

A Multi-band statistical restoration of the Aqua MODIS 1.6 micron band

Irina Gladkova^a, Michael Grossberg^a, George Bonev^a, and Fazlul Shahriar^b,

^aCity College of New York, 160 Convent Avenue, New York, NY 10031

^bCUNY Graduate Center, 365 Fifth Avenue, New York, NY 10016

ABSTRACT

Currently, the MODIS instrument on the Aqua satellite has a number of broken detectors resulting in unreliable data for 1.6 micron band (band 6) measurements. Damaged detectors, transmission errors, and electrical failure are all vexing but seemingly unavoidable problems leading to line drop and data loss. Standard interpolation can often provide an acceptable solution if the loss is sparse. Interpolation, however, introduces a-priori assumptions about the smoothness of the data. When the loss is significant, as it is on MODIS/Aqua, interpolation creates statistically or physically implausible image values and visible artifacts.

We have previously developed an algorithm to recreate the missing band 6 data from reliable data in the other 500m bands using a quantitative restoration. Our algorithm uses values in a spectral/spatial neighborhood of the pixel to be estimated, and proposes a value based on training data from the uncorrupted pixels. In this paper, we will present extensions of that algorithm that both improve the performance and robustness of the algorithm. We compare with prior work that just restores band 6 from band 7, and present statistical evidence that data from bands 3, 4, and 5 are also pertinent. We will demonstrate that the increased accuracy from our multi-band statistical estimate has significant consequences at the product level. As an example we show that the restored band 6 has potential benefit to the NASA snow mask for MODIS/Aqua when compared with using band 7 as a replacement for the damaged band 6.

Keywords: MODIS, 1.6 micron band, statistical regression, snow-cover mapping

1. INTRODUCTION

The MODerate Resolution Imaging Spectroradiometer (MODIS) aboard Aqua and Terra provides crucial earth science data which is used widely. MODIS on Aqua has a particularly unique role because it is part of a constellation of satellites known as the A-Train. Unfortunately the 1.6 micron channel (band 6) of MODIS/Aqua suffers from severe damage. In fact only 5 of the 20 detectors are fully functional, resulting in a severe striping pattern and large gaps in the data, as shown in figure 1. While usually not so severe, damaged detectors resulting in periodic line drop and striping are common. Other classic examples include damaged imagers on Landsat 4 and 5 and more recently the water vapor (WV) 6.2 micron channel on SEVIRI.

Before one can use image processing software and higher level retrieval algorithms on striped or damaged images, the missing data must be first estimated in some principled way. While providing masks for missing or damaged data is critical, leaving the task of estimating the missing data estimation to end users is far from ideal. End users may have little or no knowledge of the best practices with which to do this estimation. This is particularly problematic because end users may only download the bands directly related to their target application even though information from other bands may help in estimation of the missing information, and removing stripes.

NASA's current algorithm for filling in the missing values of band 6 MODIS/Aqua is based on a column-wise spatial interpolation. Simple interpolation methods can create statistically or physically implausible image values and visible artifacts. Structural artifacts may not even be apparent through root mean square error (RMSE) metrics. The corruption of the image by naive interpolation, however, becomes obvious from image gradients

Further author information: (Send correspondence to Irina Gladkova)
E-mail: gladkova@cs.cuny.cuny.edu, Telephone: 212-650-6295

analysis. Since 15 out of 20 band 6 Aqua/MODIS detectors are broken or noisy, there are many adjacent missing lines so there are large spatial gaps in the data making spatial interpolation inappropriate. This limits its use in many potential applications such as snow and ice. NASA had to develop a separate snow mask algorithm for Aqua³ which relies on the information in 2.1 micron (band 7) rather than 1.6 micron (band 6).

In response to the band 6 problem, our group has developed an algorithm to recreate the missing band 6 data from reliable data in the other 500m bands using quantitative restoration techniques.⁵ Our algorithm uses values in a spectral/spatial neighborhood of the pixel to be estimated, and proposes a value based on training data from the uncorrupted pixels. Due to the non-parametric nature of the estimator, we avoid the blurring inherent spatial interpolation, which have implicit smoothness priors. Our preliminary work shows this approach not only reduces the RMSE but restores the correct gradient information as well.

In this paper we extend this algorithm improving its performances and robustness. We improve the robustness by preceding our algorithm with a step to fill or replace sparsely corrupted or missing pixels in the input data. We improve the performance by using grids of tiles which overlap. The overlapping tiles provide a smooth transition between different local models. Another improvement is that in our earlier work, we were not able to apply our algorithm to pixels near the border of the image (see⁵). In the present work we have removed that restriction extending the restoration to the boundary.

We have shown that using all the neighboring channels we are able to quantitatively estimate the value at dead or noisy detectors applied to the MODIS band 6 case. We have performed evaluation using MODIS Terra data, to estimate the potential benefit to the Aqua snow product. Since the corresponding band 6 of the MODIS/Terra has no such problem, it makes it possible to evaluate the algorithm by simulating the band 6 MODIS/Aqua damage on MODIS/Terra. We have demonstrated that NASA's snow product applied to our restored band 6 better matches the true band 6 snow product than the currently used band 7 proxy. In this paper, we will demonstrate that the use of the original NASA algorithm with the restored band 6 data generates a more reliable snow mask than the modified algorithm with Band 7 data. In particular, the restored band 6 data is much more reliable under certain conditions where the band 7 algorithm tends to run into problems, including the very important border areas between snowy regions and snow-free regions.

2. RELATED WORK

There have been a few prior approaches to restoring the missing data resulting from broken detectors on MODIS/Aqua band 6. A typical image is shown in figure 1. As discussed above, the simplest approach to replacing missing data is data interpolation. Interpolation is most effective when restoring scattered isolated missing pixels of scenes at points where it can be assumed no edges or fine details are present, but this assumption does not apply to band 6 MODIS/Aqua images because they often contain fine cloud and land structural detail. In addition, rather than the lost pixels being isolated, many contiguous lines have been dropped, leaving large gaps in the data. Because whole rows of data have been dropped, NASA publishes band 6 MODIS/Aqua with the missing data filled in through linear interpolation along columns (Figure 2). The problem with this method is that there is not enough information coming from the good 25% of the detectors to restore the missing data, thus column-wise interpolation introduces significant artifacts. Because of these artifacts, the image is smoother along columns than along rows. Since the row-wise interpolation dramatically corrupts image gradients, the image is unusable to any algorithm that uses gradients as input, such as edge detectors. A more general issue is that pure mathematical interpolation methods are based on heuristic assumptions of how missing data should be filled in; without scientific validation of these assumptions, there is no guarantee that the filled in data will be accurate.

A more appropriate approach is to treat restoring the missing data as an estimation problem. In Wang et al., 2006,¹⁰ the authors exploit this by fitting, at good band 6 detectors, a cubic polynomial expressing the band 6 pixel values as a function of band 7. They then use this polynomial to fill in the missing values. They evaluate this approach using MODIS/Terra band 6 and 7 as proxies for the MODIS/Aqua band 7 and damaged band 6. They show that by using regression to find the polynomial coefficients, they obtain a restoration that is significantly better than basic interpolation. A weakness in this approach is that given the complex spectral reflectances of materials, there is no true functional relationship between band 6 and band 7.

A local cubic regression approach was proposed by Rakwatin et al. 2009.⁹ This approach drops the unrealistically strong assumption that a global relationship between band 6 and 7 exists and instead assumes that this relationship is only true locally. Thus, they let the parameters of a functional relation between band 6 and 7 vary. To do this, they define a sliding window of pixels centered at the pixel to be filled in and then they use the working sensors of band 6, and the corresponding sensors of band 7 within the sliding window, to perform a locally varying cubic regression. This local regression is then applied to the band 7 data to restore the band 6 data at the damaged sensors. Each pixel is filled in using a cubic polynomial with a potentially different set of coefficients. Note that, although a window is used to find the coefficients of the polynomial regression, the input to the regression is just the band 7 value. In addition to allowing the regression to vary across the image, they also applied histogram matching to the radiances to further improve the consistency of the regression and simultaneously reduce striping artifacts.

3. MULTI-VARIATE ESTIMATION APPROACH

As discussed above, prior work has treated band 6 as a function of band 7. The general approach we propose is to restore each missing pixel value in the damaged band using values taken from a spectral-spatial window, composed of all the nearby bands, around that pixel. The strength of relationship between two variables is often analyzed by examining the correlation or covariance. A more sensitive method to study the linkage between variables is to consider the joint probability distribution. We write the joint probability density as $P(B_i, B_j)$, where B_i, B_j are continuous random variables representing the co-registered pixel values of bands i and j respectively. If the random variables are independent then the joint distribution is a product of marginal distributions $P(B_i, B_j) = P(B_i)P(B_j)$.

As an example, it is well known that MODIS bands 6 and 7 are closely related. They are spectral neighbors and this relationship is the basis for the restoration approaches in Wang et al. 2006¹⁰ and Rakwatin et al. 2009.⁹ The joint probability density is shown in figure 3. When this figure is contrasted with the product of the marginal distributions shown in figure 4 it is clear that there is a close relationship between the bands.

Not only is band 6 related to band 7, but it is also related to the other 500 meter bands: 3, 4 and 5. Again it is clear from the figure pairs figures 5-6, figures 7-8 and figures 9-10 that none of the band 500 bands are independent of band 6 with band 5 being the most closely related band after band 7. Even when the correspondence is weaker or more complex, the combined relationship of the other bands provides information not available from just band 7.

As an example, we will present evidence that the values from band 5 provide additional information over those provided by band 7. To do this we first estimate the band 6 value using a cubic polynomial as in Wang et al., 2006.¹⁰ The difference between this estimated band 6 value and the actual band 6 value, which we call *reflectance error*, is a proxy for the part of the band 6 value not explained by band 7. The joint probability distribution of band 5 with the reflectance error of band 6 estimated from band 7 is shown in Figure 11. If band 5 were *not* informative about this error, then the distribution would be independent, that is, it would be the product of the single variable distributions, shown in Figure 12.

The clear distinction between the distribution shown in Figure 11 and Figure 12 simply illustrates that band 5 reflectance provides new information about band 6 not provided by band 7. Similarly, as we add other spatial and spectral variables, we can improve the estimation of band 6. We carefully balance the gain in information with potential over-fitting by limiting the model to a multi-linear estimator and adjusting the window size according to results obtained on independent test data. The effectiveness of this approach is validated in Section 5.

4. QUANTITATIVE IMAGE RESTORATION ALGORITHM

The outline of the quantitative image restoration (QIR) algorithm is presented in the diagram shown in Figure 13. The damaged band to be restored, for example band 6 on Modis/Aqua, are referred to as the “bad band”. The bands which we use in the QIR algorithm as input, bands 3, 4, 5 and 7 in the MODIS/Aqua example, are referred to as “good bands.” The first preprocessing step deals with the fact that there are some scattered pixels in the good bands with missing or out of valid range values. Since good band values provide input for the QIR algorithm, pixels with values outside of the valid range will wreak havoc with regression; therefore, we preprocess

out-of-range values using an adaptive mean value filter, which replaces isolated missing pixels with the mean value of the valid pixels in a window with an adaptive size. The adaptive window size is the minimum size such that the majority of the pixels in the window are within-valid range. Note that the window is limited to a fixed maximum size. If there is are too many (more than half) bad pixels in the good bands, the restoration is aborted.

The next step when working with the data is destriping the individual bands. As observed in Rakwatin et al.,⁹ destriping can significantly improve regression. In theory an image of properly calibrated radiances should not have stripes, nevertheless some striping artifacts remain and can be removed using histogram specification as is commonly done.⁸

We describe a general QIR algorithm in which we have $K - 1$ good bands and 1 bad band with broken detectors whose data must be reconstructed. In the case of MODIS/Aqua we consider all the 500 meter bands as good bands, so we have a total of $K = 5$ bands with 4 good bands. Without loss of generality, we can renumber the bands so that the good bands are numbered 1 through $K - 1$ and the band to be reconstructed is labeled K . We write the value of a pixel $p = (i, j)$ in the i th scanline (row), the j th column, and the k th band of an image I as $I_{i,j,k}$. This pixel value can be a digital count, radiance, reflectance, or other value type. The QIR algorithm does not depend on the value type and can produce a restoration function that takes inputs and produces output in any of these types. Note, we use reflectances in this paper to enable the comparison with prior work.^{9,10} However, the calculation of reflectance uses the solar angle and intensity (data outside of band 6), and involves further processing potentially amplifying or even masking errors in the restoration. Thus, we are performing the restoration on digital counts and only then converting the values to either radiance or reflectance.

The restoration function is a composite (piecewise) function built from smaller restoration functions defined on large overlapping portions of the image, which we refer to as *tiles*. The tiles are defined by first partitioning the image into a grid of non-overlapping tiles and for MODIS/Aqua we used 200×200 pixel tiles. The grids of tiles are shifted horizontally by a half-tile (100 pixels), vertically by a half-tile, and diagonally by a half-tile, as shown at the far left of the diagram in Figure 14. As a result pixels in a corner region are in only one tile, pixels near the image boundaries but not in corner regions are members of two overlapping tiles, and pixels at least a half-tile away from the boundary are covered by 4 overlapping tiles.

We determine a restoration function independently for each tile. Since a pixel may belong to several tiles, the restored value in the bad band is the average of the independent restoration functions for each tile to which it belongs. For the pixels in the corner regions the value is determined by the corner tile restoration function and there is no averaging because there is only a single function. For the other regions, such as pixels away from or near the boundary, a restored pixel value is the average of four or two, respectively, restoration functions, one for each tile containing the pixel.

QIR could be generalized to accommodate weighed averages based on distance to the boundary of a tile, but in the MODIS/Aqua case the extra complexity was unwarranted. We also considered tiles which overlapped by more than half-a-tile. An extreme case we considered was using sliding overlapping tiles centered at the band 6 pixel to be restored, providing for a per-pixel varying restoration function. In this case it might not even be necessary to average restoration functions and could potentially improve accuracy, but it is very expensive when compared to 1/2 tile overlaps. Training data is used to find a balance between accuracy and speed, empirically, when setting the tile size and the overlaps between the tiles.

For each pixel from a scanline with a broken detector in band K , the restoration function F must provide a value $I_{i_0,j_0,K} = z(q)$, with $q = (i_0, j_0)$, where we think of $z(q)$ as the dependent variable. For each pixel q , the independent variables are taken from the values in the image I for a $m \times n$ spatial window, w_q centered at q , with m and n odd, as is shown in Figure 14. The $m * n * (K - 1)$ independent variables, of w_q are

$$\mathbf{x}(q) = \{x_0(q), \dots, x_{m*n*(K-1)}(q)\} = \{I_{i,j,k}\}_{i_0 - \frac{m-1}{2} \leq i \leq i_0 + \frac{m-1}{2}, j_0 - \frac{n-1}{2} \leq j \leq i_0 + \frac{n-1}{2}, 1 \leq k \leq K-1}. \quad (1)$$

To determine F , we separately, and independently determine an F_T for each tile T . We do this by first collecting a training set made up of the set independent and dependent values $\{\mathbf{x}(p), z(p)\}_{p \in V_T}$, with V_T being the set of all pixels, p , corresponding to working detectors in the bad band as shown at the top right of Figure 14. This is indicated in Figure 13 as the box “True band 6 Value” which is $z(p)$ and as the box “Windows in Bands”

which are the variables $\mathbf{x}(p)$ from the window w_p . Every per-tile restoration function $F_T(\mathbf{x}(p))$ has a training error defined as

$$\text{Error}(F_T) = \sum_{p \in V_T} |F_T(\mathbf{x}(p)) - z(p)|^2. \quad (2)$$

To determine F_T we would like to find a function which minimizes this error without overfitting the training set F_T . We partially address the overfitting issue by restricting F_T to one of the simplest possible families of functions, multi-linear functions of the form:

$$F_{T,\alpha}(\mathbf{x}(p)) = \sum_t \alpha_t x_t(p), \quad (3)$$

where $\alpha = (\alpha_1, \dots, \alpha_{m \times n \times (K-1)})$. Therefore, for our implementation of the QIR algorithm we determine F_T by optimizing for parameters α_t that minimize the training error $\text{Error}(F_T)$ (defined in equation 2). The optimal solution for α_t is computed using a least square solver to obtain the per-tile multi-linear regression illustrated by the box at the top right of Figure 13.

QIR for a pixel q in band K (band 6/Aqua) for a damaged detector, proceeds by first determining the tile T containing q . The per-tile multi-linear estimator F_T is then applied to the associated window w_q to obtain $F_T(w_q)$, which is the per-Tile reconstruction of $\tilde{z}_T(q)$ for the pixel q in band K , as illustrated in Figure 14. Finally, as stated above, the restored value for q , $\tilde{z}(q)$ in the bad band is the average of the per/tile values $\tilde{z}_T(q)$ that contain q . This is indicated as the box ‘‘Reconstructed Band 6 Value’’ in Figure 13.

Note that in the QIR algorithm, pixels corresponding to functioning band K (band 6) detectors help determine the multi-linear estimators F_T but are not inputs to those estimators. In other words, none of the input values $\mathbf{x}(q)$ for $F_T(\mathbf{x}(q))$ come from band 6. Thus the only way the band 6 values from the good detectors influence the value of $F_T(\mathbf{x}(q))$ is through the determination of the parameters α in the large tile T .

5. IMPLEMENTATION AND ALGORITHM EVALUATION

We implemented QIR using the Python programming language, employing a number of libraries including the Numpy and Scipy scientific computing libraries and the Matplotlib visualization library.^{11–13} In order to read the MODIS data we used the PyHDF library, which is a wrapper around the C HDF library from the HDF Group. The library was fast enough to allow direct use of the raw data from disk, even when the data needed to be converted to units of radiance.

As mentioned above, although band 6 MODIS/Aqua has extensive damage, the corresponding band 6 MODIS/Terra is functioning normally. This makes it possible to evaluate QIR by simulating the damage to band 6 MODIS/Aqua on band 6 MODIS/Terra. Our evaluation method compares restored images obtained from applying each of the algorithms on the simulated damaged bands to the original undamaged band 6 MODIS/Terra (used as ground truth). The same approach was previously used to evaluate the methods to restore band 6 using a cubic polynomial function of the values in band 7.^{9,10}

We compared our restoration with the algorithms of Wang et al., 2006,¹⁰ and Rakwatin et al., 2009;⁹ it was relatively straightforward to directly implement the polynomial proposed in Wang et al.s paper, while Rakwatin et al. graciously provided their IDL code for fair comparison. We chose 10 granules with varied terrain containing snow, clouds, mountains, and vegetation in order to challenge all algorithms. These representative 10 granules were also chosen from many others because they could be restored with the prior work implementations without failure. We were unable to run the implementation Rakwatin et al provided on many granules because it was not robust to the bad data that sometimes appears in the granules. This is only a problem with their implementation, *not* their algorithm. Since we wanted to minimize any modification to their code, we restricted the evaluation to granules on which their implementation ran smoothly.

The result of our evaluation, as seen in figure 16, was that all the algorithms do reasonably well. The results compared restoration of reflectances in order to match the outputs of the algorithms of Wang et al., 2006,¹⁰ and Rakwatin et al., 2009.⁹ QIR outperforms the other methods on all granules. As reported by Rakwatin et al. in 2009,⁹ their algorithm does consistently and significantly better than that of Wang et al., 2006,¹⁰ and

our algorithm is consistently the best of the three. It is interesting to note that while our restoration is more complex in the sense that we use more bands, it is also less complex in that we use a lower order model (linear vs. cubic). We have not yet explored combining a higher order model with multiple bands.

Since we were only able to run the cross algorithm evaluation on 10 granules, we wanted to verify that the resulting range of errors was typical for QIR. To do this, we ran the QIR algorithm on 113 MODIS/Terra granules with simulated Aqua damage. Figure 17 shows the histogram of root mean square errors (RMSE) for QIR on these 113 granules. These are the errors only for the restored pixels using band 6 MODIS/Terra in which the damage was simulated. As can be seen, the RMSE varies quite a bit depending on the image content and the mean of error for QIR was just under 0.005 and is consistent with the 10 granule sample. Nevertheless, the range shows that there is considerable variation in QIR’s accuracy depending on the granule.

One important use of the MODIS 1.6 micron data is to identify snow. Hence, NASA produces a snow mask using band 6.⁴ The determination of the snow mask requires computation of the normalized difference snow index (NDSI) which is determined by the values in bands 4 and 6. Figure 21 shows a natural color image⁶ of the Terra granule on day 321 of 2009 at 18 :00 UTC. In this granule, the snow is distributed in a mountainous area of the Rocky Mountains. The NDSI image zoomed into a region on the border between a snow area and snow-free land and is shown in figure 19. Figure 20 shows the NDSI image that was obtained using restored band 6, after first being damaged to simulate band 6 on Aqua. When compared with the original NDSI image (figure 19), it is virtually impossible to discern a difference. Figure 18 shows the histogram of errors in the NDSI based on a restored band 6. The histogram includes only “clear sky pixels” and is evaluated over a set of 10 test granules. If we look at the distribution of NDSI errors that is shown in Figure 18, we can see that most of the errors are quite small. This supports the idea that the restoration may be of high enough quality for use in the snow mask.

Besides computation of NDSI values, the determination of the snow mask also requires excluding pixels where clouds are present or where there are known bodies of water. Pixels are also eliminated from consideration if the thermal MODIS bands indicate that the pixel is too hot to be snow. In addition, because trees can occlude the satellite’s view of the snow, the normalized difference vegetation index (NDVI) is also needed. If a pixel is not excluded as just described, and the NDVI and NDSI is such that it falls in the region of the plane defined by the union of a half-plane and the interior of a polygon(cf.⁷), NASA’s snow algorithm classifies that pixel as snow, otherwise it is not. For example, figure 23 shows a MODIS/Terra snow mask (MOD10.L2). The pixels in white were flagged as snow, those in black as not-snow, and those in gray were excluded because of clouds, bad pixels, or known bodies of water.

NASA had to develop a separate snow mask algorithm for Aqua, because of the severe damage to band 6. This algorithm relies on the information in the 2.1 micron band 7 rather than band 6.³ The algorithm for determining the snow involves substituting band 7 for band 6 in computing the NDSI and changing the parameters that determine the half-plane and polygon. Band 6 was originally chosen because this part of the spectrum is particularly good for distinguishing snow from clouds. Unfortunately, in band 7 the contrast between snow and clouds is not as strong. We can evaluate the effectiveness of this band 7 algorithm by applying it to MODIS/Terra, and the result is shown in figure 24, using the same color coding as in figure 23. With the original band 6 algorithm as ground truth, it is clear from the figures 24 and 23 that the band 7 algorithm overestimates the snow, i.e. it does not often miss-classify a snow pixel as non-snow (false negative), but it often miss-classifies non-snow pixels as snow (false positive).

As before, we simulated Aqua damage on a band 6/Terra and then restored it using the QIR algorithm. We then used the QIR restored band 6 to produce the snow product, using NASA’s band 6 based algorithm, and the result is shown in figure 22. There are only small differences between the QIR based snow product and ground truth.

6. CONCLUSION

We have presented an extension of the QIR algorithm described in,⁵ which uses neighboring pixels, both spectral and spatial, to quantitatively estimate missing scan lines. In the case of MODIS band 6, evaluation shows that our results out-perform results which are based purely on band 7. We tested across granules with different surface

types well separated over time and obtained a consistent improvement. We also verified the error rates of our QIR algorithm on the granules we used to compare with the prior work by using a comprehensive sample of 113 granules.

We showed that the high quality results we obtain with QIR have the potential to improve high level products such as the snow mask. While preliminary, the results on an improved MODIS/Aqua snow product are promising. In conjunction with Dorothy Hall and George Riggs of NASA, we plan to evaluate the potential use of QIR for a MODIS/Aqua snow mask product. In addition to evaluating the snow mask, we will evaluate the NDSI as well. We also plan to create a dataset of one full snow season over the Yellowstone region that will contain a restored band 6, a 500m resolution granule that contains restored band 6, and a snow product calculated using a restored band 6.

In this paper the restoration function that used data from a local spectral-spatial neighborhood to implement the QIR algorithm was multi-linear. In future work we hope to test whether extending the QIR algorithm to use polynomial input terms can further improve the results. In addition, we hope to apply this work to other cases where bands have significant damage. The QIR algorithm we have presented and evaluated provides an approach to mitigating the risk of partial loss of a band in a multi-spectral remote sensor.

7. ACKNOWLEDGMENTS

This work was supported in part by the NOAA/NESDIS Center for Satellite Applications and Research, Grant No DG133E07CQ0077. The views, opinions, and findings in this report are those of the authors and should not be construed as an official NOAA and or U. S. Government position, policy, or decision. We also would like to thank P. Rakwatin, W. Takeuchi, and Y. Yasuoka for graciously providing us with the implementation of their restoration algorithm, and providing technical advice on how to use it.

REFERENCES

- [1] D. Hall, G. Riggs, "Accuracy assessment of the MODIS snow products," *Hydrol. Process.* 21, pp. 1534-1547, 2007
- [2] Design and Performance of the EUMETSAT MSG-2 WV 6.2 Anomaly Compensation Algorithm, Doc. No. EUM/OPS-MSG/TEN/07/0131, Issue V1, Feb. 27, 2007
- [3] George A. Riggs, Dorothy K. Hall, "Snow Mapping with the MODIS Aqua Instrument," In Proceedings of the 61st Eastern Snow Conference, Portland, Maine, June 9-11, 2004; 81-84
- [4] D. Hall, G. Riggs, V. Salomonson, "Algorithm Theoretical Basis Document (ATBD) for the MODIS Snow and Sea Ice-Mapping Algorithms," <http://modis-snow-ice.gsfc.nasa.gov/>
- [5] I. Gladkova, M. Grossberg, F. Shahriar, "Quantitative Image Restoration," Proceedings of the SPIE, Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery, April 2010, Vol. 7695
- [6] "MODIS Rapid fire response system," <http://rapidfire.sci.gsfc.nasa.gov>
- [7] Andrew G. Klein, Dorothy K. Hall, George A. Riggs, "Improving snow-cover mapping in forests through the use of a canopy reflectance model," *Hydrological Processes* 12: pp. 1723-1744 (1998)
- [8] J. S. Cheng, Y. Shao, H. D. Guo, W. Wang, B. Zhu, "Destriping MODIS data by power filtering," *IEEE Trans. on Geoscience and Remote Sensing*, 2003, Vol. 49, pp. 2119-2124
- [9] P. Rakwatin, W. Takeuchi, Y. Yasuoka, "Restoration of Aqua MODIS Band 6 Using Histogram Matching and Local Least Squares Fitting," *IEEE Transactions on Geoscience and Remote Sensing*, 2009, Vol. 47(2), pp. 613-627
- [10] L. Wang, John J. Qu, X. Xiong, Y. Xie, N. Che, "A New Method for Retrieving Band 6 of Aqua MODIS," *IEEE Geoscience and Remote Sensing Letters*, 2006, Vol. 3(2), pp. 267-270
- [11] Travis E. Oliphant, *Python for Scientific Computing*, Computing in Science and Engineering 9 (90), 2007
- [12] Fernando Perez and Brian E. Granger, *IPython: A System for Interactive Scientific Computing*, Computing in Science and Engineering 9 (90), 2007
- [13] John D. Hunter, *Matplotlib: A 2D Graphics Environment*, Computing in Science and Engineering 9 (90), 2007

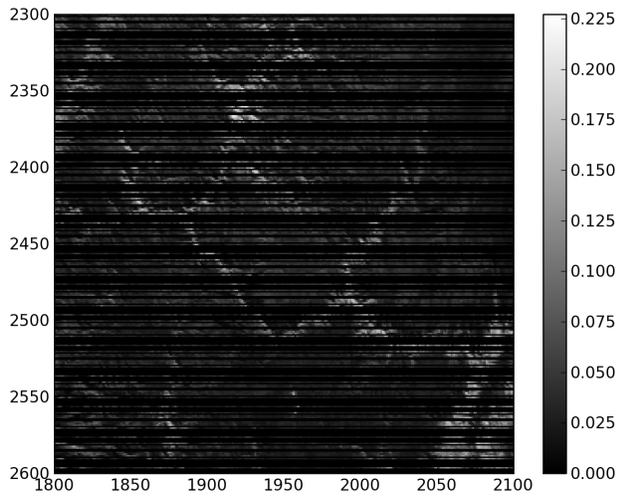


Figure 1. Band 6 of Aqua/MODIS

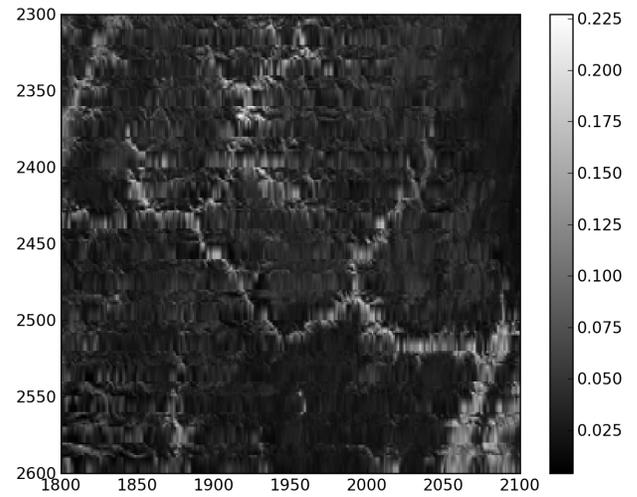


Figure 2. NASA interpolated band 6

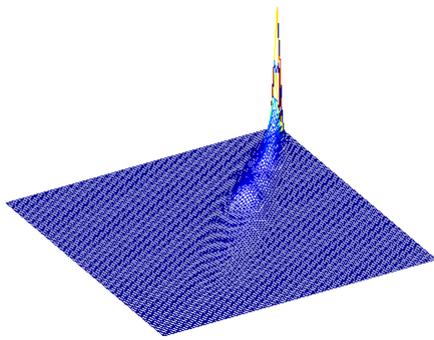


Figure 3. Joint Distribution $P(B_6, B_7)$

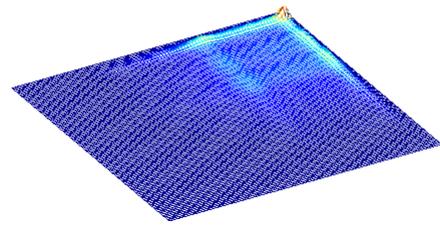


Figure 4. Product of Marginal Distributions $P(B_6)P(B_7)$

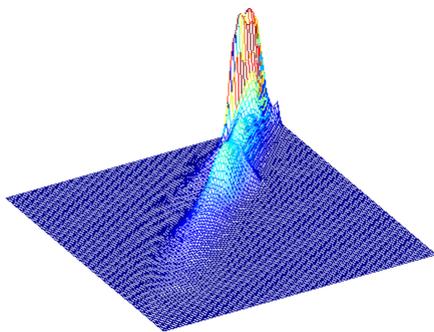


Figure 5. Joint Distribution $P(B_6, B_5)$

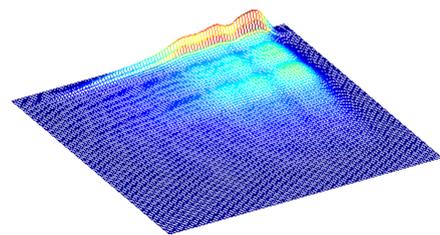


Figure 6. Product of Marginal Distributions $P(B_6)P(B_5)$

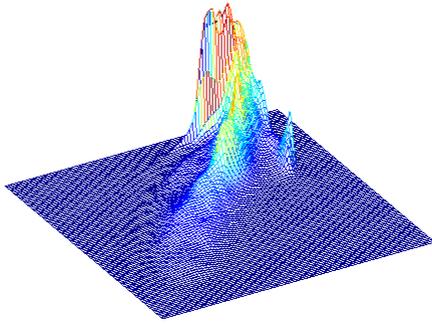


Figure 7. Joint Distribution $P(B_6, B_4)$

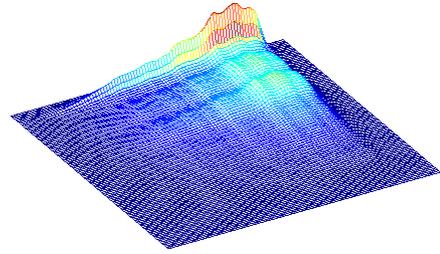


Figure 8. Product of Marginal Distributions $P(B_6)P(B_4)$

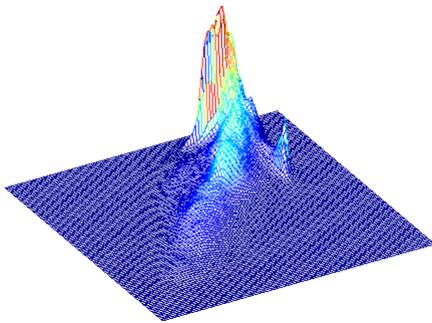


Figure 9. Joint Distribution $P(B_6, B_3)$

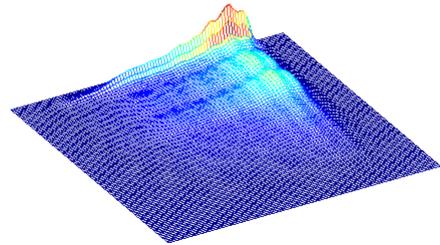


Figure 10. Product of Marginal Distributions $P(B_6)P(B_3)$

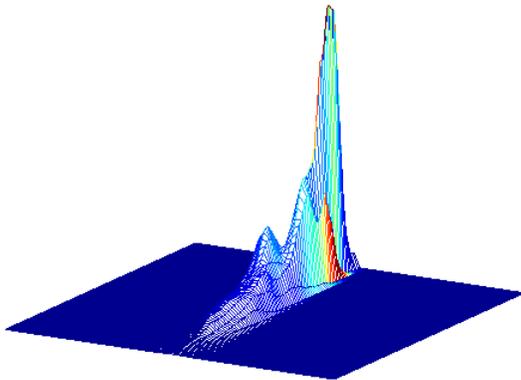


Figure 11. Joint probability distribution function (PDF) of band 5 reflectance and the residual of band 6 reflectance with the portion predicted by band 7 subtracted

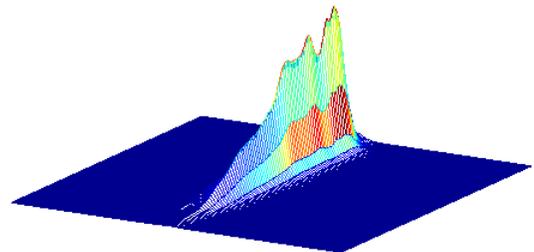


Figure 12. Product of band 5 probability and residual of band 6 radiance with the portion predicted by band 7 subtracted. The failure of this product to match the joint PDF is evidence that band 5 contains information about band 6 not captured by band 7

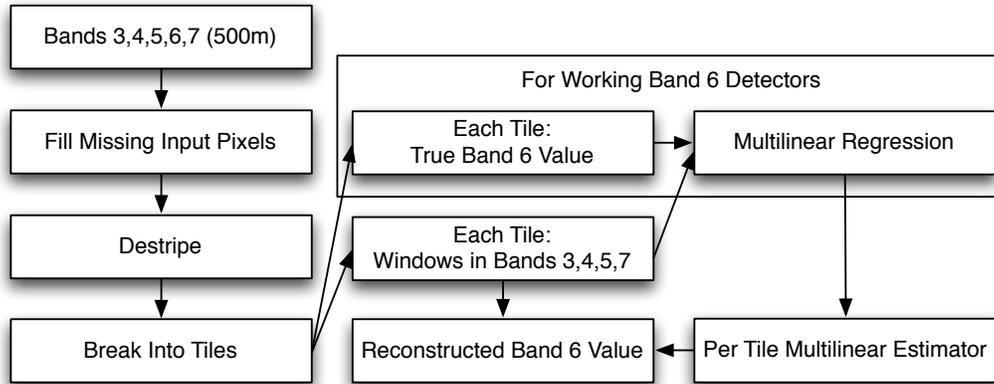


Figure 13. Block diagram showing the overall structure of the quantitative image restoration algorithm (QIR). The granule is preprocessed by filling in sparse missing values using an adaptive mean filter and then destriping. The image is then covered by overlapping tiles. For each pixel on a working-detector scanline within the target band, the value in band 6 and a spatial-spectral window of values surrounding that are taken as an example of output and respectively input variable examples for reconstruction. All such pixels within a tile are used to form a training set for that tile. The training set is used to build a restoration function for the tile using multi-linear regression. Spatial-spectral windows surrounding the pixels with broken detectors are input into the per-tile restoration functions to get a restored pixel value for each tile containing the pixel. The QIR restored value is obtained by averaging the per-tile restoration functions

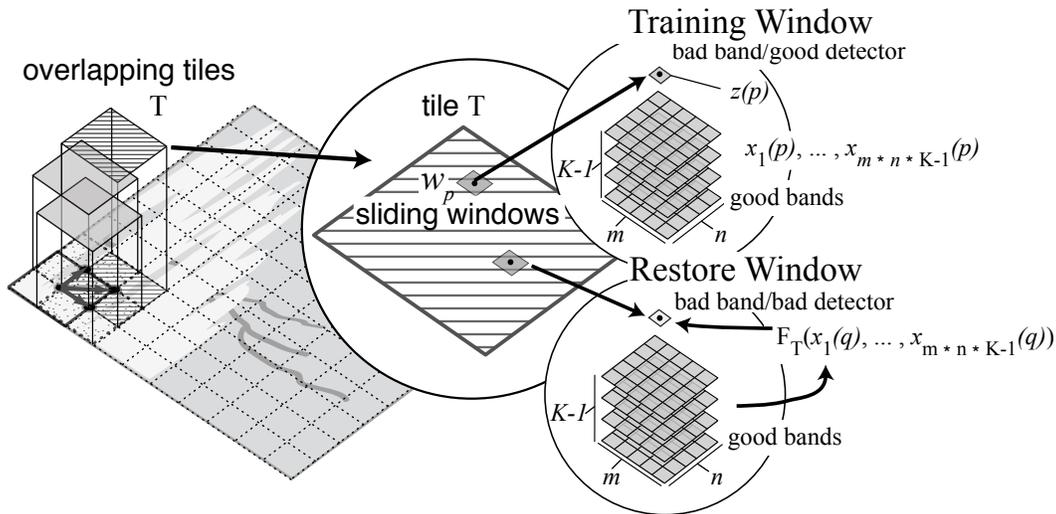


Figure 14. Diagram showing the process of determining the restoration function for a tile. On the far left, half-tile shifted tiles are shown, making up a grid over overlapping tiles. The example tile T is shifted a half-tile horizontally, vertically and diagonally. A spatial window w_p surrounds every pixel p , which is associated with a working detector in the bad band. Correspondingly, surrounding every pixel q , which signifies a broken detector, is a window w_q . At the top right of the diagram, the values in all good bands in the spatial window w_p , along with the working value in the bad band $z(p)$ give one element of a training set which is used to build a restoration F_T function for T . On the bottom left, the restoration function F_T takes the values from the good bands in w_q as input to produce a restored value for that tile. The QIR restored value for q in the bad band is the average over all the tiles T containing q of the restored values F_T at q

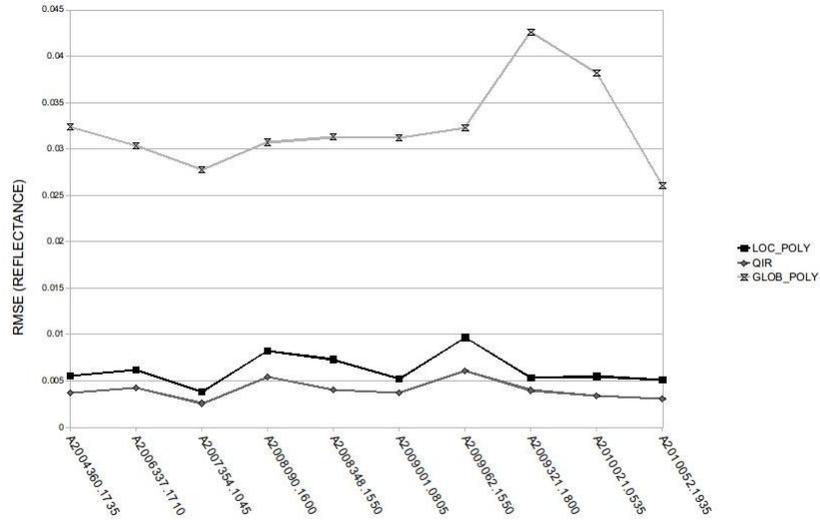


Figure 15. RMSE, in reflectances, of previous algorithms for 10 test granules.

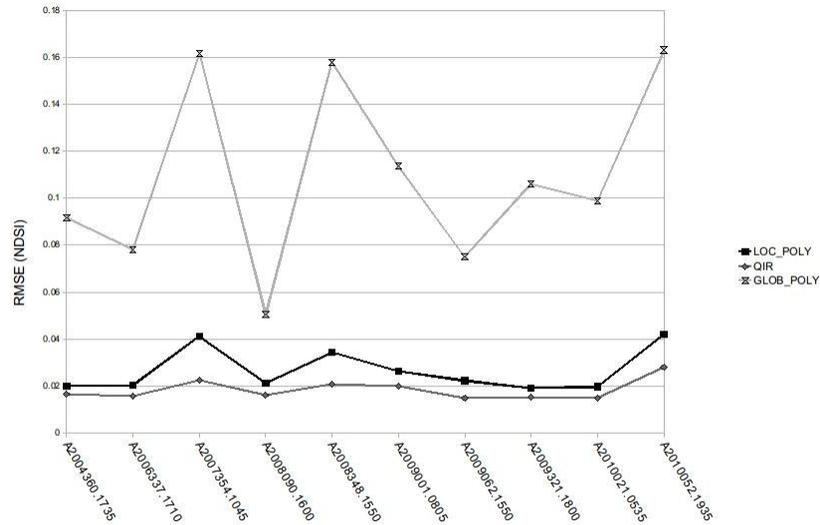


Figure 16. RMSE of NDSI for 10 test granules.

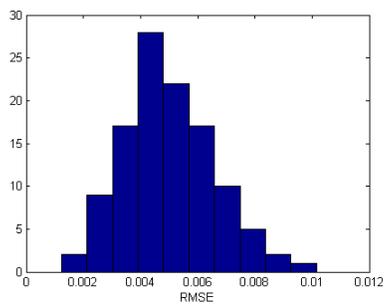


Figure 17. Distribution over 113 granules of the QIR RMSEs in reflectances.

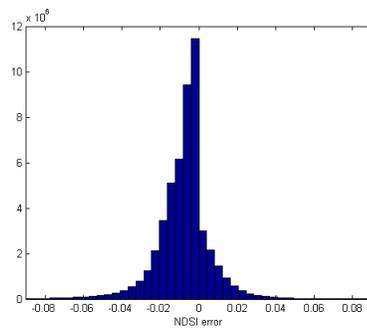


Figure 18. Histogram of NDSI errors using restored band 6 over 10 test granules.

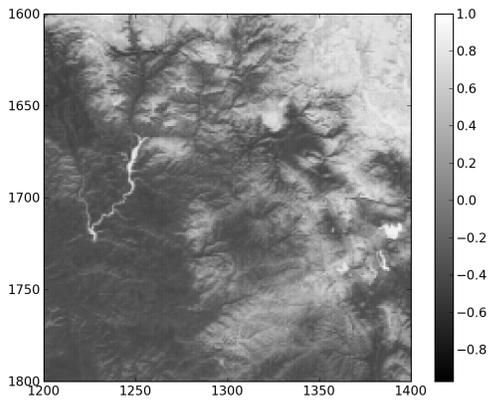


Figure 19. NDSI using original band 6

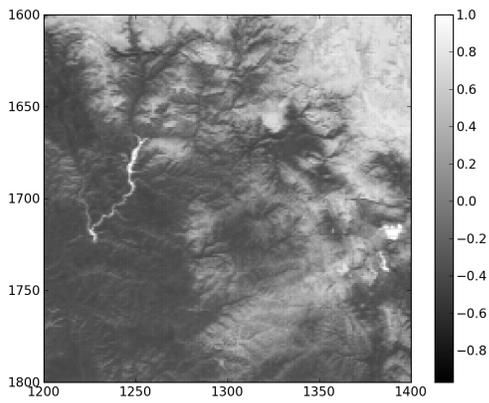


Figure 20. NDSI using restored band 6



Figure 21. 2009/321 - 11/17 at 18 :00 UTC Snow in the Rocky Mountains.⁶ In this granule, the snow is distributed in a mountainous area.

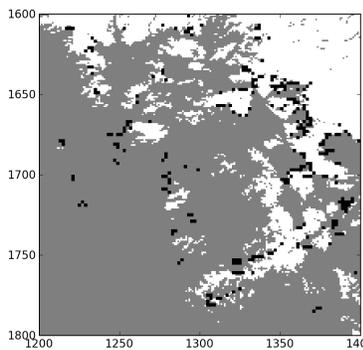


Figure 22. QIR snow mask

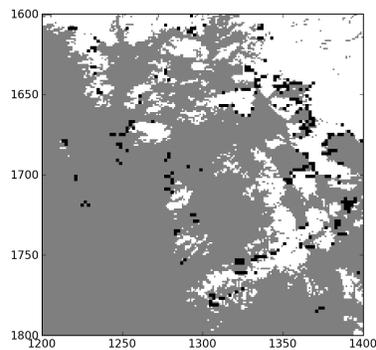


Figure 23. NASA's MOD10 snow mask

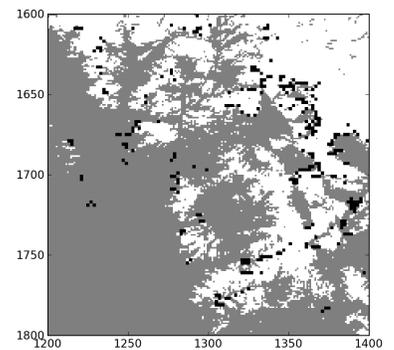


Figure 24. Band 7 snow mask