FFI-rapport 2009/01351

A method for restoring data in a hyperspectral imaging system with large keystone without loss of spatial resolution

Gudrun Høye (FFI) og
Andrei Fridman (NEO)

FFI Forsvarets forskningsinstitutt

FFI-rapport 2009/01351

FORTROLIG
iht. Beskyttelsesinstruksen
jf. offentleglova § 14

A method for restoring data in a hyperspectral imaging system with large keystone without loss of spatial resolution

Gudrun Høye (FFI) og Andrei Fridman (NEO)

Forsvarets forskningsinstitutt/Norwegian Defence Research Establishment (FFI)

18 August 2009

FFI-rapport 2009/01351

1142

Keywords

Hyperspektral avbildning

Hyperspektralt kamera

Keystone korreksjon

Bildebehandling

Avbildende spektroskopi

Avbildende spektrometri

Approved by

Svein Erik Hamran	Project Manager
Johnny Bardal	Director

English summary

This report proposes a new method of keystone correction in the postprocessing of hyperspectral images. Unlike conventional resampling the proposed method does not introduce any loss of resolution. A hardware modification of the hyperspectral camera, which will be necessary for the practical implementation of the method, is also briefly discussed.

The potential advantages of the proposed method are large. When keystone correction is no longer required in hardware, it will be possible to design significantly sharper and/or faster optics. In addition, such optics may be both smaller and cheaper than the optics of the current hyperspectral cameras.

We suggest a joint FFI-NEO project with the goal of developing the method further and building a new hyperspectral camera based on it. Such a project would benefit from NEO's expertise in design of hyperspectral cameras and FFI's expertise in processing of hyperspectral data. The outcome of the project could be a rather impressive instrument which will perform substantially better than the current generation of hyperspectral cameras.

Sammendrag

Denne rapporten presenterer en ny metode for keystone korreksjon i etterprosesseringen av hyperspektrale data. Den foreslåtte metoden bevarer dataenes oppløsning, i motsetning til tradisjonell resampling som gir dårligere oppløsning. En hardware modifikasjon av det hyperspektrale kameraet, som vil være påkrevet for implementeringen av metoden, diskuteres også kort.

Metoden har potensielt store fordeler. Når det ikke lenger er nødvendig å korrigere for keystone i hardware vil det være mulig å designe vesentlig skarpere og/eller mer lyssterk optikk, som også vil kunne bli både mindre og billigere enn det man finner i dagens hyperspektrale kameraer.

Vi foreslår å opprette et felles FFI-NEO prosjekt med det formål å utvikle metoden videre og lage et nytt hyperspektralt kamera basert på metoden. Prosjektet vil dra nytte av NEOs ekspertise innenfor design av hyperspektrale kameraer og FFIs ekspertise innenfor prosessering av hyperspektrale data. Resultatet vil kunne bli et hyperspektralt kamera med vesentlig forbedret ytelse sammenlignet med dagens kameraer.

Contents

Preface

Hyperspectral cameras are increasingly used for various military, scientific, and commercial purposes. Users constantly demand better optical resolution and higher light sensitivity. However, the current generation of instruments already has extremely tight tolerances for optical aberrations compared to other imaging systems. For this reason, the development of hyperspectral cameras has more or less converged to a couple of “standard” layouts, each of them with some inherent limitations such as minimum possible F-number, maximum possible spatial resolution, etc.

As a result of cooperation between Forsvarets forskningsinstitutt (FFI) and Norsk Elektro Optikk AS (NEO) a novel concept has emerged. It offers a new method of data processing which in combination with hardware modifications will allow the creation of a substantially better instrument.

1 Introduction

Anyone would like to have a hyperspectral camera with very high resolution, very high sensitivity and very low smile and keystone effects. As the new sensors with higher pixel count become available, camera manufacturers try to develop the optics for these sensors. This optics must be sharper and faster in order to justify the use of a newer better sensor. And since the requirements for smile and keystone effects are normally set relative to the pixel size, the absolute smile and keystone errors must decrease. This makes development of new optics increasingly difficult.

An additional problem is the fact that optimization for smile/keystone requires a lot more calculations than optimization just for image sharpness. In fact, the difference is huge; even a simplified smile/keystone optimization routine runs at least 20-30 times slower than a routine which takes into account only image sharpness. It means that even if it is possible to design an optical system which is acceptable in terms of aperture, sharpness and the smile/keystone error (and this is not always the case), designing such a system will take 20-30 weeks instead of 1 week, or maybe 20-30 months instead of 1 month.

It is very tempting to ask the following question: is it really necessary to have so tight requirements for smile and keystone? All the data is there even with smile and keystone present – it's just not arranged in the same «neat» way. A possible solution, with smile and keystone present in the system, would be to use resampling to restore the datacube (Figure 1.1). However, quite some resolution would be lost in the process.

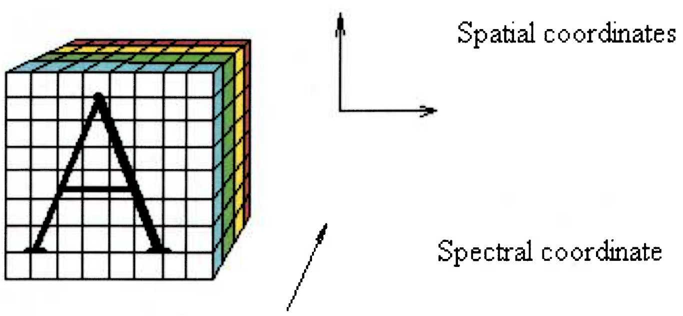

Figure 1.1 The datacube.

Of course, sharper optics allows to use sensors with smaller pixel pitch, which will compensate the loss of resolution caused by resampling. But is there another way? Can it be that it is possible to convert the distorted data to the required "cubic" shape without loss of resolution? We have earlier proposed one possible way to do that [1]. By making certain assumptions about the scene, it is possible to avoid resampling of small objects and borders in a scene, and resample just the areas which do not contain small features. Potentially this method may deliver high resolution data with no measurable smile and keystone. However, the method relies on the presence of relatively large areas with no small details. Whenever a part of a scene is too "busy", i.e. the small objects and the borders are packed too densely, this method will "retreat" to resampling in that area. This method is a *lossy* way to convert data with keystone into a proper datacube.

In this paper we will propose a *lossless* way to convert data with large keystone into a datacube. The potential benefits of such an approach are large. As already mentioned, it is very much easier to design optics with no smile/keystone correction. Such a system may collect at least 3 times more light compared to the widely used Offner design. At the same time the system will have similar optical resolution and be both smaller and cheaper. Note that "3 times more light" is not a fundamental limitation – it is just a real world example. It may very well be possible to design even faster optics.

Eventually it will be necessary to have both smile and keystone corrected, but for now let us look at the keystone correction only.

2 The method

2.1 A simple example

The method we propose to use for correcting the keystone can be explained by a simple numerical example.

Let us consider one spatial line which is 8 pixels long. We will be looking at one half of it only; from the center of the field of view to the edge – it will be 4 pixels. Let us also say that at a certain wavelength there is 1 pixel keystone. The 4 pixels from the scene are then recorded into 5 pixels on the sensor.

Let us say that in case of this particular scene at this particular wavelength the pixels in the scene have the following values:

$$\begin{aligned} S_1 &= 10 \\ S_2 &= 30 \\ S_3 &= 100 \\ S_4 &= 50 \end{aligned} \qquad (2.1)$$

This gives the following values for the recorded pixels, see Figure 2.1:

$$\begin{aligned}
S_1^R &= \tfrac{4}{5}\cdot S_1 = \tfrac{4}{5}\cdot 10 = 8 \\
S_2^R &= \tfrac{1}{5}\cdot S_1 + \tfrac{3}{5}\cdot S_2 = \tfrac{1}{5}\cdot 10 + \tfrac{3}{5}\cdot 30 = 20 \\
S_3^R &= \tfrac{2}{5}\cdot S_2 + \tfrac{2}{5}\cdot S_3 = \tfrac{2}{5}\cdot 30 + \tfrac{2}{5}\cdot 100 = 52 \\
S_4^R &= \tfrac{3}{5}\cdot S_3 + \tfrac{1}{5}\cdot S_4 = \tfrac{3}{5}\cdot 100 + \tfrac{1}{5}\cdot 50 = 70 \\
S_5^R &= \tfrac{4}{5}\cdot S_4 = \tfrac{4}{5}\cdot 50 = 40
\end{aligned} \tag{2.2}$$

We have here assumed that the intensity distribution over each pixel in the scene is uniform.

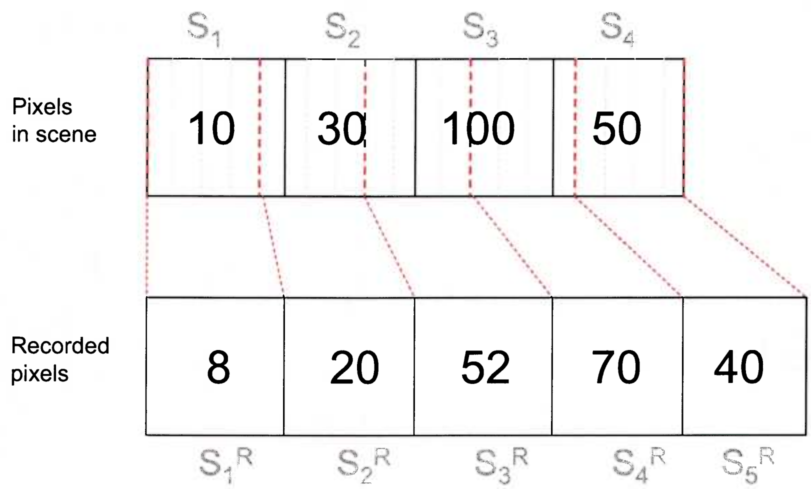

Figure 2.1 Scene pixels with known values and corresponding recorded pixels.

Of course, in real life we will not know the actual values of the scene pixels. This situation is shown in Figure 2.2.

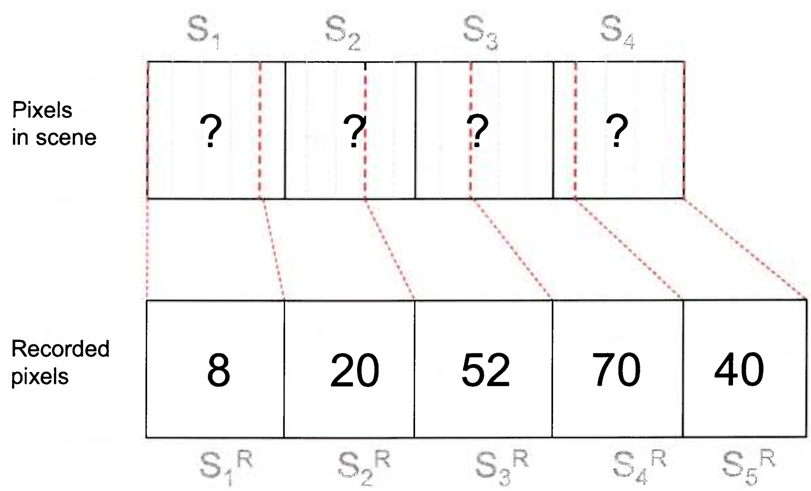

Figure 2.2 Recorded pixels with known values and corresponding scene pixels with unknown values.

In order to determine the values of the scene pixels, we set up the following set of equations:

$$\begin{aligned} \tfrac{4}{5}\cdot S_1 &= S_1^R = 8 \\ \tfrac{1}{5}\cdot S_1 + \tfrac{3}{5}\cdot S_2 &= S_2^R = 20 \\ \tfrac{2}{5}\cdot S_2 + \tfrac{2}{5}\cdot S_3 &= S_3^R = 52 \\ \tfrac{3}{5}\cdot S_3 + \tfrac{1}{5}\cdot S_4 &= S_4^R = 70 \\ \tfrac{4}{5}\cdot S_4 &= S_5^R = 40 \end{aligned} \qquad (2.3)$$

The system can easily be solved for the unknowns S_1, S_2, S_3, and S_4, giving the following values for the scene pixels:

$$\begin{aligned} S_1 &= 10 \\ S_2 &= 30 \\ S_3 &= 100 \\ S_4 &= 50 \end{aligned} \qquad (2.4)$$

which are identical to the actual values in the scene pixels as given in Equation (2.1). We have now managed to restore the true values of the 4 pixels in the scene based only on the information about the values of the 5 recorded pixels and the size of the keystone (1 pixel).

This was a simple numerical example. Now let us look at the general case.

2.2 The general equations

We will assume again that the intensity distribution over each pixel is uniform. In any real scene this is almost never true, but we will take care of this problem later.

Let us now consider the situation where we want to restore *N* pixels in the scene from *M* recorded pixels (*M*>*N*). The keystone is then equal to (*M*-*N*) pixels. This situation is shown in Figure 2.3.

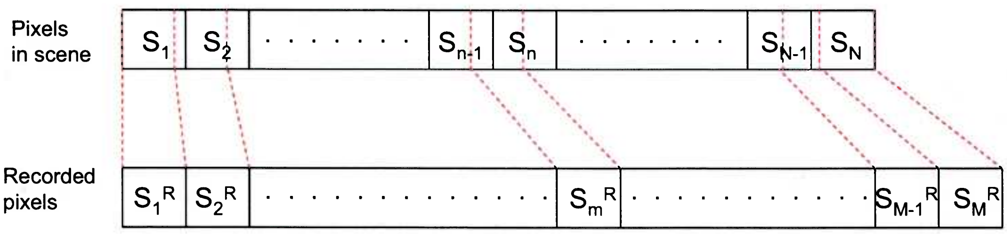

Figure 2.3 Scene pixels and corresponding recorded pixels for the general case.

We can now set up the following set of equations:

$$S_m^R = \sum_{n=1}^{N} q_{mn} S_n, \qquad m = 1, 2, ..., m, ..., M-1, M \tag{2.5}$$

where:

S_n	-	Signal level for pixel $\#n$ in the scene.
S_m^R	-	Signal level recorded in pixel $\#m$ on the sensor.
q_{mn}	-	Percentage of the signal level in pixel $\#n$ in the scene that contributes to the signal level recorded in pixel $\#m$ on the sensor.
N	-	Total number of pixels in the scene.
M	-	Total number of pixels recorded on the sensor.

Here S_m^R is known (measured) and q_{mn} is known (can be calculated) for all m and n when the keystone for the system is known.

Equation (2.5) can be written in matrix form:

$$\begin{pmatrix} q_{11} & \cdots & q_{1n} & \cdots & q_{1N} \\ \vdots & & \vdots & & \vdots \\ q_{m1} & \cdots & q_{mn} & \cdots & q_{mN} \\ \vdots & & \vdots & & \vdots \\ \vdots & & \vdots & & \vdots \\ q_{M1} & \cdots & q_{Mn} & \cdots & q_{MN} \end{pmatrix} \cdot \begin{pmatrix} S_1 \\ \vdots \\ S_n \\ \vdots \\ S_N \end{pmatrix} = \begin{pmatrix} S_1^R \\ \vdots \\ S_m^R \\ \vdots \\ \vdots \\ S_M^R \end{pmatrix} \tag{2.6}$$

Since no more than two scene pixels (and sometimes just one) contribute to each recorded pixel, most of the coefficients q are equal to zero. The matrix will then typically have the form:

$$\begin{pmatrix} q_{11} & & & & & \\ q_{21} & q_{22} & & & & \\ & \ddots & \ddots & & & \\ & & q_{m(n-1)} & q_{mn} & & \\ & & & \ddots & \ddots & \\ & & & & q_{(M-1)(N-1)} & q_{(M-1)N} \\ & & & & & q_{MN} \end{pmatrix} \cdot \begin{pmatrix} S_1 \\ \vdots \\ \vdots \\ S_n \\ \vdots \\ \vdots \\ \vdots \\ S_N \end{pmatrix} = \begin{pmatrix} S_1^R \\ \vdots \\ \vdots \\ \vdots \\ S_m^R \\ \vdots \\ \vdots \\ \vdots \\ S_M^R \end{pmatrix} \tag{2.7}$$

where the coefficients q_{mn} are nonzero only along the diagonals and zero everywhere else. We have here assumed that the keystone is an integer number of pixels and that the first pixel in the scene starts exactly at the beginning of the first recorded pixel. However, the method works just as well when the keystone is not an integer number of pixels and/or if the first pixel in the scene does not start exactly at the beginning of the first recorded pixel.

Note that:

$$\sum_{n=1}^{N} q_{mn} = \frac{N}{M} \qquad (2.8)$$

and

$$\sum_{m=1}^{M} q_{mn} = 1 \qquad (2.9)$$

A fraction N/M of a scene pixel corresponds to 1 recorded pixel sizewise (see Figure 2.3) and Equation (2.8) states that for any m the sum of q_{mn} over all N pixels in the scene must equal this fraction. This makes sense since this sum represents the total contribution from all pixels in the scene to the recorded pixel #m. Equation (2.9) expresses that the total content (100%) of pixel #n in the scene has been used when summed over all recorded pixels[1].

The matrix system (2.7) can now be solved for the unknowns S_n. Note that the system has more equations than unknowns ($M>N$). In fact, each extra pixel of keystone gives one extra equation. However, for the ideal case when there is no noise in the system, the matrix system is compatible, i.e., can be solved. For a real system with noise, the system is overdetermined and an optimization method such as for instance the least squares method must be used to obtain the solution.

The process described above can be repeated for all spectral lines. This means that even though the spectral lines have different length (i.e. different keystone) they will all be converted to the same final grid with no loss of resolution.

Note that solving the matrix system (2.7) for S_m^R (unknown values of the recorded pixels) when S_n is known, corresponds to traditional (linear) resampling. Instead we do the inverse process; we solve the same matrix system for S_n (unknown values of the scene pixels) when S_m^R is known. The method we have suggested can therefore be described as *inverse* resampling; we find the original values of the scene pixels from the data resampled by the sensor.

[1] This is true except for the end pixel in the scene if the keystone is not an integer, or if the keystone is an integer but the first pixel in the scene does not start exactly at the beginning of the first recorded pixel.

3 The optics

It has already been pointed out that the method only works if the light distribution in each single scene pixel is uniform. Since in most cases this assumption is not valid, something has to be done about it. A possible solution would be to separate the camera slit into small chambers (1 chamber for each scene pixel) and to mix the light inside each of these chambers (see Figure 3.1).

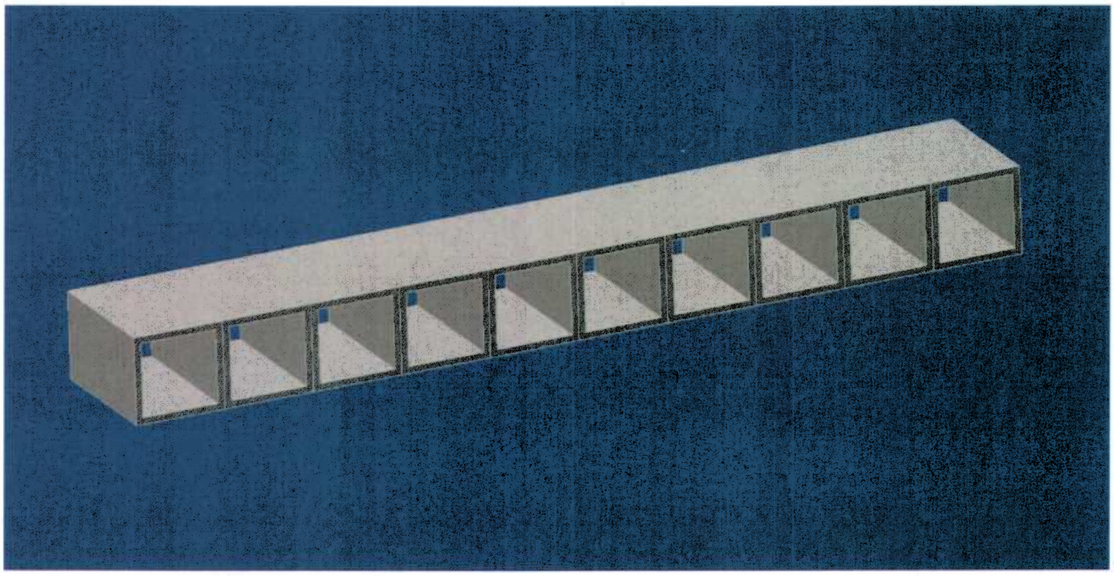

Figure 3.1 The slit chambers.

Achieving reasonably even light distribution across each of these chambers is not a straightforward task. Manufacturing such a slit may not be very easy either, since each chamber is only ~3-10 times larger than a pixel on the optical sensor. However, a brief analysis of the problem has shown that it should be possible both optically and technologically to make such a slit. Diffraction in the optics, optical aberrations or slight defocus of the collimator after the slit will mix the light inside each scene pixel even more. The design of such a slit will be a crucial part of the camera design.

4 A brief noise analysis

In order to understand how useful this method of eliminating keystone will be, it is important to understand how noise is transferred from the recorded pixels to the restored scene pixels. When it is acceptable to have large keystone in the system, it is possible to design much faster optics. However, getting more light onto the sensor is a good thing only as long as it means having less noise in the final data. Traditional resampling gives less noise in the data, but since the suggested method is in fact inverse resampling, we may get more noise in the final data than we had directly from the sensor.

One way to see what happens with the sensor noise during the data processing is to take a look at an example.

Let us say that the scene consists of 1000 scene pixels. Further, let us say that 220 of these pixels are relatively bright (value 100) and the remaining 780 pixels are noticeably darker (value 50). These two types of pixels are arbitrarily distributed in the scene.

We will check the noise performance for three different keystone values: 1, 10, and 100 pixels (i.e., the number of recorded pixels will be 1001, 1010 or 1100). This is done by performing the following steps for each case:

1) Calculate the values for the recorded pixels (without noise) based on the known (true) values for the scene pixels.
2) Add noise with standard deviation 0.5 to the recorded pixels.
3) Calculate the (noisy) values for the scene pixels from the noisy recorded pixels.
4) Calculate the relative noise for the scene pixels. For each pixel the relative noise is calculated as the difference between the noisy and the true value, divided by the true value.
5) Compare the relative noise for the scene pixels with the relative noise for the recorded pixels.

The calculations were performed in Matlab, which applies the least-squares method to solve overdetermined matrix systems.

Figure 4.1a) shows the relative noise for the recorded pixels. The standard deviation for the relative noise is 0.010 and is shown with a red dotted line. The number of recorded pixels will of course vary according to the keystone in the system. In this particular case, 1000 scene pixels and 100 pixels keystone yield 1100 recorded pixels. If the keystone is smaller, the number of recorded pixels will decrease accordingly. The question now is, how will the noise in the recorded pixels affect the quality of the restored data?

Figure 4.1b) shows the relative noise in the restored scene pixels when the keystone in the recorded data is 1 pixel. The noise at the beginning and at the end of the restored scene pixel array appears to be of the same size as the noise in the recorded pixels. However, there is a pronounced peak in the noise level in the middle of the restored scene pixel array. If we increase the keystone in the recorded data to 10 pixels the number of peaks also increases to 10, see Figure 4.1c). The peaks are smaller and narrower, but can still be seen.

Figure 4.1d) shows the relative noise in the restored scene pixels when the keystone in the recorded data is 100 pixels. In this case, the noise in the restored data appears fairly random and the standard deviation is found to be 0.013 (shown with red dotted line). This is 1.3 times larger than the relative noise in the recorded pixels. The method therefore seems to increase the noise somewhat. However, the increase in noise is relatively small, and the expected increase in the optical system's light gathering capacity when allowing for large keystone certainly outweighs it.

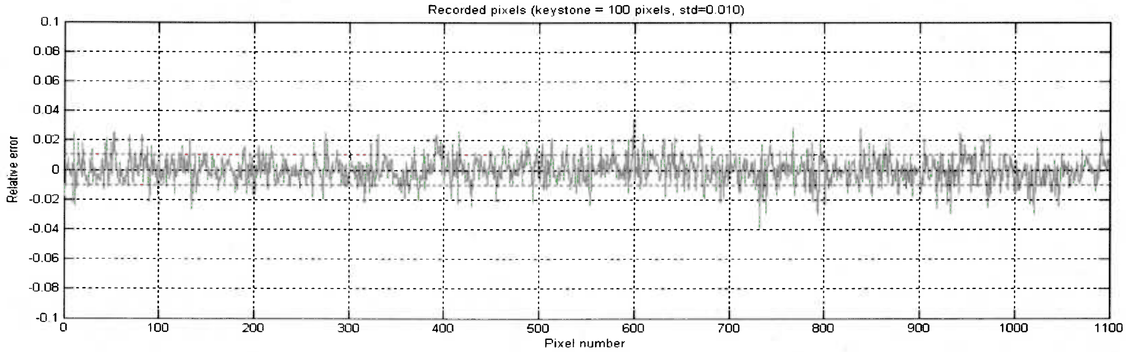

a)

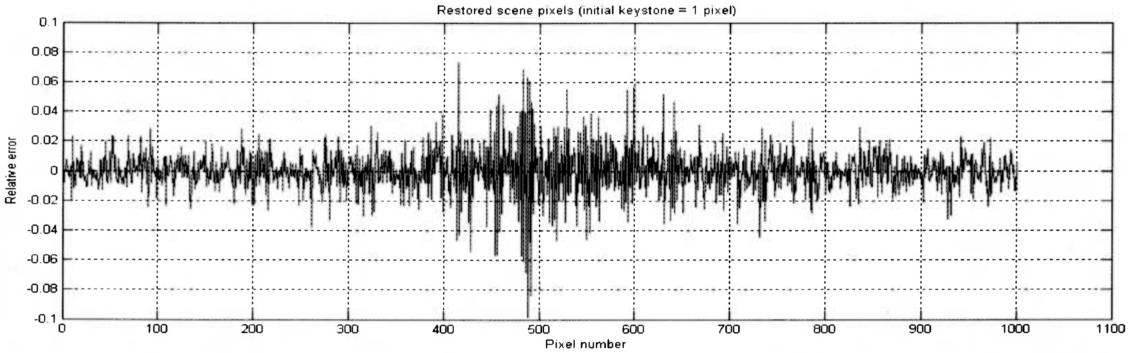

b)

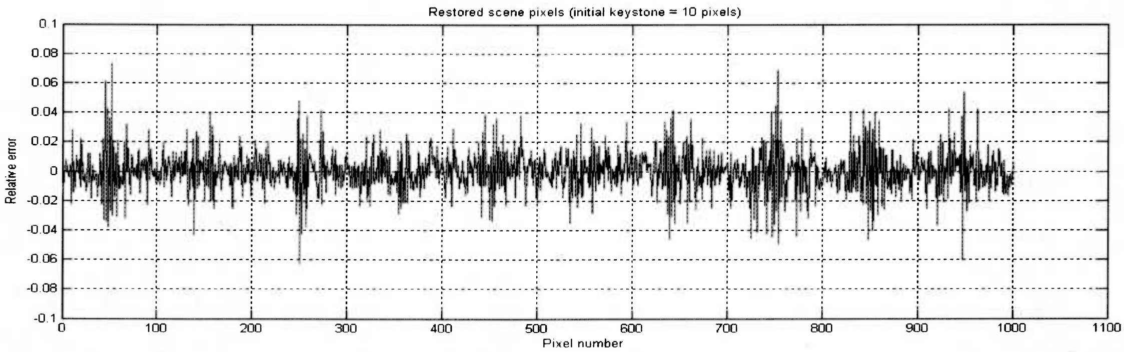

c)

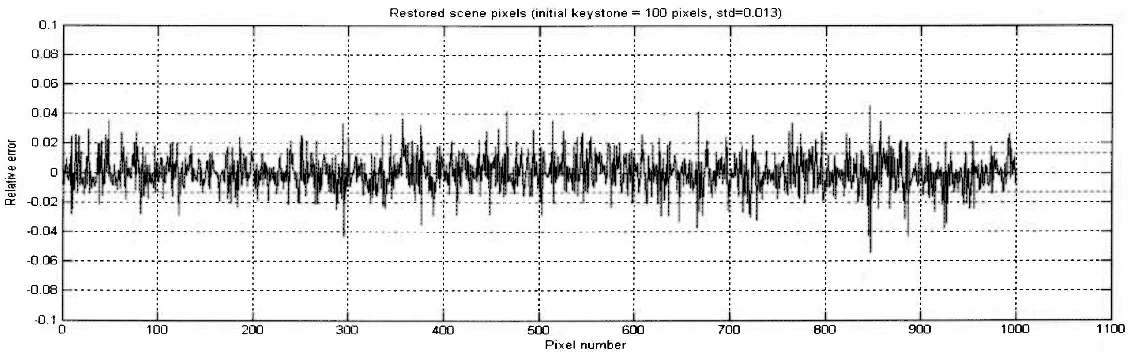

d)

Figure 4.1 Relative noise in recorded pixels and in scene pixels with different initial keystone.

A potential problem is the non-uniformity of the noise distribution over the restored scene pixels (with peaks as described above). A simple solution would be to use a relatively large keystone, since the noise distribution then becomes more uniform. It may also be possible to optimize the solution of the matrix system for a more uniform noise distribution, for instance by applying different weights to the different pixels during the optimization process.

In addition to sensor noise, other possible error sources in the restored scene pixels are:

- Not precise enough keystone characterization of the system.
- Non-uniform light distribution in the pixels in the scene.

Errors in the keystone characterization introduce noise in the coefficients q_{mn} in the matrix system (2.6), which will affect the resulting noise level in the restored scene pixels. Non-uniform light distribution in the scene pixels will similarly affect the calculated values of the restored scene pixels, and may possibly also be treated as noise in the coefficients q_{mn}. The effect of these error sources on the noise level in the restored scene pixels should be investigated further.

5 Advantages and disadvantages of the method

Compared to the current Offner cameras, the described method will allow to design hyperspectral cameras which are smaller and most likely cheaper, while at the same time being able to collect more light and/or providing higher resolution images. The expected increase of noise level in the new system will be compensated by a much higher light level. Also, when comparing the new system's performance with the current cameras with hardware corrected keystone, it is important to remember that we are comparing *noise* in the new system to *noise* and *keystone* in the traditional one.

Compared to the hyperspectral cameras that use resampling for keystone correction, the new system has a lower signal-to noise-ratio. However, the new system will have the capability to utilize the full spatial resolution of the sensor, yielding up to 2 times higher spatial resolution than a camera with the same sensor that uses resampling to correct the keystone. The higher spatial resolution does, however, come at the cost of a more complicated slit design.

It is very important that a hyperspectral camera has similar point spread function for each pixel on the sensor. Probably the requirements for uniformity of the point spread function across the sensor should be as tight as the requirements for smile and keystone [2]. In traditional cameras this is very difficult to achieve if the optics is diffraction limited (which is the case more and more often), since the size of the diffraction pattern is very wavelength dependent. The optics in the proposed camera, on the other hand, may be made to have quite uniform point spread function across the sensor. This is possible because the optics will be faster (and therefore will have much weaker diffraction effects) and because the point spread function in this case will be the result of mixing the light in each slit pixel in exactly the same way.

The method is at present not able to correct smile effect. After eliminating the keystone, smile will have to be corrected by resampling in the spectral direction. This means that in order to achieve the same spectral resolution as an Offner camera, the new camera will need 2 times more pixels in the spectral direction. It may or may not be a problem in real life.

So far it looks like the new system both has the potential to give higher spatial resolution than a camera with resampling and to collect much more light than the state of the art Offner cameras.

6 What is next?

In order to build a new hyperspectral camera based on the described method, several things should be investigated more closely.

Until now we have been assuming that the intensity transition between scene pixels is instant. This is almost exactly right for a system with very large pixels and very sharp and fast optics. However, the method should be extended so that the type of transition can be taken into account. Certainly, it should not be strictly necessary to have a uniform light distribution in each scene pixel in order to calculate the values of the scene pixels from the recorded pixels. Having *known* distribution in each scene pixel should be good enough.

A solution for the slit design should be found. The chambers in the slit plane should mix the light from each scene pixel very well. Preferably the mixing process should not increase the numerical aperture after the slit. A possibility for slight defocus of the collimator after the slit for better mixing should be considered. The consequences of such defocus for the validity of the model as well as for image sharpness should be investigated. During the design process it is important to remember that the size of the slit chambers is very small (only a few times larger than sensor pixels), and some promising designs may therefore not be very production friendly.

The method seems to increase the noise, which may diminish somewhat the benefit of the larger light gathering capacity of the new camera. In this report we briefly checked how the sensor noise is transferred to the final data, and the noise increase seems to be very moderate. However, there are other potential sources of noise in the system, such as for instance non-uniform light distribution in the scene pixels. The influence on the noise in the restored scene pixels from such error sources should be thoroughly checked. However, when the new system's performance is being judged it is important to remember that we are comparing *noise* in the new system to *noise* and *keystone* in the traditional one.

The method should take into account the fact that pixels have less than 100% fill factor. This concerns both the sensor pixels and the slit pixels.

It should be investigated whether it is possible to use a continuous function (such as for instance cubic splines) to describe the light distribution in the slit plane instead of the step function. This would remove the need for light mixing chambers in the slit plane.

7 Conclusion

We have discussed a new method for keystone correction in postprocessing, as well as the hardware modification which is necessary for implementation of the method. The potential advantages of the method are large.

We suggest a joint FFI-NEO project with the goal of developing the described method further and building a new hyperspectral camera based on it. Such a project would benefit from NEO's expertise in design of hyperspectral cameras and FFI's expertise in processing of hyperspectral data. The outcome of the project could be a rather impressive instrument which will perform substantially better than the current generation of hyperspectral cameras.

References

[1] A. Fridman and G. Høye (2009): "Use of large keystone effect for high quality hyperspectral imaging", Norsk Elektro Optikk AS, January 2009.

[2] P. Mouroulis (1999): "Spectral and spatial uniformity in pushbroom imaging spectrometers", Proceedings Paper, SPIE Vol.3753, pp. 133-141, October 1999.

~~FORTROLIG~~

~~iht. Beskyttelsesinstruksen jf. offentleglova § 14~~

FFI Forsvarets forskningsinstitutt

Norwegian Defence Research Establishment

www.ffi.no

Norwegian Defence Research Establishment (FFI)
P O Box 25
N-2027 Kjeller

Office address:
Instituttveien 20
N-2007 Kjeller

Telephone: +47 63 80 70 00
Telefax: +47 63 80 71 15
E-mail: ffi@ffi.no

Forsvarets forskningsinstitutt
Postboks 25
2027 Kjeller

Besøksadresse:
Instituttveien 20
2007 Kjeller

Telefon: 63 80 70 00
Telefaks: 63 80 71 15
E-post: ffi@ffi.no

Design/produksjon FFI