Radiometric calibration of digital cameras using Gaussian processes

Martin Schall, Michael Grunwald, Georg Umlauf, Matthias O. Franz

Institute for Optical Systems, University of Applied Sciences Konstanz, Brauneggerstr. 55, 78462 Konstanz, Germany

ABSTRACT

Digital cameras are subject to physical, electronic and optic effects that result in errors and noise in the image. These effects include for example a temperature dependent dark current, read noise, optical vignetting or different sensitivities of individual pixels. The task of a radiometric calibration is to reduce these errors in the image and thus improve the quality of the overall application. In this work we present an algorithm for radiometric calibration based on Gaussian processes. Gaussian processes are a regression method widely used in machine learning that is particularly useful in our context. Then Gaussian process regression is used to learn a temperature and exposure time dependent mapping from observed gray-scale values to true light intensities for each pixel. Regression models based on the characteristics of single pixels suffer from excessively high runtime and thus are unsuitable for many practical applications. In contrast, a single regression model for an entire image with high spatial resolution leads to a low quality radiometric calibration, which also limits its practical use. The proposed algorithm is predicated on a partitioning of the pixels such that each pixel partition can be represented by one single regression model without quality loss. Partitioning is done by extracting features from the characteristic of each pixel and using them for lexicographic sorting. Splitting the sorted data into partitions with equal size yields the final partitions, each of which is represented by the partition centers. An individual Gaussian process regression and model selection is done for each partition. Calibration is performed by interpolating the gray-scale value of each pixel with the regression model of the respective partition. The experimental comparison of the proposed approach to classical flat field calibration shows a consistently higher reconstruction quality for the same overall number of calibration frames.

Keywords: Radiometric calibration, flat field correction, digital imaging, Gaussian processes, dark current, bias value, vignetting, fixed pattern noise, photon noise, read noise, camera characteristic

1. INTRODUCTION

Images taken from digital cameras with charge-coupled devices (CCD) or complementary metal-oxide-semiconductor (CMOS) sensors are subjected to a large variety of error sources. Most errors are created by the sensor noise which has four fundamental sources *photon* or *shot noise*, *Fano noise*, *fixed pattern noise* and *read noise*. Photon shot noise and Fano noise are caused by photon interaction with a semiconductor. Fixed pattern noise is non-uniform pixel-to-pixel sensitivity. This noise does not vary between images, but leads to a consistent pattern. Read noise summarizes all noise sources that are independent of signal strength.[1] Further sources of errors are related to optical vignetting. Vignetting refers to brightness decrease starting from the image center[2] and is influenced by focal distance and aperture. Errors and noises include a bias value, dark current, shot noise, read noise, non-linear camera characteristics and different sensitivities of individual pixels. The described errors are mostly dependent on the brightness of the observed scene, exposure time and sensor temperature while others are constant for any given pixel. While the noise is random, error types like the bias value, dark current or a non-linear characteristics are determined by the illumination of the observed scene, exposure time and sensor temperature.[3]

Further author information: (Send correspondence to Michael Grunwald)
Martin Schall: E-mail: Martin.Schall@htwg-konstanz.de, Telephone: +49 (0) 7531 206 273
Michael Grunwald: E-mail: M.Grunwald@htwg-konstanz.de, Telephone: +49 (0) 7531 206 498
Matthias O. Franz: E-mail: mfranz@htwg-konstanz.de, Telephone: +49 (0) 7531 206 651
Georg Umlauf: E-mail: umlauf@htwg-konstanz.de, Telephone: +49 (0) 7531 206 702

Radiometric calibration quantifies these errors and noises. This is done individually for each camera setup in order to reduce influence on the images. Fixed pattern noise, varied sensitivities of individual pixel and vignetting can be reduced using radiometric calibration of the camera.

The classic approach to these problems is dark frame and flat field correction. Bias value and dark current are pixel-wise fixed offsets. The bias value is constant, whereas for a stable temperature each pixel's dark current is constant. Removing this pixel-wise fixed offset is done by subtracting the dark frame. Dark frames cover both dark current and bias value. Only the bias values are measured in a bias frame. To correct transmission variation – depending on optical vignetting or different sensitivities of each pixel – flat field correction is used by applying a multiplicative factor to individual pixels. This requires a flat field image of a uniformly illuminated surface.[4]

Flat field correction builds on a linear model and thus does not reflect the non-linearity of the dark current. Another source of non-linearity in camera characteristics is the full well capacity. Non-linear dependency of the dark current on sensor temperature does not allow linear interpolation in temperature. The proposed approach utilizes a regression method which adapts to the non-linearities.

Regression models can be used for radiometric calibration by estimating the correct intensities of the image pixels based on the observed intensities and known camera parameters. Constant camera parameters – focal distance and aperture – are not included in our regression model. The used regression model approximates a function $f : \mathbb{R}^3 \to \mathbb{R}$ with the independent variables being the observed intensity of the pixel, exposure time and sensor temperature. The dependent variable is the correct intensity of the pixel. The main part of this work is concerned with building one or more regression models for the pixels of a given digital camera setup in order to produce a good estimate of the pixel intensities but still stay within reasonable runtime constraints.

This paper uses a non-parametric approach utilizing machine learning techniques to model arbitrary pixel characteristics dependent on exposure time and sensor temperature. The learning approach infers pixel characteristics from training examples. Regression was done using Gaussian processes to implement a radiometric calibration. A Gaussian process is a machine learning method which can be used for extrapolation, regression and interpolation. In order to do this, an input data set with known light intensities is used to train the process. Based on learned information, the Gaussian process can give a prediction for the relation of unseen data. The unseen data has to be generated by the same system as the known data in order to achieve a suitable prediction. An advantage of the Gaussian process is the confidence interval, that is provided for each predicted value, which gives information about the range within which the value can vary.

Gaussian processes were chosen for regression because they adapt well to non-linear functions with added noise, as described by Rasmussen and Williams.[5] This is made possible by the ability to use an arbitrary number of base functions in weight-space or alternatively a covariance function in function-space. In addition, Gaussian processes include only a small number of hyper-parameters and allow using gradient descent to optimize these.

Section 2 describes radiometric calibration with Gaussian processes. Section 3 describes the hardware setup for generating evaluation data. Results are presented in Section 4, conclusions and outlook given in Section 5.

2. RADIOMETRIC CALIBRATION WITH GAUSSIAN PROCESSES

2.1 Proposed approach

To build a regression model, Gaussian processes need a set of sample observations from the function to be approximated. The images used are gray-scale with one channel. These were created using an integrating sphere as a uniform light source with the camera setup attached. Also attached to the integrating sphere was a LED cold-light source and a photometer. This setup was automatically controllable and measurements were automated. Observations for the regression model were then created by sampling the parameter space consisting of light intensity, exposure time and sensor temperature in uniform steps. The photometer measured the real light intensity within the integrating sphere while the observed light intensities were provided by the camera itself. Also temperature control and measurement was included in the used camera hardware.

Results of previous approaches showed that neither one regression model per pixel, nor one global model for all pixels lead to practical solutions. These approaches need either a high computational effort or produce

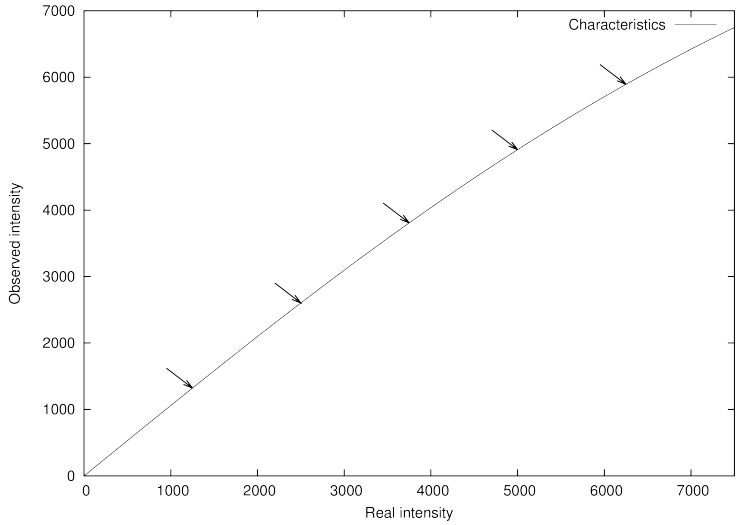

Figure 1. Sampling the pixel characteristic at five different intensities. The characteristic curve was generated by first sorting the training data frame-wise by ascending real intensity.

restored images with a high *Mean Squared Error* (MSE). Simple approaches show conflicting properties of either low runtime or low MSE in the restored images.

The proposed solution partitions all pixels such way that pixels with equal or very similar behavior (characteristic and errors) are in the same partition. The partitioning is done using no spatial information. One suitable representative pixel is chosen per pixel partition. The final step is to compute one regression model for each representative pixel and using this regression model for restoring all the pixels within the associated partition. This approach reduces the computational effort by reducing the number of necessary regression models while maintaining the quality of the restored images. The following four steps of the proposed solution are described in the following sections:

1. Partitioning of all pixels.
2. Choosing one representative pixel per partition.
3. Computing the regression model for each representative pixel.
4. Image restoration using these regression models.

2.1.1 Partitioning of pixels

In order to do a radiometric calibration with only a few regression models, the pixels are partitioned in a way that pixels with similar behavior in their error rates and characteristic are clustered into the same partition. Partitioning the pixels is done by ordering all pixels in such way that pixels with the same or similar characteristic are neighbors. The resulting ordered list is then split into partitions with maximal one pixel difference in size. This way pixels with similar characteristic end up within the same partition.

Sorting of the pixels is done by placing them into a one-dimensional list. Table 1 illustrates this. As the sorting criteria the characteristic of each pixel is sampled at five equally spaced points, as shown in Figure 1. In the best case these five samples per pixel are only dependent on the predictable errors like the bias value, but in reality also random noise influences these observed values. To counter this effect, not the observed intensities at the arrow tips alone are used as sorting criteria, but the mean observed intensity over the neighboring ± 2 intensities in the characteristic. The list of all pixels is then sorted lexicographical by these five sampled observed intensities per pixel. Lexicographic sorting is done with the lowest intensity as the first criteria and then ascending

to accommodate for the fact that the non-linearity of the characteristic has lower impact on smaller intensities. Table 2 shows the ordered example pixel list.

Table 1. Example list of pixels as input for the partitioning.

Coordinates		**Sorting criteria**				
X	**Y**	**1**	**2**	**3**	**4**	**5**
0	0	100	500	1200	1800	3300
0	1	98	499	1200	1802	3307
0	2	98	499	1199	1803	3305
0	3	100	500	1201	1802	3305
0	4	101	503	1208	1801	3303
0	5	99	500	1201	1803	3306

Table 2. Ordered list of pixels with similar pixels neighboring. Spatial coordinates are ignored for sorting.

Coordinates		**Sorting criteria**				
X	**Y**	**1**	**2**	**3**	**4**	**5**
0	2	98	499	1199	1803	3305
0	1	98	499	1200	1802	3307
0	5	99	500	1201	1803	3306
0	0	100	500	1200	1800	3300
0	3	100	500	1201	1802	3305
0	4	101	503	1208	1801	3303

Using only a limited number of frames from the training data in this step reduces the time necessary for loading them into memory since this is time consuming for large quantities of frames with high spatial resolution. Sorting the frames in the training data is done while only loading the meta-data from the header of each file, stored in FITS[6] format.

Final partitioning of the pixels is done by splitting the ordered list of all pixels into partitions with maximal one pixel difference in size. For the experiments, 4000 pixel partitions were generated for a digital camera with 4008×2672 pixels spatial resolution. Thus, each of the 4000 partitions contained 2677 or 2678 pixels. Table 3 shows the partitioning of the six example pixels into two partitions.

Table 3. Partitioned list of pixels as the final partitioning output.

Coordinates		**Sorting criteria**				
X	**Y**	**1**	**2**	**3**	**4**	**5**
0	2	98	499	1199	1803	3305
0	1	98	499	1200	1802	3307
0	5	99	500	1201	1803	3306
0	0	100	500	1200	1800	3300
0	3	100	500	1201	1802	3305
0	4	101	503	1208	1801	3303

Partitioning of the pixels is improved by a preprocessing step in which a suitable bias frame is subtracted from the frames before generating the five sorting criteria. This eliminates most of the bias value and thus reduces the difference in the sorting criteria to the influence of the other error types.

2.1.2 Choosing representative pixels

In order to apply this information to radiometric calibration with Gaussian processes, one representative pixel is chosen per partition. Such that its values from the training data frames are used as input for a regression model suitable to be used for all pixels in the partition while minimizing the resulting error.

The choice of the representative pixel is based on the mean intensity and the intensity variance of each pixel in the partition. The mean intensities of the pixels of one partition give a rough estimate for the different pixel characteristics within it which is sufficient because pixels within one partition are already known to have similar characteristics. Intensity variance is used as a second criterion to ensure that no noisy pixel is chosen as a representative. Also the inclusion of variance prevents pixels hit by cosmic rays in the training data to be chosen as representative pixels. The following equations use μ_i for the mean intensity and σ_i^2 for its intensity variance per pixel. Since calculating the mean and variance are very basic operations with low memory usage, those values are obtained quickly for all pixels while using all available frames from the training data.

The optimal representative pixel minimizes the maximal distance in both mean and variance between itself and all other pixels within the same partition. This suggests that the representative pixel can be artificially generated by calculating the mean over those two values according to $\mu_{Part} = \frac{1}{n}\sum_i^n \mu_i$ and $\sigma_{Part}^2 = \frac{1}{n}\sum_i^n \sigma_i^2$ for n pixels. One actual existing pixel is selected as the representative. This way the necessary training data for the regression model is extracted from one specific pixel within the frames without the need for more computations.

Selecting the representative pixel is done as close as possible to the optimal representative. The optimal representative is the one with mean intensity μ_{Part} and variance σ_{Part}^2. Similarity of the pixels is determined using the Euclidean distance between this optimal representative and the pixels within the partition. The pixel with the minimal Euclidean distance to the optimal representative is chosen as the actual representative pixel for the partition. One representative pixel is selected for each of the previously generated pixel partitions. $repr_{Part}$ contains the index of the pixel to be used as the representative,

$$repr_{Part} = \arg\min_i \sqrt{(\mu_i - \mu_{Part})^2 + (\sigma_i^2 - \sigma_{Part}^2)^2}. \tag{1}$$

2.1.3 Computing regression models

For each of the selected representative pixels, we compute a separate regression model that maps the three-dimensional input $\mathbf{x} \in \mathbb{R}^3$ (observed intensity, exposure time and sensor temperature) to the true intensity value y. The regression algorithm needs training data $D = \{(\mathbf{x}_1, y_1), \ldots, (\mathbf{x}_N, y_N)\}$ consisting of N corresponding input-output pairs. These data are taken from flat fields recorded at varying input intensities, exposure times and sensor temperatures as inputs together with corresponding photometric measurements as outputs (see Section 3). The number of training data points varies from 100 to 250 in our experiments. For all regression models, the independent variables are normalized to a range between zero and one. Similar to the partitioning step, a suitable bias frame is subtracted from all frames in the training data to eliminate the bias value. This is done for both computing the regression model and the later restoring of images.

The pixel characteristic is estimated by Gaussian process (GP) regression which works similar to standard least squares by fitting a nonlinear function consisting of a linear combination of Gaussian radial basis functions (RBFs) centered at the training inputs

$$f(\mathbf{x}) = \sum_{i=1}^{N} w_i e^{-\alpha(\mathbf{x_i}-\mathbf{x})^2} \tag{2}$$

to the input-output pairs in the training data. In GP regression, the functional form of the basis functions is referred to as *covariance function*. Our choice of Gaussian RBFs as covariance function is motivated by its universal approximation capability which allows for modeling arbitrary nonlinear pixel characteristics by an appropriate choice of the regression weights w_i. In contrast to least squares, GP regression is a regularized estimation method, i.e. it automatically prefers smooth regression functions over non-smooth ones by assuming a Gaussian prior on the regression weights w_i with zero mean which leads to a tendency of choosing small regression weights. The resulting smooth estimates for the pixel characteristics have the advantage of being less

sensitive to noise and outliers as compared to least squares estimates. More details about the fitting process in GP regression and the choice of other covariance function can be found in Rasmussen and Williams.[5]

The degree of smoothness of the solutions and the parameter α in the RBFs are so-called *hyperparameters* in GP regression since the values of these parameters cannot be inferred from the regression process alone. The hyperparameters have to be found in an additional model selection step. Here, we use gradient descent on the log likelihood of the found regression models. The log likelihood can be interpreted as the probability that the model is correct given the training data, so the model selection process chooses the most likely hyperparameters that lead to the observed training set. Again, we refer to Rasmussen and Williams[5] for more details on model selection. As gradient descent method, we use a variant of conjugate gradients with a limit of maximally 25 iterations.

After the hyper-parameters are selected and the necessary sample observations known, the Gaussian processes regression models are ready to use for radiometric calibration.

2.1.4 Image restoration

For image restoration, a digital raw image is given and the radiometric calibration is done by estimating the correct intensity for each pixel using the Gaussian processes regression model of the containing partition. As detailed before, the regression models approximates a function $f : \mathbb{R}^3 \rightarrow \mathbb{R}$ mapping the observed intensity, exposure time and sensor temperature to the correct intensity. Exposure time and sensor temperature are given globally for the whole image, the observed intensity is individual for each pixel. During radiometric calibration, this estimation of the correct intensity is done for each pixel in the image. This results in individually estimated intensities for the pixels even while using only one regression model per pixel partition.

2.2 Runtime of the proposed algorithm

The runtime behavior of the proposed algorithm is of interest. For this purpose, an analysis of the algorithm has been done under the following assumptions:

- Runtime for arithmetic operations on scalar variables is constant.
- Runtime for reading an image file is linear to its spatial resolution.
- Runtime for reading meta-data from an image file is constant.
- Runtime for comparison and swapping of two list elements is constant.
- Runtime for calculating a matrix inverse is $O(n^3)$ for a matrix of size $n \times n$.

The following variables were used for the analysis:

- n_{Train} for the number of frames in the training data.
- n_{Pix} for the number of pixels within one image.
- n_{Part} for the number of pixel partitions that will be used.
- $n_{Criteria}$ for the number of sort criteria. The number of frames for smoothing these is constant ± 2.

Analysis of the algorithm can then be broken down into smaller steps as detailed in Table 4. The runtime analysis of the smaller steps of the algorithm is included in the table. The overall runtime analysis is obtained by summing up the steps since they are executed in sequential order.

Overall runtime behavior of the proposed algorithm is $O(n_{Train} \times log(n_{Train}) + n_{Pix} \times log(n_{Pix}) + n_{Pix} \times n_{Train} + n_{Part} \times n_{Train}^3)$ which reduces to $O(n_{Pix} \times log(n_{Pix}))$ under the assumption that the number of frames in the training data n_{Train} and the number of used pixel partitions n_{Part} are constant. These simplifications take into consideration that $n_{Criteria}$ is always less or equal than n_{Train}. Gradient descent was not included into

Table 4. Algorithm broken down into steps with runtime analysis.

Step	**Runtime**
Sorting the training data	$O(n_{Train} \times log(n_{Train}))$
Extracting pixel sort criteria	$O(n_{Pix} \times n_{Criteria})$
Sorting the pixel list	$O(n_{Pix} \times log(n_{Pix}))$
Splitting the pixel list	$O(n_{Pix})$
Choosing representative pixels	$O(n_{Pix} \times n_{Train} + n_{Part})$
Computing regression models	$O(n_{Part} \times n_{Train}^3)$

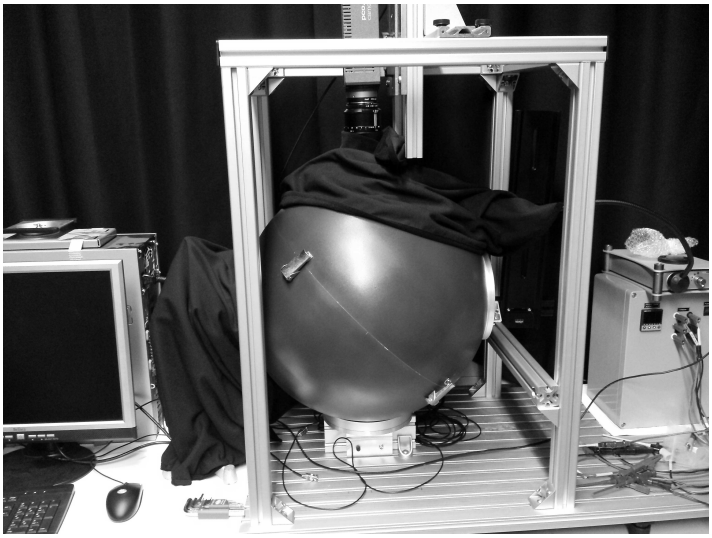

Figure 2. Hardware setup consisting of integrating sphere, camera, light source and photometer.

this analysis because both the maximum number of iterations and the number of frames used are limited, thus the gradient descent is enforced to an upper runtime boundary.

That using n_{Train} and n_{Part} as constant is a valid simplification is shown in Section 5. The only unbound variable remaining in the analysis is the number of pixels in the digital images. It is also to be expected that this number will increase in the future. Runtime behavior of the algorithm is only mixed logarithmic and linear in dependency of this size.

Runtime for calculating a matrix inverse for a matrix of size $n \times n$ is assumed as $O(n^3)$ in this analysis. An algorithm with $O(n^{2.3727})$ runtime has been published by Williams.[7]

3. HARDWARE SETUP

The hardware setup used for obtaining the images for comparison consisted of an integrating sphere — 60 centimeter diameter and coating for usage in the visible light spectrum — with an attach LED cold-light source providing 6500 kelvin color temperature and up to 725 lumen luminous flux. Light intensity within the sphere was measured using an i1Display Pro photometer for reference. The used digital camera system, a PCO4000s* with 4008×2672 spatial resolution, 14 bit dynamic range and a cooled CCD sensor was mounted on the top of the sphere. A Carl Zeiss Makro-Planar T* 2/100† optics with $97,5mm$ focal length and $f/2,0 - f/22$ aperture

*`http://www.pco.de/fileadmin/user_upload/pco-product_sheets/BR_pco_1600_2000_4000_102_online.pdf`

†`http://www.zeiss.com/content/dam/Photography/new/pdf/en/downloadcenter/datasheets_slr/makroplanart2100.pdf`

range was utilized. Figure 2 shows the overall setup which was automatically controllable to allow for producing large series of images with different parameters.

The optics were used with focal distance and aperture settings that did not produce significant optical vignetting and imaged the interior of the integrating sphere in a blurry way. Blurry imaging was of use here because the integrating sphere used was uneven and a blurry imaging prevented localized strong deviations from the reference light intensity. Black cloth were used to cover the hardware setup in order to prevent scattered light from influencing the automated imaging (Figure 2).

4. RESULTS

4.1 Procedure

Part of this work is a comparison of the proposed algorithm to flat field correction. The implementation of the flat field correction used does reduce the noise by calculating the mean over several frames. Dark current is removed by subtracting a suitable dark frame. Multiplication of the pixel intensities with individual coefficients, calculated out of a flat field, does counter optical vignetting, fixed pattern noise and individual sensitivities.

The flat field correction needs several frames, obtained from the actual camera setup, for radiometric calibration. These frames include dark frames and flat fields. Flat fields are independent of exposure time and temperature since they will be corrected by subtraction of a mean dark frame. Dark frames on the other hand need to be parameterized with exposure times and sensor temperatures from the same range that will later be used for images to be restored. Dark frames can be interpolated in the exposure time but not in sensor temperature. Other parameters like the focal distance and aperture were kept constant during the comparisons.

The goal of the comparisons was to find out how the MSE behaves in both the flat field correction and the proposed algorithm based on Gaussian processes when the number of available reference frames is varied. Frames for flat field correction were obtained from the hardware setup by first generating a fixed number of flat fields and bias frames, followed by evenly spacing the remaining number of dark frames in a grid over both the exposure time and sensor temperature ranges. This allows interpolation of missing dark frames. Training data for the Gaussian processes were generated by taking reference frames with exposure time, sensor temperature and light intensity evenly spaced over all three parameters. For the Gaussian processes, only one frame was obtained for each parameter set, thus allowing for a denser grid while sampling the whole parameter space.

To allow for a comparison, each test also included a set of digital raw images obtained using an integrating sphere. Frames from this integrating sphere were assumed to be uniform in light intensity per pixel and thus the attached photometer could be used as the error free reference value (see Section 3 for details).

Several tests were done during this comparison, each of them repeated using a variable exposure time, sensor temperature or both. This resulted in four different types of tests. To determine the behavior of both algorithms using an increasing number of calibration frames, tests were done using cropped images with 500×500 pixels spatial resolution. Tests with full ten megapixels spatial resolution, but a constant number of 250 calibration frames, provide insight into the remaining errors of the algorithms for realistic applications. Measures for comparing both algorithms is the MSE between the restored image and the exact image.

4.2 Results using increasing numbers of calibration frames

Comparison of the behavior of both radiometric calibration algorithms with a variable number of calibration frames was done by obtaining sets of images with 500×500 pixel spatial resolution. These sets contained 100 to 250 calibration frames in steps of ten frames. The set of 100 raw images for restoration and comparison was fixed during the whole test series. This was repeated for four different test series with variable exposure time, sensor temperature or both. Figure 3 shows the results for the series with constant sensor temperature and variable exposure time, a common use-case for digital cameras in scientific environments.

The results show a decrease in MSE in the images that were restored with the proposed algorithms between 100 and 160 available calibration frames. MSE remains within a small range above 220 calibration frames. The use of more than 130 calibration frames resulted in lower MSE of the proposed algorithm in comparison to the flat field correction. These observations were reproduced in the other test series or with a lower MSE in the proposed algorithm with more than 100 calibration frames.

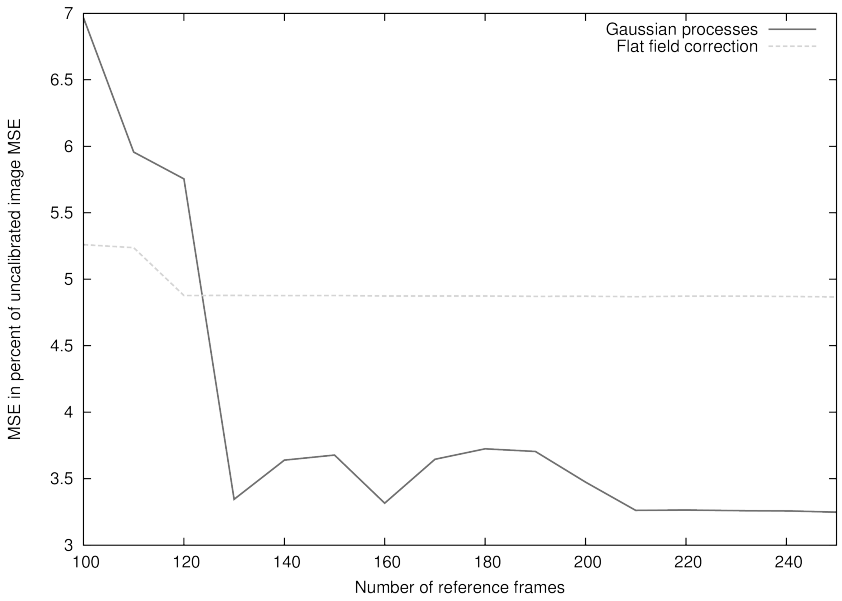

Figure 3. MSE for the test series with constant sensor temperature and variable exposure time.

4.3 Results using full-resolution images

Comparisons on full resolution (4008×2672 pixel) images but with a constant number of 250 calibration frames were done using variable exposure time, sensor temperature or both. The results of these tests are shown in Table 5.

Table 5. Results using full spatial resolution images. Comparing flat field correction (FFC) and the calibration using Gaussian processes (GP). MSE is in percent of the MSE before restoration.

Parameters & results	**Series 1**	**Series 2**	**Series 3**	**Series 4**
Temperature	$0°C$	$0°C$	$0°C - 2°C$	$0°C - 2°C$
Exposure time	$100ms$	$25ms - 150ms$	$100ms$	$25ms - 150ms$
FFC MSE	$4,28\%$	$3,73\%$	$4,54\%$	$4,26\%$
GP MSE	$3,12\%$	$3,12\%$	$3,11\%$	$3,56\%$
GP runtime regression	$990s$	$990s$	$1018s$	$1071s$
GP runtime per restoration	$121s$	$116s$	$120s$	$120s$

The same results as with low resolution images are observed in this case. Calibration using Gaussian processes did produce images with lower MSE than images restored with flat field correction. Restoring ten megapixel images did take around 120 seconds per image.

5. CONCLUSIONS

Results in Sections 4.2 and 4.3 show that the radiometric calibration with Gaussian processes produces a consistently smaller MSE in restored images when compared to the flat field correction. MSE decreases steadily for calibrations using between 100 and 200 training images. Above 200 training images, no further decrease in MSE could be achieved. Runtime analysis and measurement shows that application of the proposed algorithm to full

resolution images using 250 calibration frames is possible for offline use-cases. Measured runtimes of 120 seconds per image disqualifies the algorithm for real-time online applications on standard hardware.

The test series using full resolution images, in Section 4.3, include four different use-cases with variable exposure time, sensor temperature or both. All four test series show a lower remaining MSE in the images restored with Gaussian processes in comparison to the flat field correction. This leads to the conclusion that a calibration with Gaussian processes is more reliable in interpolating error rates for unknown values of the exposure time or sensor temperature. This makes the proposed algorithm flexible in applications.

The proposed algorithm partitions the pixels by splitting the ordered list of pixels into slices of equal size. The size of each partition is dependent on the spatial resolution and the number of partitions generated. With higher spatial resolution, the number of pixels per partition will increase. Since the order of the pixel list is determined by the sizes of the errors and not the spatial resolution, even with more pixels per partition the individual pixels are expected to retain or decrease their difference to the representative pixel of the partition.

This will lead to better results of the radiometric calibration with new hardware generations even while spatial resolution increases and the number of frames in the training data and number of partitions generated are kept constant.

The resulting restored images using the proposed algorithms rely on the partitioning of the pixels and selection of a representative pixel per partition. Partitioning splits the list of pixels into parts of equal size. It could be possible that some error sources (e.g. pixel sensitivities) do not have a continuous distribution over the whole value range. Assuming such a distribution, partitioning could be improved by splitting the pixel list at positions showing large differences in the observed values. This would lead to partitions with different sizes but higher correlation of the included pixels.

6. ACKNOWLEDGMENT

We thank Pascal Laube for reading the draft of this work and providing valuable input.

REFERENCES

[1] Janesick, J. R., [*Photon Transfer DN to lambda*], SPIE (2007).

[2] Hecht, E., [*Optik*], Addison-Wesley, 5th ed. (2009).

[3] Dierks, F., "Sensitivity and image quality of digital cameras," tech. rep., Basler Vision Technologies (2004).

[4] Gendler, R., [*Lessons from the Masters: Current Concepts in Astronomical Image Processing*], Springer (2013).

[5] Rasmussen, C. E. and Williams, C. K. I., [*Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*], The MIT Press (2006).

[6] Pence, W. D., Chiappetti, L., Page, C. G., Shaw, R., and Stobie, E., "Definition of the flexible image transport system (fits), version 3.0," *Astronomy & Astrophysics* **524** (2010).

[7] Williams, V. V., "Breaking the coppersmith-winograd barrier," tech. rep., UC Berkeley and Stanford University (2011).