these methods—median threshold alignment, pixel value feature matching, and radiant power feature matching—depend on an overlap of well-exposed regions in the image pair. None of these techniques will succeed if the exposure difference maps underexposed pixels in one image to overexposed pictures in the second image.

It is important to understand, then, which of these methods performs best for aligning images acquired at different exposure settings. Because a greater EV difference between images yields a greater potential dynamic range in a composite image, the method that provides the least alignment error at the greatest EV difference is the most suitable for HDR applications.

4.2. Alignment Evaluation

We use synthetic images to evaluate how robust each method is to differences in exposure setting. We recover a high quality HDR radiant power image of a stationary scene using many exposures from a single camera. The radiant power map and a forward camera response function (Equation 2, $z = f(\phi t)$) enables synthesizing ideal pixel value images at different virtual shutter speeds. Gaussian noise is added to each pixel to simulate camera noise. Ideal pixel values are quantized and clipped to the range $[0 \dots 255]$.

A pair of images synthesized with different virtual shutter speeds are translated randomly and cropped to obtain images with a known ground-truth displacement. Figure 1 displays images generated in this manner. Only translation transforms are evaluated in this analysis because the median threshold alignment does not account for rotation.

For each exposure value difference, we generate 30 synthetic image pairs with random displacements of no more than 50 pixels along each axis. Image alignment is performed on

the image pair with each of the alignment techniques, and the error with respect to the ground-truth displacement is recorded.

Table 1 summarizes the root mean squared (RMS) alignment error and failure rates obtained over a range of EV differences for the three alignment methods. An alignment is considered a failure if the error is greater than 5 pixels. RMS errors are computed only for successful alignments so that egregious failures do not skew the error for an otherwise successful set of alignments. A method is considered a complete failure at an EV difference if more than 20% of the alignments are unsuccessful.

Table 1 indicates that median threshold alignment is prone to occasional failure at all EV differences. This occurs when the coarse-to-fine optimization gets caught in a local minimum on one level and steered away from the globally optimal solution. The feature-based alignment techniques, however, are robust as long as there are sufficient well-exposed features in both images to enable building a transform model in RANSAC.

Figure 2 summarizes the RMS alignment errors graphically. The two feature-based alignment techniques are well matched up to exposure differences of 2 EVs. At greater EV differences, the descriptors extracted from radiant power images prove to be more accurate and robust to exposure differences. The median threshold alignment technique is robust to exposure differences but is less accurate than the feature-based techniques. No method operates reliably beyond a 4 EV difference.

From this analysis, we select matching image feature descriptors extracted from radiant power images as the preferred alignment technique for images acquired with different exposure settings. We reiterate that this result is applicable to alignment in any HDR imaging application, not just the microscopy application described in Section 5.

5. HDR Video Microscope

We now describe an HDR video microscope that enables observation of moving specimens. Although HDR imaging has been used in microscopy, the use has been limited to stationary specimens. The key innovation here is to use a pair of cameras attached to the microscope to simultaneously acquire images at different shutter speeds. The images are then aligned and combined in an HDR image.

Figure 3 shows the HDR video microscope implemented on a Zeiss Axio Lab.A1 bright field microscope. A pair of Point Grey Research Flea (FL2G-13S2C) IEEE 1394 video cameras are attached to 0.5× video couplers and mounted in place of the ocular lenses. (Another option would have been to attach a beam splitter and a pair of cameras through the trinocular photoport tube. Our solution was adopted because it is less expensive and the cameras we used were light enough to not introduce any instability or mechanical stress.) The cameras are attached to a stabilizing bracket that minimizes, but does not eliminate, orientation differences and ensures the cameras do not move during an experiment.

The majority of our image processing system was implemented in Python. C libraries are used for the camera control and image acquisition. We use the Debevec and Malik approach to calibrate the camera response functions [5] and perform color and flat-field correction as described in Section 3. Image alignment uses the radiant power feature matching technique discussed in Section 4 to recover an affine transform.

Combining images of a moving specimen acquired from the camera pair requires that the frames are synchronized. Otherwise, specimen motion between frames will lead to ghosting effects in the composite HDR image. The cameras that we use (as well as many camera

models from other manufacturers) automatically synchronize frame acquisition to within 125 μ s when the cameras are connected to the same IEEE 1394 bus and are configured to use the same color mode, frame size, and frame rate.

Taking advantage of the hardware synchronization, however, requires that the software keeps up with the acquisition frame rate on both cameras—the frame buffers cannot be allowed to fill or frames will be dropped and possibly fall out of sync. The time required to apply the camera response function, color-correction, flat-fielding, image transformation, HDR image composition, and tone mapping is too great to be able to process frames at the acquisition rate of 15 fps.

The HDR video application therefore uses two processing threads. One thread displays live images from both cameras, which enables focusing and field of view selection in addition to keeping the frame buffers flushed. The other thread applies all HDR image processing from the two most recently acquired frames. While this means that the HDR video frame rate is less than the camera acquisition rate, one could obtain a full frame rate video by saving the raw video frames to disk and performing HDR composition off line.

Figure 4 shows a set of microscopy images of a living *Daphnia* specimen—a small crustacean found in fresh water lakes and rivers. Figure 4a is an image from a single camera taken with auto exposure and auto white balance. This image demonstrates the base level of detail available from a single camera. Note that the camera's built-in color balance cannot eliminate the color cast from the illuminating lamp.

The remaining images in Figure 4 are an HDR sequence assembled from image pairs acquired with shutter speeds that differ by 2.5 EVs. The image alignment is determined using feature descriptor matching on the radiant power images and an affine transform motion model. Each composite HDR image is tone-mapped for display using the technique of Durand and Dorsey [6]. This tone mapping procedure is a local contrast adjustment that makes it possible to better represent the full range of detail present in the image on limited dynamic range display devices, such as computer monitors or print. It is stressed that the HDR radiant power image has more detail than can be displayed, so some local contrast adjustment is needed for visualization.

The absence of ghosting in these images demonstrates that the frames are synchronized despite specimen motion. In this experiment, which uses a 3.75 μ m photosite size and 5 $\times$ magnification, specimen motion at rates up to 6000 μ m/s will result in no more than 1 pixel displacement between frames. The activity of the *Daphnia* during this sequence was continuous—the microscope stage was constantly adjusted to keep the specimen in the field of view, and the specimen's arms moved rapidly.

The bright regions along some image borders are areas where only the long exposure image contributes to the composite HDR image, and those pixels are overexposed. Although cropping would eliminate such regions, they are left in to illustrate the alignment.

There is a visibly greater level of detail in the HDR images, particularly within the dark organs in the abdomen of the *Daphnia*. These details are recovered without losing information in lighter regions, such as the arms. That is, the HDR images recover detail in regions of the specimen that are lost in single camera images.

Figure 5 demonstrates that the HDR video microscopy technique also improves dark field microscopy. Dark field illumination captures only scattered light rays, which is useful for observing nearly transparent specimens. Figure 5a is a dark field image acquired from a single camera; Figure 5b displays a tone-mapped HDR composite of two images.

Longer shutter speeds are required to capture well-exposed images at the lower light level present in dark field illumination. This leads to more frequent motion blur when the specimen moves during a single exposure. This is not a problem in the HDR technique, but rather a problem in any application where the shutter speed is too long.

6. Discussion

We have presented a high dynamic range video microscopy system that enables recording images of live specimens with a greater level of detail than possible with a single camera. This innovation employs a pair of synchronized cameras from which images can be combined with consistent image content.

The system relies on an alignment technique that is robust to differences in exposure. Because the system combines only two images at a time, it is desirable to push the camera shutter speeds far apart to yield the greatest dynamic range. We demonstrate that matching feature descriptors extracted from radiant power images provides a more accurate and robust alignment than median threshold alignment and feature matching with pixel value descriptors.

If the EV difference is pushed too far, however, there will be a gap in the range of accurately measured radiant powers. Adding more cameras is one way to extend the dynamic range capabilities even further. If the specimen is moving slowly relative to the cameras' frame rate, another option would be to acquire multiple images at different shutter speeds with each camera. This could leverage the bracketing modes available on some cameras that cycle through a set of shutter speeds.

The result of our image alignment analysis is instructive for other HDR applications. For example, HDR imaging features are offered in a growing range of consumer digital cameras and image processing software. Our approach would enable composing images with a broad dynamic range from fewer exposures taken with a hand-held camera. This arrangement reduces the possibility of motion in the scene and thereby reduces ghosting artifacts in the composite image.

The multiple camera system described here has broader implications for imaging systems. One of the major factors influencing the cost of a digital camera is the dynamic range of its image sensor. Combining images from multiple cameras enables extending the dynamic range. This poses the question of whether a pair of moderately priced cameras could produce images with a quality that matches or even surpasses significantly more expensive cameras. Microscopy offers a rich application area in which to investigate these issues.

Acknowledgments

The project described was supported by the Howard Hughes Medical Institute Science Education Grant number 52006288, the National Center for Research Resources award number P20RR016463, and Colby College. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Center for Research Resources or the National Institutes of Health. Thanks to Dr. Bruce Maxwell, Charles Jones, and Grovenia Perryman for their contributions to this project.

References

- [1]. Beauchemin SS, Barron JL. The computation of optical flow. *ACM Computing Surveys*. 1995; 27(3):433–466.
- [2]. Bell AA, Brauers J, Kaftan JN, Meyer-Ebrecht D, Boecking A, Aach T. High dynamic range microscopy for cytopathological cancer diagnosis. *IEEE Journal of Selected Topics In Signal Processing*. 2009; 3(1):170, 184.

- [3]. Bell, AA.; Seiler, C.; Kaftan, JN.; Aach, T. Noise in high dynamic range imaging; IEEE International Conference on Image Processing; 2008; p. 561-564.
- [4]. Brown, M.; Szeliski, R.; Winder, S. Multi-image matching using multi-scale oriented patches; IEEE Conference on Computer Vision and Pattern Recognition; 2005; p. 510-517.
- [5]. Debevec, PE.; Malik, J. Recovering high dynamic range radiance maps from photographs; ACM Conference on Computer graphics and interactive techniques, Proc. SIG-GRAPH 1997; 1997; p. 369-378.
- [6]. Durand, F.; Dorsey, J. Fast bilateral filtering for the display of high-dynamic-range images; ACM Transactions on Graphics, Proc. SIGGRAPH 2002; 2002; p. 257-266.
- [7]. Hasinoff, SW.; Durand, F.; Freeman, WT. Noiseoptimal capture for high dynamic range photography; IEEE Conference on Computer Vision and Pattern Recognition; 2010;
- [8]. Healey G, Kondepudy R. Radiometric ccd camera calibration and noise estimation. IEEE Transactions on Pattern Analysis and Machine Intelligence. 1994; 16(3):267, 276.
- [9]. Inoué, S.; Spring, KR. Video Microscopy : The Fundamentals. 2nd edition. Springer; 1997.
- [10]. Kang SB, Uyttendaele M, Winder S, Szeliski R. High dynamic range video. ACM Transactions on Graphics. Jul.2003 22:319–325.
- [11]. Uyttendaele, M.; Eden, A.; Szeliski, R. Eliminating ghosting and exposure artifacts in image mosaics; IEEE Conference on Computer Vision and Pattern Recognition; 2001. p. 509-516.
- [12]. Ward G. Fast, robust image registration for compositing high dynamic range photographs from handheld exposures. Journal of Graphics Tools. 2003; 8:17–30.
- [13]. Wilburn, B.; Joshi, N.; Vaish, V.; Talvala, E-V.; Antunez, E.; Barth, A.; Adams, A.; Horowitz, M.; Levoy, M. High performance imaging using large camera arrays; In ACM Transactions on Graphics, Proc. SIGGRAPH 2005; 2005. p. 765-776.

Figure 1.
Image pair acquired with a virtual shutter speed difference of 1.5 EV, synthesized from an HDR image of octopus muscle tissue.

Figure 2.
Image alignment error as a function of exposure value difference between images.

Figure 3.
Upright bright field microscope fitted with two IEEE 1394 video cameras in place of ocular lenses.

Figure 4. Images of a live *Daphnia* specimen; (a) image from a single camera with auto exposure and white balance; (b - f) composite HDR images with tone-mapping.

Figure 5. Images of a live *Daphnia* specimen in dark field illumination; (a) image from a single camera with auto exposure and white balance; (b) composite HDR image with tone mapping.

Table 1

Error and failure rate for alignment methods: median threshold alignment, feature matching with pixel value descriptors, and feature matching with radiant power descriptors. RMS error is measured in pixels.

ΔEV	median		feature-pixel		feature-radiant	
	RMSE	% fail	RMSE	% fail	RMSE	% fail
0.0	0.168	0.0	0.024	0.0	0.029	0.0
0.5	0.128	3.3	0.035	0.0	0.035	0.0
1.0	0.218	6.7	0.045	0.0	0.044	0.0
1.5	0.294	0.0	0.077	0.0	0.072	0.0
2.0	0.360	3.3	0.126	0.0	0.153	0.0
2.5	0.581	10.0	0.327	3.3	0.202	0.0
3.0	0.783	13.3	—	26.7	0.519	10.0
3.5	1.111	13.3	—	63.3	0.678	10.0
4.0	—	50.0	—	100.0	—	6.7