Optical Engineering

[SPIDigitalLibrary.org/oe](https://spiedigitallibrary.org/oe)

Resampling in hyperspectral cameras as an alternative to correcting keystone in hardware, with focus on benefits for optical design and data quality

Andrei Fridman
Gudrun Høye
Trond Løke

Resampling in hyperspectral cameras as an alternative to correcting keystone in hardware, with focus on benefits for optical design and data quality

Andrei Fridman,* Gudrun Høy, and Trond Løke

Norsk Elektro Optikk, P.O. Box 384, Lørenskog 1471, Norway

Abstract. Current high-resolution hyperspectral cameras attempt to correct misregistration errors in hardware. This severely limits other specifications of the hyperspectral camera, such as spatial resolution and light gathering capacity. If resampling is used to correct keystone in software instead of in hardware, then these stringent requirements could be lifted. Preliminary designs show that a resampling camera should be able to resolve at least 3000–5000 pixels, while at the same time collecting up to four times more light than the majority of current high spatial resolution cameras. A virtual camera software, specifically developed for this purpose, was used to compare the performance of resampling and hardware corrected cameras. Different criteria are suggested for quantifying the camera performance. The simulations showed that the performance of a resampling camera is comparable to that of a hardware corrected camera with 0.1 pixel residual keystone, and that the use of a more advanced resampling method than the commonly used linear interpolation, such as high-resolution cubic splines, is highly beneficial for the data quality of the resampled image. Our findings suggest that if high-resolution sensors are available, it would be better to use resampling instead of trying to correct keystone in hardware. © The Authors. Published by SPIE under a Creative Commons Attribution 3.0 Unported License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI. [DOI: [10.1117/1.OE.53.5.053107](https://doi.org/10.1117/1.OE.53.5.053107)]

Keywords: hyperspectral; imaging; resampling; high-resolution cubic splines; co-registration; misregistration; keystone; virtual camera.

Paper 131941P received Dec. 27, 2013; revised manuscript received Mar. 27, 2014; accepted for publication Apr. 1, 2014; published online May 7, 2014.

1 Introduction

Hyperspectral cameras, also called imaging spectrometers, are used in various fields, such as geology, military, forensics, food industry, and so on. The key feature of these instruments is their ability to acquire two dimensional (2-D) images where each pixel contains spectral information of the corresponding small area of the depicted scene. In order to ensure accuracy of the spectral data, these cameras must have good spatial and spectral co-registrations, i.e., the signal for each spectral band should be collected from the same scene area for every spatial pixel (spatial co-registration), and for any spatial pixel inside the field of view every specific spectral channel N should correspond to the same range of wavelengths (spectral co-registration).¹ A typical example of spectral misregistration would be “smile”: a change of the central wavelength of a specific spectral channel as a function of position in the field of view. A typical example of spatial misregistration would be “keystone”: a change in the position of the same spatial pixel on the scene as a function of wavelength.

The current approach for achieving good co-registration in high-end cameras is to do the necessary corrections as well as possible in hardware (HW corrected cameras). Resampling is sometimes used to further reduce errors in the data captured by such cameras.² However, there are cameras where correcting smile and keystone in hardware is not an option, and where resampling has to be used instead in order to get sensible data. In the case of push-broom

cameras, hardware correction of smile and keystone is normally set up as a requirement, because cameras with large smile and keystone that require resampling are believed to give inferior data quality. We are going to examine the validity of this assumption.

In cases where resampling is considered at all, the most common method seems to be linear resampling. This method gives better results than the nearest neighbor method, and also seems to be intuitively “right” or “physically correct.” However, several more advanced resampling methods have been developed and are being used in traditional image processing.³ We are going to show the advantage of using a more advanced interpolating method, also in hyperspectral imaging.

It is often assumed that the resampling reduces the spatial resolution by a factor of 2. Again, this seems to be intuitively correct, since in the case of linear resampling two pixels are used to form one pixel in the resampled image. We are going to examine the quality of resampled data when there is no significant downsampling. Avoiding significant downsampling is especially important for infrared sensors because of their limited spatial resolution.

Very stringent requirements for the optics in terms of smile and keystone correction severely affect other optical specifications, such as the light gathering ability and spatial resolution. Some of the optical aberrations, such as point-spread-function (PSF) center of gravity position and PSF shape, in the current hyperspectral cameras have to be corrected with a small fraction of a pixel precision.¹ It is clear that lifting such stringent requirements would allow a great (i.e., not a few percent but a few times) increase in the spatial

*Address all correspondence to: Andrei Fridman, E-mail: fridman@neo.no

resolution of the optics, and, in most cases, significantly (a few times) increase the amount of light reaching the sensor. Increased light gathering ability reduces the influence of photon and readout noises (or the camera can be operated at higher frame rate than before), and higher spatial resolution (with an appropriate sensor) reduces misregistration errors for large objects and, of course, makes it possible to detect and identify smaller objects. The advantages of the increased light gathering capacity with respect to data quality will be demonstrated in this article.

2 Virtual Camera Simulations

A virtual camera software was developed⁴ in order to evaluate and compare the performance of various types of push-broom cameras, such as resampling cameras, mixel cameras,⁵ and HW corrected cameras. The virtual camera software simulates the performance of a hyperspectral camera and uses the hyperspectral data of a real scene (captured by a real hyperspectral camera) as input. The virtual camera distorts the input data somewhat in accordance with the modeled optical distortions, sensor characteristics, and photon noise. Then, by comparing the data at the input of the virtual camera with the data at the output of the virtual camera, we are able to evaluate the performance of the camera. The virtual camera software models various aspects of camera performance, such as keystone, PSF of the optics, photon and readout noises, and so on.

3 Resampling Methods

The following resampling methods will be evaluated:

1. linear resampling
2. high-resolution cubic splines

Linear resampling is a fast and straightforward way to resample an image (or a line of an image) from one grid to another. The method applies a linear interpolating function and uses the two nearest pixels to calculate the value of the new grid point. This method seems to be intuitively “right” in preserving the data when resampling to a different grid.

Perhaps for this reason, it is often used when there is a need to resample hyperspectral data.

However, in traditional imaging applications, high-resolution cubic splines have been shown to introduce smaller errors than other resampling methods, such as nearest neighbor resampling, linear resampling, and cubic b-splines resampling, when preservation of different spatial frequencies is used as the criterion.³ The high-resolution cubic splines method used in this article utilizes an interpolating function that is applied over the four nearest pixels:³

$$\begin{aligned} f_1(x) &= (a+2)x^3 - (a+3)x^2 + 1 \quad [0, 1] \\ f_2(x) &= ax^3 - 5ax^2 + 8ax - 4a \quad [1, 2]. \end{aligned} \quad (1)$$

The function f_1 is applied over the interval $x = 0$ to $x = 1$, while f_2 is applied over the interval $x = 1$ to $x = 2$. The interpolating function is symmetric about $x = 0$. Frequently used values for the parameter a in the literature are -1 , $-3/4$, and $-1/2$. We have used $a = -3/4$ in our simulations, which ensures that the second derivatives of the two cubic polynomials f_1 and f_2 are equal at $x = 1$.

4 Camera Performance Analyses: One-Dimensional Example Scene

A hyperspectral data set containing 1600 spatial pixels, originally captured using a HySpex VNIR1600 hyperspectral camera,⁶ forms the “continuous” 1-D scene (blue curve in Fig. 1) to be captured by the virtual camera. The virtual camera is set to have significantly lower resolution (320 pixels) than the resolution of the scene, so that five spatial pixels from the HySpex VNIR1600 data set form one scene pixel. By doing this, we simulate the fact that any real scene contains smaller details than the resolution of the camera to be tested (in this case, details as small as $1/5$ of a scene pixel are present).

Figure 1 shows the number of photons in the signal from the scene for one spectral band. The signal contains large areas with slowly changing brightness, relatively sharp borders between such areas, and some quite small objects which are significantly different in intensity compared to the background. This scene will therefore allow us to examine how

Fig. 1 The reference scene consisting of 320 scene pixels. The blue curve shows the photon number density with spatial details as small as $1/5$ of a scene pixel (see the enlarged part of the graph), while the corresponding scene pixel values are shown in red.

the cameras perform on different scene features. Further, the number of pixels is large enough that some conclusions can be drawn based on statistics. A more extensive statistical analysis will be performed in Chap. 5.

As the starting point of our evaluation, we will compare the following four virtual cameras:

1. HW corrected camera, 0.1 pixel keystone (the camera that many people aim for during the design phase)
2. HW corrected camera, 0.3 pixel keystone (which is easier and cheaper to achieve)
3. resampling camera that uses linear resampling
4. resampling camera that uses high-resolution cubic splines.

The HW corrected camera is simulated by shifting the scene pixels to the left relative to the sensor pixels by an amount equal to the maximum residual keystone. This is in a way the worst case scenario, since a real camera will never have so large keystone in every spatial pixel in every spectral band. However, this assumption ensures that we will be able to examine the effect of having maximum residual keystone also in the most difficult areas of the image, where adjacent pixels are significantly different from each other.

For our virtual resampling camera, we assume a keystone of 32 pixels, i.e., the content of the 320 scene pixels is spread over 352 pixels when recorded onto the sensor. A keystone as large as 10% of the image size was chosen for several reasons. First of all, this keystone is definitely large enough to provide much greater flexibility in optical design compared to the HW corrected cameras. Also, resampling produces larger errors when the output pixel is positioned exactly between two input pixels,³ and in case of so large keystone this will occur several times along a single image line. This will make the conclusions drawn from the simulations more reliable. On the other hand, no significant downsampling is performed, i.e., the spatial resolution of the sensor is more or less preserved in the final data. The keystone is assumed to be linear across the image, changing from zero on the left side of the image to 32 pixels on the right side. The recorded pixels are then resampled onto the scene pixel grid (using a linear or cubic splines interpolation) to give the final data.

When evaluating the performance of the cameras, we calculate the error in the final data relative to the input. The relative error, dE , is given by

$$dE = \frac{(E_{\text{final}} - E_{\text{init}})}{E_{\text{init}}}, \quad (2)$$

where E_{init} is the scene pixel value (number of photons) and E_{final} is the calculated value of the same pixel after the signal has been processed by the camera. We can then find the standard deviation of dE over the 320 pixels, and we can also determine the maximum relative error. Both are important parameters when evaluating the performance of the cameras.

We will first compare the misregistration errors for the cameras in Sec. 4.1. Then, in Sec. 4.2, we will include photon noise in the simulations, before moving on to compare the performance of the cameras in low light in Sec. 4.3. Finally, in Sec. 4.4, we will explore the possibility of

downsampling the data for the resampling camera and see how this reduces the errors compared to a HW corrected camera with the same spatial resolution.

4.1 Misregistration Errors

Figure 2 shows the misregistration errors (i.e., when photon noise is not included in the calculations) for the HW corrected and resampling cameras. We see that both HW corrected cameras [Figs. 2(a) and 2(b)] show distinct error peaks in the areas with high local contrast, and that the camera with 0.3 pixel keystone has about three times larger errors than the camera with 0.1 pixel keystone. It is also clear from the figure that linear resampling [Fig. 2(c)] gives larger errors than the use of high-resolution cubic splines [Fig. 2(d)] with standard deviation of the error 3.1% versus 2.8% and maximum error 28.1% versus 18.4%.

Both resampling methods give significantly more precise data than the HW corrected camera with 0.3 pixel keystone, which has standard deviation of the error 5.4% and maximum error 46%. However, the HW corrected camera with 0.1 pixel keystone outperforms both resampling cameras with standard deviation of the error 1.9% and maximum error 14.8%.

So far, it looks like the best option is to use a HW corrected camera with 0.1 pixel keystone. However, if you are not able to build a camera with keystone significantly less than 0.3 pixel, it may be better to skip correcting keystone in hardware and rather concentrate your efforts on correcting it in postprocessing by resampling.

For all the cameras tested here, large brightness variations on pixel and subpixel scales cause large misregistration errors. In a real camera, the signal will be somewhat blurred before being sampled by the sensor. Blur will make the signal smoother, which will reduce resampling errors. Blur will also reduce local contrast, which should reduce misregistration errors in HW corrected cameras. In all further simulations, the signal was blurred by convolving a real PSF of a HySpex VNIR1600 camera with the signal. This PSF corresponds to a modular transfer function of 0.44 at Nyquist frequency. The PSF was first scaled to the scene pixel size used in the simulations. After the optical blur was applied to the signal, the contrast of small details was greatly reduced as shown in Fig. 3 (gray versus black curves). The entire signal with applied optical blur is shown in Fig. 4 together with the corresponding scene pixel values. Please note, that, when calculating the errors for the blurred signal [Eq. (2)], optical blur is applied both to the input and the output signals. This is done because here we do not investigate errors caused by the optical blur. We investigate errors caused by resampling or keystone applied to a blurred signal, with another blurred signal (which is not resampled and is keystone-free) used as the reference.

Let us examine the performance of the four cameras when the input signal is blurred by the optics (Fig. 5). We see that the errors for all cameras are now smaller. The two HW corrected cameras and the camera that uses linear resampling still perform the same relative to each other, but the resampling camera that utilizes the high-resolution cubic splines method now performs better than any of the other three cameras.

So far, resampling with high-resolution cubic splines has given better results than linear resampling. For the remaining

Fig. 2 Misregistration errors for (a) a HW corrected camera with 0.1 pixel keystone, (b) a HW corrected camera with 0.3 pixel keystone, (c) a resampling camera that uses linear resampling and (d) a resampling camera that uses high-resolution cubic splines. The standard deviation of the error is marked by a dashed line. Photon noise is not included.

of this chapter, we will therefore focus on comparing the performance of the two HW corrected cameras to the performance of the resampling camera that uses high-resolution cubic splines. We will, however, get back to linear resampling in Sec. 5, where the camera performance is analyzed based on large amounts of data.

4.2 Errors when Photon Noise is Included

Misregistration errors are not the only errors in a real system – photon and readout noises will further degrade the signal. If we set as a requirement that the signal-to-noise ratio (SNR) should be at least 10, then we would need to have at least 100

Fig. 3 Intensity variations in the scene before (gray) and after (black) blurring in the optics. Note that the gray curve shown here corresponds to the blue curve in Fig. 1, while the black curve shown here corresponds to the blue curve in Fig. 4.

photons per pixel (in the case of 100% quantum efficiency). In this case, the photon noise would be 10 electrons RMS. The best modern sensors have readout noise less than 2 electrons RMS. The contribution from readout noise will therefore be more or less negligible even for the lowest useable light level. For this reason, we have not included readout noise in the calculations. However, if a more noisy sensor is used in a camera, a similar type of analysis can (and should) be done with the readout noise taken into account.

Figure 6 shows the performance for the HW corrected and resampling cameras when photon noise is included in the simulations. The scene in Fig. 4 (blurred signal) was used for the calculations. Error peaks, corresponding to sharp brightness variations, are clearly visible above the photon noise. However, the areas with smaller brightness variations seem to contain larger errors than before (if compared with Fig. 5) due to the presence of photon noise. The errors in those areas could be significantly reduced if the optical system was able to collect more light.

We have looked into the design of a hyperspectral camera that uses resampling, exploring the possibilities that open up when correction of keystone in hardware is no longer necessary. Even the first design attempts show that such a camera should be able to resolve at least 3000–5000 spatial pixels while at the same time collecting up to four times more light compared to the majority of current high-resolution

HW corrected cameras. For example, the optical system shown in Ref. 5 is suitable for a resampling camera if the array of light mixing chambers is replaced with a traditional slit. Having a four times higher signal reduces noticeably the errors in the areas with smaller brightness variations, as shown for the resampling camera in Fig. 7. However, the error peaks (corresponding to sharp brightness variations) remain more or less the same.

The best recent Dyson designs,^{7,8} which correct keystone in hardware, also have very low F-number, i.e., are able to collect a lot of light. However, these designs contain an expensive large concave grating and should have very precise component centration in order to keep misregistration errors reasonably low. They require telecentric foreoptics which should have the same low F-number as the Dyson relay itself. In addition to the required telecentricity and low F-number, the foreoptics needs to have good keystone correction. These three factors may make the design of the foreoptics extremely difficult, and the manufacturing process challenging (i.e., expensive). The design and manufacturing of resampling cameras appear to be easier and cheaper, because the focus of the design is shifted from keystone correction to keystone characterization. Also, with a resampling camera it seems to be possible to have low F-number, very high spatial resolution, and a large field of view at the same time in one single instrument.

Fig. 4 The reference scene when the signal is blurred. The blue curve shows the photon number density (with spatial details as small as 1/5 of a scene pixel), while the corresponding scene pixel values are shown in red.

Fig. 5 Misregistration errors when the signal is blurred for (a) a HW corrected camera with 0.1 pixel keystone, (b) a HW corrected camera with 0.3 pixel keystone, (c) a resampling camera that uses linear resampling and (d) a resampling camera that uses high-resolution cubic splines. The standard deviation of the error is marked by a dashed line. Photon noise is not included.

4.3 Low Light

The advantages of a resampling camera, with its ability to collect considerably more light compared to the majority of existing high-resolution cameras, are clearly visible in low light conditions. In the following example, we have used the scene in Fig. 4 (blurred signal), but reduced the intensity of the signal by a factor of 10.

Figure 8 shows the performance in low light when all cameras collect the same amount of light. Both the HW corrected camera with 0.1 pixel keystone [Fig. 8(a)] and the resampling camera [Fig. 8(c)] seem to be limited by photon noise with standard deviation of the errors close to 7% in both cases. However, the HW corrected camera with 0.3 pixel keystone [Fig. 8(b)] still shows noticeable

Fig. 6 Camera performance when photon noise is included for (a) a HW corrected camera with 0.1 pixel keystone, (b) a HW corrected camera with 0.3 pixel keystone and (c) a resampling camera that uses high-resolution cubic splines. The standard deviation of the error is marked by a dashed line.

misregistration errors, having a maximum error of almost 40% compared to around 20% for the other two cameras.

Figure 9 shows that, as expected, the ability of the resampling camera to collect four times more light significantly improves its performance in low light situations. With standard deviation of the error 3.3% and maximum error 9.6%, the resampling camera that collects four times more light performs considerably better than the HW corrected camera

with 0.1 pixel keystone (standard deviation of the error 6.8%, maximum error 20.3%).

4.4 Minimization of Misregistration Errors in a Resampling Camera by Downsampling

In some applications, the requirements regarding misregistration errors may be more stringent than what can be

Fig. 7 Camera performance when photon noise is included for a resampling camera that uses high-resolution cubic splines and collects four times more light. The standard deviation of the error is marked by a dashed line.

Fig. 8 Camera performance in low light for (a) a HW corrected camera with 0.1 pixel keystone, (b) a HW corrected camera with 0.3 pixel keystone and (c) a resampling camera that uses high-resolution cubic splines. The standard deviation of the error is marked by a dashed line.

achieved with the high-resolution cubic splines resampling method. In such cases, the misregistration errors can be significantly reduced by pixel binning of the resampled data in the spatial direction. Since resampling cameras can be built to have significantly higher spatial resolution than HW corrected cameras, binning or downsampling in the spatial direction may be acceptable.

Figure 10 compares the performance of a resampling camera that applies a downsampling factor of 2 in the spatial direction with a HW corrected camera with the same resolution. The resampling camera (with 352 pixels resolution resampled to 320 pixels and then binned by a factor of 2 to 160 pixels) shows significantly smaller errors [Fig. 10(a)] than the HW corrected camera with 160 pixels

Fig. 9 Camera performance in low light for a resampling camera that uses high-resolution cubic splines and collects four times more light. The standard deviation of the error is marked by a dashed line.

Fig. 10 Camera performance for (a) a resampling camera that uses high-resolution cubic splines (with 352 pixels resolution resampled to 320 pixels and then binned by a factor of 2 to 160 pixels), (b) a HW corrected camera with 0.1 pixel keystone and 160 pixels resolution and (c) a resampling camera that uses high-resolution cubic splines and collects four times more light (with 352 pixels resolution resampled to 320 pixels and then binned by a factor of 2 to 160 pixels). The standard deviation of the error is marked by a dashed line. Photon noise is included.

resolution and keystone 0.1 pixel [Fig. 10(b)]. In addition to much lower standard deviation of the error (1.4% for the resampling camera versus 2.9% for the HW corrected camera), the resampling camera has very low maximum error (4.3% versus 12.4%). Also, the error peaks in the areas of high local contrast are barely visible in the case of the resampling camera. These results indicate that a resampling camera with large keystone that uses a high spatial resolution sensor, should be able to deliver data of significantly higher quality than a HW corrected camera with a residual keystone of 0.1 pixel that uses a lower spatial resolution sensor. Indeed, if the high resolution raw data from the resampling camera are downsampled to match the low spatial resolution of the HW corrected camera, the resampled data will have the same spatial resolution but lower misregistration errors. Alternatively, if the high resolution raw data from the resampling camera are resampled without any significant downsampling, the misregistration errors per pixel of the resampling camera will be comparable to the errors of a

good HW corrected camera, but there will be much more spatial information in the resampled data.

Of course, a resampling camera would also be able to benefit from the ability of the optics to collect significantly more light. Then, the advantages of a resampling camera compared to a HW corrected camera will be even more noticeable; the standard deviation of the error in this case is reduced to 0.84% and the maximum error to 2.9% [Fig. 10(c)].

If it is possible to collect even more light from the scene, the misregistration errors of the resampling camera with downsampling factor approximately 2 will eventually at some point become visible above the photon noise. In such cases, an even larger downsampling factor could be used to lower misregistration errors further.⁴

5 Camera Performance Analyses Based on Large Amount of Data

Analyses of a single 1-D signal show in a very intuitive way the advantages and disadvantages of different approaches to

building cameras, since the connection between the scene features and the errors is apparent in this case. In order to verify some of the findings of the previous chapter, the performance of a HW corrected camera has also been compared to the performance of a resampling camera by using a large amount of data as input. We used the following approach. A 2-D scene of $1600 \times 12,233$ pixels (Fig. 11) was scanned by a virtual HW corrected camera and a virtual resampling camera. Instead of one 1-D signal with 1600 different values (as in the previous chapter), we now have 12,233 such 1-D signals to draw conclusions from.

Both tested cameras should produce a spatial resolution of 320 pixels as before. The calculations will be done for the 697 nm wavelength. The virtual resampling camera has 32 pixels keystone at the tested wavelength, so that it produces 352 pixels output which should then be resampled to 320 pixels in the final data. This is similar to before.

The HW corrected camera will be modeled differently now. Instead of simulating keystone by shifting the entire row of pixels on the sensor by a fraction of a pixel, we will now say that there is no sensor shift, and that pixel #1 is positioned perfectly. However, due to keystone, the image at the tested wavelength is linearly expanded by 0.3 pixels, so that the entire image is projected onto 320.3 pixels instead of 320 pixels. The keystone in pixel #1 is then nearly 0, in pixel #160 it increases to 0.15 of a pixel, and in pixel #320 the keystone reaches 0.3 pixels. Simulating such a keystone distribution makes it possible to take advantage of the large amount of data available and to check the resulting errors for different keystone values.

5.1 Misregistration Errors

We will first look at the misregistration errors. Figure 12(a) shows the standard deviation of the misregistration errors observed in each of the 320 pixels of the output for a HW corrected camera (blue curve) and two resampling cameras that use high-resolution cubic splines (green curve) and linear resampling (red curve), respectively. The standard deviation for each pixel is calculated over the 12,233 pixels that have the same position on the sensor as that particular pixel of interest (for instance, pixel #53 represents one such position and corresponds to a keystone of 0.05 pixels for the HW corrected camera).

Figure 12(a) shows that the standard deviation of the misregistration errors for the HW corrected camera increases almost linearly as the keystone increases from 0 to 0.3 pixels from left to right in the figure. For the resampling cameras, the errors vary periodically as a function of pixel number. The periodic variations occur because resampling gives the smallest errors when the positions of the resampled pixel and the corresponding recorded pixel are nearly

identical, while the largest errors occur when a pixel of the resampled image is positioned right between two pixels of the recorded image,³ which here happens 32 times due to the 32 pixels keystone. Comparing the curves for the three cameras, we see that the misregistration errors in the resampled image are equivalent to the misregistration errors in a HW corrected system with 0.19 pixel keystone when linear resampling is used (the blue curve crosses the peaks of the red curve around pixel #200 in the figure), and to a HW corrected system with 0.1 pixel keystone when high-resolution cubic splines are used for the resampling (the blue curve crosses the peaks of the green curve around pixel #110 in the figure). Please note that we compare a HW corrected camera with a resampling camera in the areas of the image where resampling gives the least accurate results.

The standard deviation of the errors gives a good indication of how accurate the spectra of the majority of the pixels will be. However, there are tasks when it is required to capture high quality spectra of a few particular objects in the scene. For example, if we are looking for an object in the forest, we need a reasonably accurate spectrum of that object. If the camera fails to capture that particular spectrum accurately enough, while giving low standard deviation of the error on the rest of the forest, then this camera is not good enough for this particular task. This is the reason why we were monitoring not only the standard deviation of the errors, but also the maximum errors, when testing the cameras with the 1-D signal in the previous chapter. However, when checking the camera performance on large amounts of data (in this case approximately 20 million objects which are depicted by approximately 4 million pixels) there are situations for all cameras when the largest errors are very large. We believe that a discussion whether a camera with 70% maximum error is significantly better than a camera with 100% maximum error (which we have seen examples of during simulations) is not particularly useful. Therefore, when dealing with large amounts of data, we instead suggest setting up a threshold for the maximum acceptable error, and then using the number of pixels with errors above this threshold as a criterion for the camera performance.

Let us say that only pixels with error less than 10% of the signal can be considered useful for further analysis. A maximum acceptable error of 10% is by no means the ultimate criterion, but it seems to be an adequate and practically relevant criterion for high-end scientific hyperspectral imaging systems. Let us take a look at Fig. 12(b). Instead of showing the standard deviation of the errors, as was the case in Fig. 12(a), the vertical axis now shows how many depicted scene pixels (relative to the total number of scene pixels captured by that particular pixel on the sensor) have

Fig. 11 The two-dimensional scene, $1600 \times 12,233$ pixels in size, providing a large amount of data for performance analysis of hyperspectral cameras. The band with 697 nm central wavelength is shown.

Fig. 12 (a) Standard deviation of the misregistration errors and (b) relative number of pixels with misregistration errors larger than 10% of the signal, for a HW corrected camera (blue curve) and two resampling cameras that use high-resolution cubic splines (green curve) and linear resampling (red curve), respectively. For the HW corrected camera, the keystone varies from 0 to 0.3 pixels from the left to the right part on the sensor. The calculations were done for the 697 nm wavelength. Photon noise is not included.

misregistration errors above 10% of the signal. We see that for the resampling cameras the number of such pixels is less than 5% when linear resampling is used and less than 1% when high-resolution cubic splines are used. For the HW corrected camera, the number is less than 1% up to about 0.1 pixel keystone (around pixel #110 in the figure) and increases approximately linearly from there and up to 13% at 0.3 pixel keystone. This means that in the areas of the sensor where the keystone is 0.3 pixels, only seven out of eight pixels give usable information for the HW corrected camera. Comparing the curves for the three cameras, we see that linear resampling gives number of pixels with large errors roughly equal to a HW corrected camera with 0.17 pixels keystone (the blue curve crosses the peaks of the red curve around pixel #185 in the figure), while a resampling camera that uses high-resolution cubic splines has number of pixels with large errors roughly equivalent to a

HW corrected camera with 0.1 pixel keystone (the blue curve crosses the peaks of the green curve around pixel #110 in the figure). The results in Fig. 12 confirm that the use of high-resolution cubic splines gives significantly smaller errors than linear resampling.

5.2 Misregistration Errors: Three Different Wavelengths

Light of shorter wavelengths is normally scattered more in the atmosphere than light of longer wavelengths. Light with stronger scattering is expected to give smaller errors both for a HW corrected camera and a resampling camera. In order to verify the performance of the cameras for different amount of scatter, we looked at the data in the same way as in Fig. 12, but this time for three different wavelengths (483, 697, and 865 nm). Results for linear resampling are omitted in order to

avoid cluttered graphs (we have already seen that use of linear resampling introduces significantly larger errors than use of high-resolution cubic splines). Figure 13(a) shows the standard deviation of the error, while Fig. 13(b) shows the relative number of pixels with error larger than 10% of the signal. We see that the misregistration errors for the shortest wavelength are significantly smaller than for the other two wavelengths. Nevertheless, the misregistration errors in the resampling camera that uses high-resolution cubic splines are still equivalent to a HW corrected camera with approximately 0.1 pixel keystone (all three curves for the HW corrected camera cross the peaks of the corresponding curves for the resampling camera around pixel #110 in both figures).

At the signal levels used in these simulations, the influence of photon noise would be relatively minor compared to the misregistration errors, and a 4× gain in signal level, provided by the faster optics of the resampling camera, would not be very noticeable. Separate graphs, where photon noise

is included, are therefore not shown here. Since, in general, the main errors in the case of a stronger signal will be generated by keystone rather than photon noise, the use of sensors with larger full well (i.e., larger peak SNR) does not necessarily improve the quality of the hyperspectral data. Any large keystone errors that are present in the data, and that are noticeable above the photon noise, will remain regardless of how much more light the sensor is able to collect.

5.3 Low Light

In order to check the performance of the cameras in low light conditions, the same test scene (Fig. 11) was reduced in brightness by a factor of 10. Photon noise was now included in the calculations and it was also taken into account that a resampling camera is capable of collecting four times more light than a HW corrected camera. Linear resampling was excluded from the analyses. Figures 14(a) (standard deviation of the error) and 14(b) (relative number of pixels

Fig. 13 (a) Standard deviation of the misregistration errors and (b) relative number of pixels with misregistration errors larger than 10% of the signal, for the HW corrected camera and the resampling camera that uses high-resolution cubic splines. For the HW corrected camera the keystone varies from 0 to 0.3 pixels from the left to the right part on the sensor. Results are shown for three different wavelengths: 483 nm (blue curve), 697 nm (green curve), and 865 nm (red curve). Photon noise is not included.

with errors larger than 10% of the signal) show clearly that the resampling camera (green curve) is more suitable for low light applications than the HW corrected camera (blue curve) and confirms the findings in Sec. 4.3. In fact, for the chosen signal level, the resampling camera performs significantly better than the HW corrected camera even when the keystone of the HW corrected camera is almost 0 [the left part of Figs. 14(a) and 14(b) where pixel #1 is positioned]. This is due to the fact that the misregistration errors of the resampling camera are negligible compared to the photon noise at the light levels used by the HW corrected camera in this case. Note, however, that the misregistration errors for the resampling camera are still visible (periodic variations in the green curve) at the light levels used by the resampling camera itself. Also, note that the presence of photon noise has lifted the curves (i.e., larger errors) for both cameras compared to Fig. 12 where only misregistration errors were considered.

5.4 Partial Correction of Keystone in Hardware and Resampling of Residual Keystone

If the residual keystone of a HW corrected camera is precisely characterized, then it is possible to try to further reduce the misregistration errors from such a camera by resampling.

In order to investigate the effectiveness of this approach, we will make graphs similar to Fig. 12. However, this time the resampling cameras will have only 2 pixels keystone, i.e., the image will have to be resampled from 322 to 320 pixels. Both high-resolution cubic splines and linear resampling will be used. The HW corrected camera will have 320.3 pixels as before. Figure 15(a) shows the standard deviation of the error for all three cameras, and Fig. 15(b) shows the relative number of pixels with errors exceeding 10% of the signal. The errors of the resampling cameras show periodic behavior as before, but now with only two periods due to the 2 pixels keystone. Slightly lower errors for the resampling cameras

Fig. 14 Camera performance in low light. The figure shows (a) the standard deviation of the errors and (b) the relative number of pixels with errors larger than 10% of the signal, for the HW corrected camera (blue curve) and the resampling camera that uses high-resolution cubic splines and collects four times more light (green curve). For the HW corrected camera, the keystone varies from 0 to 0.3 pixels from the left to the right part on the sensor. The calculations were done for the 697 nm wavelength.

Fig. 15 (a) Standard deviation of the misregistration errors and (b) relative number of pixels with error larger than 10% of the signal, for a HW corrected camera (blue curve) and two resampling cameras that use high-resolution cubic splines (green curve) and linear resampling (red curve), respectively. For the HW corrected camera, the keystone varies from 0 to 0.3 pixels from the left to the right part on the sensor. The resampling cameras have 2 pixels keystone. The calculations were done for the 697 nm wavelength. Photon noise is not included.

in the right part of the image (pixels 161–320) compared to the left part of the image (pixels 1–160) can be explained by a slight difference in spatial content of the scene (presence of small details with large contrast). Small differences in errors for different areas of the scene are also visible in Figs. 12 and 13.

The resampling camera has 0 keystone at pixel #1, 0.1 pixel keystone at pixel #16, 0.2 pixel keystone at pixel #32, 0.25 pixel keystone at pixel #40, and so on. The keystone reaches 0.5 pixel value at pixel #80. This keystone value causes the largest possible error during resampling. Keystone values larger than 0.5 pixel again give smaller errors, since the center of the resampling pixel now is moving away from the border between the two recorded pixels and becomes closer in position to one of them. The error drops toward zero again at pixel #160, where the keystone is 1 pixel large. Here, the resampled pixel is practically in the same position and occupies the same area as the corresponding pixel from the recorded image. After that, the error

increases again and reaches its second maximum at pixel #240, where the keystone reaches 1.5 pixel value. This long and somewhat trivial explanation illustrates that, when resampling is used, the largest possible errors are caused by keystone 0.5, 1.5, 2.5 pixels etc., because, with respect to the magnitude of errors, 0.7 pixel keystone is equivalent to 0.3 pixel keystone, 1 pixel keystone is equivalent to 0 pixel keystone, and so on.

Figure 15(a) gives us insight into what errors we can expect if a residual keystone of a HW corrected camera is further corrected by resampling. We will focus on resampling with high-resolution cubic splines and not discuss linear resampling any further here. If the original data has 0.1 pixel residual keystone, for the green resampling curve this corresponds to pixel #16, #143, #175 or #302 (marked by vertical black arrows in the figure), the standard deviation of the error after resampling (green curve) will be approximately 1.2% (marked by horizontal dashed red line). For a HW corrected camera with no resampling (blue curve), this

corresponds to approximately 0.047 pixel keystone (pixel #50, marked by a vertical red arrow). This means that if we have a HW corrected camera where:

- a. the keystone of the camera is precisely known, and
- b. the maximum value of the keystone is 0.1 pixel,

then, after resampling with the high-resolution cubic splines method, we should expect that the misregistration errors are reduced, so that they are similar to those from a HW corrected camera with only 0.047 pixel residual keystone. This example shows that it may be possible to significantly reduce residual misregistration errors in HW corrected cameras by postprocessing, without altering the hardware in any way.

However, there are some limitations to this method:

1. If the residual keystone of a HW corrected camera exceeds approximately 0.25 pixel, resampling of the data will yield very minor improvement compared to resampling of data with very large keystone. This can be seen directly from Fig. 15(a). The green resampling curves are relatively flat in the areas where the keystone is 0.25–0.5 pixels (pixels #40–120 and pixels #200–280 in the figure), indicating that any efforts to reduce the keystone in hardware in this range will only result in minor reductions in the errors in the final data cube after resampling. Achieving a residual keystone smaller than 0.25 pixel in a high-resolution camera is, of course, easier than aiming for 0.1 pixel keystone, but it is still a challenge compared to building a camera without any keystone correction.
2. The method requires a stable keystone. It is generally more difficult to keep a small keystone stable than a larger one. This may make it difficult to use the method, since the initial residual keystone should preferably be smaller than approximately 0.25 pixel.
3. It is much more difficult to create a high-resolution system with high light gathering capacity when a very low keystone is required compared to when a large keystone is acceptable.

We believe that, very often, it is more beneficial to focus the design effort on achieving very high spatial resolution and very low F-number, leaving keystone correction to resampling. Such a camera would collect higher quality data simply because of much higher spatial resolution and lower photon noise than the cameras with hardware correction of keystone. However, if high resolution sensors are not available (as is the case for most of the IR region) or there is already a camera (with low, well characterized, and stable keystone) which must be used in the best possible way, then high-resolution cubic splines resampling in addition to HW correction of keystone may be the way to go. We find it hard to recommend linear resampling for this task, since it gives larger errors, unless the processing speed is the limitation (which is rarely the case for low-resolution cameras).

6 Subjective Quality of Resampled Images

So far, we have shown that a resampling camera that uses high-resolution cubic splines performs well compared to

Fig. 16 Subjective quality of resampled image: (a) reference image and (b) resampled image. The high-resolution cubic splines method was used for the resampling.

traditional cameras where keystone is corrected in hardware. However, we have to make sure that the subjective quality of the image does not suffer too much because of resampling. Those, who have seen linearly resampled images, probably remember their slightly blurred appearance. Figure 16(b) shows part of an image which was captured by a 352 pixels virtual sensor and resampled by use of high-resolution cubic splines to 320 pixels. We have chosen a scene with various features, such as point sources, small objects, borders, and so on. Even a very careful subjective examination of this image shows more or less no sign of blur compared to the reference image captured with a 320 pixels virtual camera with 0 keystone [Fig. 16(a)].

The high-resolution cubic splines resampling method is optimized to deliver the highest objective quality of the resampled data.³ One could of course apply two different resampling methods to the raw datacube with large keystone: one optimized for objective measurements and one for subjective evaluations. However, Fig. 16 shows that there is no need to do so when high-resolution cubic splines are used for resampling the data with large keystone, since the subjective impression is that the resampled image is virtually identical to the reference image.

7 Practical Implementation of Resampling

In order to utilize the high-resolution cubic splines resampling method (or any other resampling method) for correcting the keystone of a real camera, the keystone must be characterized across the field of view for every spectral channel. This characterization can be done in the lab as the final step of camera manufacturing. Typically, the keystone changes are very slow across the field of view. Therefore, it may not be necessary to characterize the keystone for every pixel. Instead, the keystone could be measured in several field points placed densely across the field of view and interpolated between the measurement points. Of course, in order to take full advantage of the resampling method used, the keystone has to be characterized quite precisely (significantly more precise than 0.1 pixel in the case of the high-resolution cubic splines resampling method).

In order to simplify the simulations we have assumed a linear keystone distribution across the field of view. However, this is not a requirement. Since high-resolution cubic splines resampling is performed individually for

every spatial pixel, and uses data only from the four nearest pixels of the input image [Eq. (1)], this resampling method can be used for more or less any practically relevant keystone distribution across the field of view.

Before attempting to perform resampling, the spatial coordinates of the pixels in the output image should be known. These spatial coordinates can be defined differently, depending on what is most suitable for the application. For example, the coordinates of the pixels in the output image could be an evenly spaced grid or the input data with large keystone could be resampled directly into the pixel coordinates of a georeferenced image. Alternatively, if a resampling camera was used in combination with a HW corrected camera (a typical example would be two hyperspectral cameras capturing the same scene in two different spectral ranges), the data from the resampling camera could be resampled to match the pixel coordinates of the HW corrected camera.

The high-resolution cubic splines resampling method is suitable for real-time processing. The method does not rely on having data from the whole image; therefore, the processing can be done on the incoming data already during image acquisition. In order to generate one pixel of the output image, only four multiplications and three additions are required. Also, calculations for different output pixels of the same image line can be distributed across multiple central processing unit and graphics processing unit cores, if necessary.

8 Conclusion

Current high-resolution hyperspectral cameras attempt to correct misregistration errors in hardware. Usually, it is required that aberrations in the optical system must be controlled with precision 0.1 pixel or smaller. This severely limits other specifications of the hyperspectral camera, such as spatial resolution and light gathering capacity, and often requires very tight tolerances. If resampling is used to correct keystone in software instead of in hardware, then these stringent requirements could be lifted. Preliminary designs show that a resampling camera should be able to resolve at least 3000–5000 pixels, while at the same time collecting up to four times more light than the majority of current high spatial resolution HW corrected cameras.

A virtual camera software was used to compare the performance of resampling cameras and traditional HW corrected cameras. The performance was measured by comparing the resulting image (after being processed by the virtual camera) to the “ideal” input image and calculating the corresponding errors. The simulations showed that the performance of a resampling camera is comparable to that of a HW corrected camera with 0.1 pixel residual keystone. It is important to note that this level of performance was achieved with virtually no downsampling. This opens up a possibility to design and build hyperspectral cameras based on resampling with very high spatial resolution and fairly low misregistration errors. In low light, the advantages of the resampling camera, with its ability to collect about four times more light, became very visible: the errors

were significantly lower than a HW corrected camera, even in the case of zero keystone.

We have also shown that the use of a more advanced resampling method than the commonly used linear interpolation, such as for instance high-resolution cubic splines, is highly beneficial for the data quality of the resampled image. In addition to giving significantly smaller misregistration errors, the subjective quality of the resampled image does not seem to suffer when high-resolution cubic splines are used for the resampling.

We have suggested a new criterion for evaluating camera performance. In addition to looking at the standard deviation of the error, we suggest to use also the relative number of pixels where the error exceeds a certain threshold value as a criterion.

Our findings in this article suggest that if high-resolution sensors are available, it would be better to use resampling instead of trying to correct keystone in hardware. We believe that a similar approach as described for keystone correction, could (and should) be used for smile correction, where over-sampling in the spectral direction normally easily can be performed.

References

1. P. Mouroulis, R. O. Green, and T. G. Chrien, “Design of pushbroom imaging spectrometers for optimum recovery of spectroscopic and spatial information,” *Appl. Opt.* **39**(13), 2210–2220 (2000).
2. D. Schlöpfer, J. Nieke, and K. Itten, “Spatial PSF non-uniformity effects in airborne pushbroom imaging spectrometry data,” *IEEE Trans. Geosci. Remot. Sens.* **45**(2), 458–468 (2007).
3. J. A. Parker, R. V. Kenyon, and D. E. Troxel, “Comparison of interpolating methods for image resampling,” *IEEE Trans. Med. Imaging* **2**(1), 31–39 (1983).
4. G. Høye and A. Fridman, “Performance analysis of the proposed new restoring camera for hyperspectral imaging,” FFI-rapport 2010/02383, declassified on 25 March 2014, Norwegian Defence Research Establishment (FFI), Kjeller, Norway (2010).
5. G. Høye and A. Fridman, “Mixel camera – a new push-broom camera concept for high spatial resolution keystone-free hyperspectral imaging,” *Opt. Express* **21**, 11057–11077 (2013).
6. “VNIR 1600 camera specifications,” Norsk Elektro Optikk, Norway, <http://www.hyspex.no/products/hyspex/vnir1600.php> (20 April 2014).
7. P. Mouroulis et al., “A compact, fast, wide-field imaging spectrometer system,” *Proc. SPIE* **8032**, 80320U (2011).
8. R. Lucke and J. Fisher, “The Schmidt-Dyson: a fast space-borne wide-field hyperspectral imager,” *Proc. SPIE* **7812**, 78120M (2010).

Andrei Fridman is an optical designer at Norsk Elektro Optikk. He received his MSc degree in optics from the Technical University of Fine Mechanics and Optics in St. Petersburg, Russia, in 1994. In addition to his main optical design activities, his interests include image processing.

Gudrun Høye is a researcher at Norsk Elektro Optikk in addition to her main employment as a principal scientist at the Norwegian Defence Research Establishment. She received her MSc degree in physics from the Norwegian Institute of Technology in 1994 and her PhD degree in astrophysics from the Norwegian University of Science and Technology in 1999. Her current research interests include hyperspectral imaging and maritime surveillance.

Trond Løke is a senior research scientist at Norsk Elektro Optikk. He received his MS degree in photonics from the Norwegian University of Science and Technology in 2003. Since 2003, he has been working in the hyperspectral (HySpex) department at Norsk Elektro Optikk.