

Error Mitigation for CCSD Compressed Imager data

Irina Gladkova^a and Michael Grossberg^a and Srikanth Gottipati^a

^aCCNY, NOAA/CREST, 138th Street and Convent Avenue, New York, NY 10031, USA

ABSTRACT

To efficiently use the limited bandwidth available on the downlink from satellite to ground station, imager data is usually compressed before transmission. Transmission introduces unavoidable errors, which are only partially removed by forward error correction, and packetization. In the case of the commonly used CCSD Rice-based compression, it results in a contiguous sequence of dummy values along scan lines in a band of the imager data. We have developed a method, which is capable of using the image statistics to provide a principled estimate of the missing data. Our method outperforms interpolation yet can be performed fast enough to provide uninterrupted data flow. The estimation of the lost data provides significant value to end users who may use only part of the data, may not have statistical tools, or lack the expertise to mitigate the impact of the lost data. Since the locations of the lost data will be clearly marked as meta-data in the HDF or NetCDF header, experts who prefer to handle error mitigation themselves will be free to use or ignore our estimates as they see fit.

1. INTRODUCTION

Compression of space-based imager data is both necessary and desirable for making most efficient use of precious bandwidth. Compression is particularly important for the limited bandwidth downlink from satellite to ground station. Not only is this link the most bandlimited but it is also subject to unavoidable transmission noise. Such noise will necessarily introduce errors into the raw data stream. To reduce the number of such errors, the data is typically encoded with a forward error correcting code. This partially protects the data from error but it is impossible to prevent all errors from finding their way into the data stream.

Compression both helps and aggravates the problem of protecting the data from error. It helps indirectly because the reduced size of the compressed data stream allows for the extra room needed by forward error correcting encoding. Compression directly aggravates errors because even single bit errors can cause many more bits to contain errors after decompression. This process is called error propagation. One way to limit error propagation is to break the data into independently compressed packets. Breaking the data into packets guarantees that the worst case scenario for a single bit error is the loss of a packet. This is the strategy used by the CCSD Rice-based compression scheme developed by Pen Shu of NASA-Goddard, see .¹

The packets in the CCSD Rice-based compression are usually defined by breaking the data up along scanlines. Thus a lost, or partially lost packet results in errors in starting at pixel and continuing forward along the scan lines until the packet ends. While the true value of these corrupted pixels is lost, the identity of the corrupted pixels is known. The CCSD standard calls for the corrupted pixels to be filled in with dummy values upon decoding. It is often the case that locations of errors are also stored in the HDF or NetCDF header as well.

Leaving the data with dummy values is problematic for several reasons. It puts the burden upon the data product consumers to estimate the missing data despite the fact that they may have little or no knowledge of how to do that. It is important to note that many end users will often work with partial data, selected bands or regions of interest. This both magnifies the impact of missing data and makes it harder to accurately compute a principled estimate.

We will show that using a statistical regression, we can reconstruct the missing data with a high degree of accuracy. We will show our approach outperforms either spatial or spectral interpolation. Because our estimation will only be applied to the small amount of data, which would otherwise be lost, we will demonstrate that the throughput of the algorithm should be fast enough to insure uninterrupted data flow. It is essential to note that any estimated data will be indicated as such in the HDF or NetCDF header so it can be ignored if desired.

Further author information: (Send correspondence to I. Gladkova)
E-mail: gladkova@cs.cuny.cuny.edu, Telephone: 1 212 650 6261

2. RELATED WORK

In this section we review the field of *digital image inpainting* from the computer-vision literature which is closely related to our problem of error correction of multispectral images. *Inpainting* is the art world’s term for what researchers in image processing call “image interpolation”. The objective of inpainting algorithms is to restore an incomplete image by filling in the missing parts based on the partial observations. Sapiro gave a detailed account of the motivations and numerous applications of image inpainting and it was his group who first borrowed the term inpainting from museum restoration artists.

Inpainting algorithms can roughly be grouped into three types: variational/partial differential equation (PDE), statistical, and exemplar based methods. The variational/PDE based methods solve optimization problems constrained on the known data on the boundary of the hole. In variational methods the problem is defined through a functional of the original image and the restored image. An inpainting solution is obtained by finding the restored image which minimises the functional. This optimisation problem can be solved by using calculus of variation leading to the specific Euler-Lagrange PDE which has to be solved numerically. A few examples of this type of methods can be found in the work by Ballester et al. ² and Chan et al. ^{3,4} There also exists PDE based methods in which a PDE is defined directly ^{5,6} This type of methods are able to reconstruct regions with evidence of large scale deterministic geometric structures at the boundary of the hole, such as an intensity edge going through the hole. These methods generally work best on small sized regions. However, in the case where the surroundings of the hole consist of stochastic textures, such methods tend to produce blurred reconstructions and not to reproduce the texture.

The statistically based methods,⁷⁻¹¹ in which statistical models of the image content are constructed and used for sampling a new image content in the hole. These methods are mainly based on techniques for texture synthesis and as such work well for regions with stochastic and irregular textures. However, for regions with large scale deterministic geometric structures, such methods have a tendency to fail or converge slowly toward a good solution.

Exemplar based methods ^{12,13} search for similar image patches, usually in the image surrounding the hole, and paste a filling of the hole using the found patches. For a specific point on the boundary, exemplar based methods consider a patch consisting of neighboring pixels. The patch will consist of pixels which are either in the hole or on the boundary and the surrounding image. A similarity measure is used for comparing the known part of the current patch with other patches in the known image. The reconstructed data flows from the boundary into the hole. Usually a filling or visiting order strategy is employed. Criminisi et al. ¹³ choose the filling order based on a gradient measure on the boundary to the hole, which leads to an order where patches at high gradient edges going into the hole are visited first. Exemplar based methods combining ideas from statistical texture methods with search for exemplar patches also exist ^{14,15} Exemplar based methods can faithfully reconstruct small regions with both stochastic texture and deterministic geometric structure. In the methods by Criminisi et al. ¹³ and Efros et al. ^{14,15} once a patch has been pasted into the hole the reconstructed data cannot be changed at a later stage. The impact of this is that if a less optimal choice of patch is made, the algorithm can never recover from it. This can for instance be seen in the case where an edge is being occluded by the hole. In this case two edges will flow in from opposite sides of the hole, and this effect makes it difficult to obtain a reconstruction where the two edges meet. This issue is especially visible for large holes with large scale structures meeting at the middle, but it also occurs for small scale textures. This one-shot approach makes the choice of filling order crucial in order to obtain a good result, however this choice is non-trivial to make.

The difficulty applying prior inpainting work to the current case is that it typically does not take into account both spectral and spatial correlations. We attempted using available implementations of these methods using purely spatial inpainting, for example the inpainting algorithms implemented in OpenCV. For our data and the single line data loss scenario we consider, we found the performance of inpainting algorithms on a par with standard interpolation. In order to achieve a high quality reconstruction of the lost data, we must use all the spatial and spectral information the imager data provides.

3. MULTISPECTRAL IMAGER DATA

As a proxy for GOES-R we looked at imager data from the Meteosat Second Generation (MSG) satellite’s main payload, an optical imaging radiometer called the Spinning Enhanced Visible and Infrared Imager (SEVIRI).

SEVIRI is a 50 cm-diameter aperture line-by-line scanning radiometer which provides image data in four Visible and Near-InfraRed (VNIR) channels and eight InfraRed (IR) channels. The 12 SEVIRI channels consist of 8 InfraRed (IR) detector packages (3 detectors each), 1 High Resolution in the Visible (HRV) channel (9 detectors), 2 Visible and 1 Near-IR (3 detectors each) channels. The spectral range for the 4 visible/NIR channels is 0.4-1.6 μm , and for the 8 IR channels is from 3.9-13.4 μm . SEVIRI has unique capabilities for cloud imaging and tracking, fog detection, measurement of the Earth-surface and cloud-top temperatures, tracking of ozone patterns, and many other improved measurements. SEVIRI data granules were obtained through EUMETSTAT (European Organisation for the Exploitation of Meteorological Satellites).

SEVIRI, like the future GOES-R, have passive sensors that detect and measure energy reflected and/or emitted from the Earth's surface in different wavelengths of the electromagnetic spectrum. Different spectral bands are used to gain information about various materials on the Earth's surface, such as water, vegetation and soil, as well as the composition of the atmosphere. Because each material uniquely absorbs and reflects energy throughout the spectrum, these materials produce their own distinct spectral responses or signatures. Since these spectral responses correspond to a limited set of materials from the surface and atmosphere, the bands of the imager exhibit strong spectral correlations.

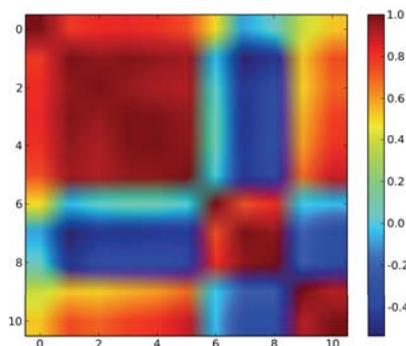


Figure 1. The above image shows the degree of correlation among the various channels of SEVIRI. It shows a greater degree of correlation among channels that belong to the same group of channels. The groups that we are looking at are the Infra-red (channels 1-6), near Infra-red (channel 7), visible (channels 8-9), and water vapor (channels 10-11).

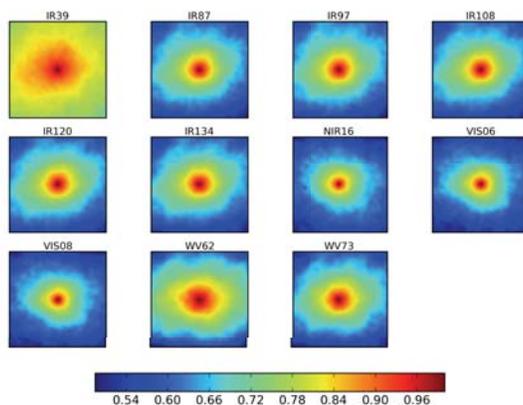


Figure 2. The above graph shows the degree of spatial correlation in each of the Seviri channels, using 50x50 spatial windows to compute the spatial correlation matrices. We found that our error correction algorithm works better when we take spatial windows as our feature vectors, rather than just individual pixels.

Figure 1 shows the spectral correlation among the various channels. There is a noticeable correlation within the 4 groups of channels and poor correlation across groups. The visible channels are negatively correlated with the other groups and the near Infra-red channel is not correlated with other channels. Figure 2 shows the spatial correlation for each individual channel. Figure 3 shows the degree of spatial cross-spectral correlation in a representative channel from the IR, NIR, VIS, and WV group of channels with respect to the remaining channels. As expected when two channels are poorly correlated, the spatial correlation of one channel with respect to the mid-pixel of the second channel vanishes rapidly as one moves far from the mid-pixel.

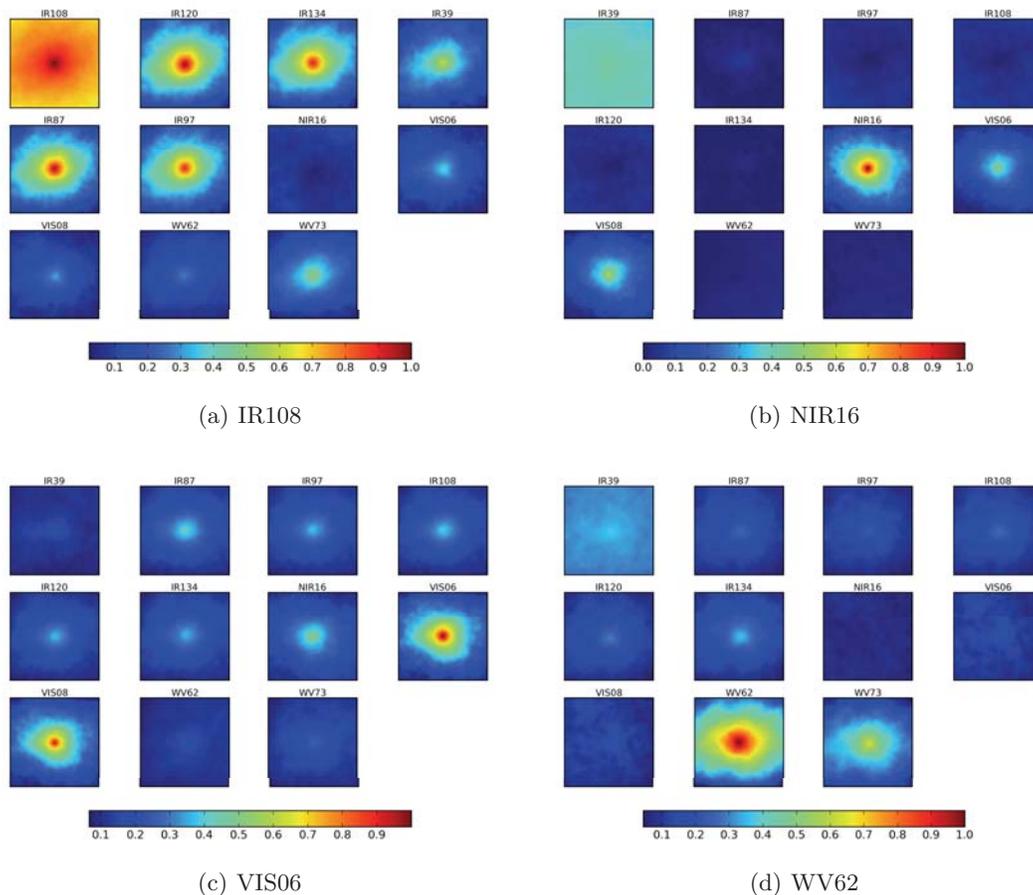


Figure 3. The above graph shows the degree of spatial cross spectral correlation in a representative channel from the IR, NIR, VIS, and WV group of channels with respect to the remaining channels. 50×50 spatial windows were used to compute the spatial cross spectral correlation matrices.

4. PREDICTION ALGORITHM

When multispectral imager data is transmitted to the ground, the signal contains noise. Forward error correcting encoding and packetization are often used to provide protection against loss during transmission. Despite this precautions, there are occasions when a packet of data will be dropped. The data is transmitted in order of spatial raster scans for each band. The consequence of the loss of a single packet is often the dropping of a scan line in a single band. In fact, it is possible that more than one scan line may be dropped, or that multiple lines in multiple bands are missing at different scan positions. Nevertheless, we make the simplifying assumption that a single scan line is lost, and we restore pixels for all columns, except those a small distance (less than a window width) from the left and right margins. These assumptions are not fundamental to the algorithm, and can also be handled easily with the same framework, however the bookkeeping of pixel indices becomes more intricate.

We assume the imager produces a multi-spectral images with K bands, I scanlines, J columns. We represent a discrete image as an integer-valued function $g(i, j, k)$, with values in the range is $[0, 2^b - 1]$, with b the bit-depth of the instruments, and $1 \leq i \leq I$, $1 \leq j \leq J$, and $1 \leq k \leq K$. With this representation, an image from the k th band is the restricted function $G_k(i, j) = g(i, j, k)$. From this point of view, $G = \{G_k\}_{k=1}^K$ can be considered as a stack of 2D images, one for each band k .

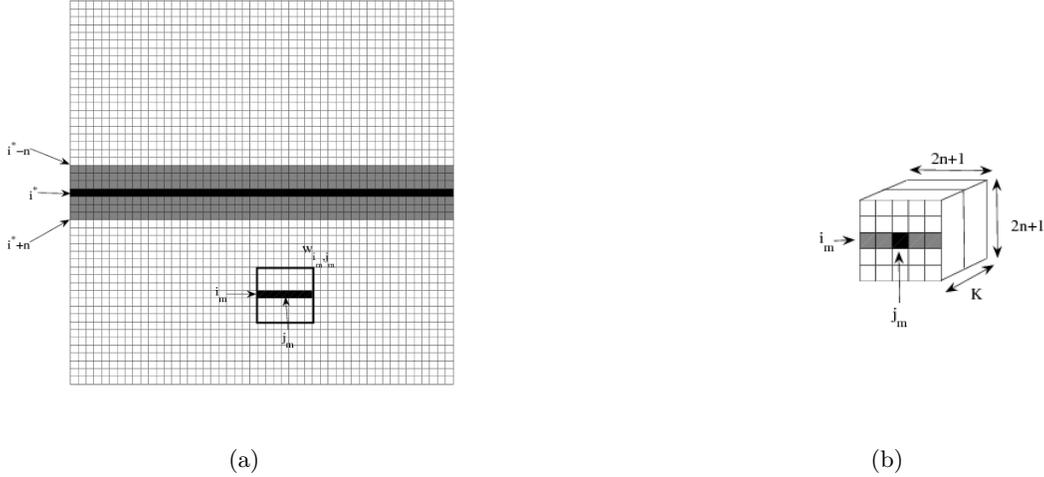


Figure 4. Diagram of the spatial arrangement of pixels used for prediction (a) as well as a view multi-spectral/spatial neighborhoods considered for prediction (b) The dark black line in the row i^* shown in (a) is taken to be the row of lost data. Windows of size $s \times s \times K$, with $s = 2n + 1$, centered along the lost line (shown as a gray strip) are considered to decide how to fill in the missing value. Windows of the same size but with complete data (no loss) (a,b), and center at (i_m, j_m) within n scan lines of the missing line (a) are used fit a regression to predict the missing values along i^* .

The data loss takes the form of a fixed missing row $G_k(i^*, j)$, for a fixed row denoted i^* , corresponding to the missing packet. (cf. Figure 4 a). Given the strong correlations in both the spectral and spatial domain, our approach is to predict each missing value $g(i^*, j, k)$ along the row i^* using a spatial and spectral neighborhood surrounding the missing pixel. For computational simplicity and without loss of generality we assume the neighborhood is a square window of width and height s pixels along i, j and includes all bands k . We take s to be odd so that it can be written $s = 2n + 1$ for a positive integer n . For notational convenience we can reorder the bands, without loss of generality, so that the missing row appears in the first band image G_1 in the considered stack G . The multi-spectral window is shown as a block in Figure 4 b.

The statistical regression strategy works by collecting other windows in the granule, which do not include the missing row to train the regression. Training on all windows in the granule is both computationally expensive and does not yield good performance since a distant window may contain a substantially different scene. For this reason we only consider images within a large but fixed number of pixels from the missing line.

The correlations in nearby the spectral and spatial values the window centered at the location (i_m, j_m) provide correlations which can be used to predict center pixel as confirmed by looking at Figure 3 in the previous section. We represent the local context as a window \mathbf{w}_{i_m, j_m} defined by

$$\mathbf{w}_{i_m, j_m} = \{g_{i, j, k} | i \in S^0(i_m, n), j \in S(j_m, n), 1 \leq k \leq K\}, \quad (1)$$

where (i_m, j_m) represent a middle pixel, $S(i_m, n) = \{i_m - n : i_m + n\}$ is an index set of $2n + 1$ consecutive integers ranging from $i_m - n$ to $i_m + n$, $S^0(i_m, n) = S(i_m, n) \setminus i_m$ is a set $S(i_m, n)$ without middle index i_m , see Figure 4 b.

Figure 5 illustrates another evidence that statistically based prediction can be effective for estimation of missing pixels. It shows the first two principal components of vectors of values on the surrounding window \mathbf{w}_{i_m, j_m} along two horizontal axis and the value of g at (i_m, j_m) along the vertical axis. The blue cloud of points in the figure seems to form a 2-D surface with the green cloud as the projection in the first two components. This indicates that if values in the surrounding pixels are known, for example a the red points in the plane formed by the PCA components, up onto the blue surface, will give good predictions for the true mid-pixel values. More formally, we use a linear regression scheme to discover that relationship based on the collection of windows from the rest of the image, where the value of the middle pixel is available. We will denote the set of such windows by

$$\{\mathbf{w}_{i,j}, i \in S(i^*, M) \setminus I(i^*, n), j \in \{J_0 : J_1\}\}, \quad (2)$$

where M is a number of rows above and below the missing row i^* , and J_0, J_1 are the first and the last column of the patch that will be used for estimation of the function in question.

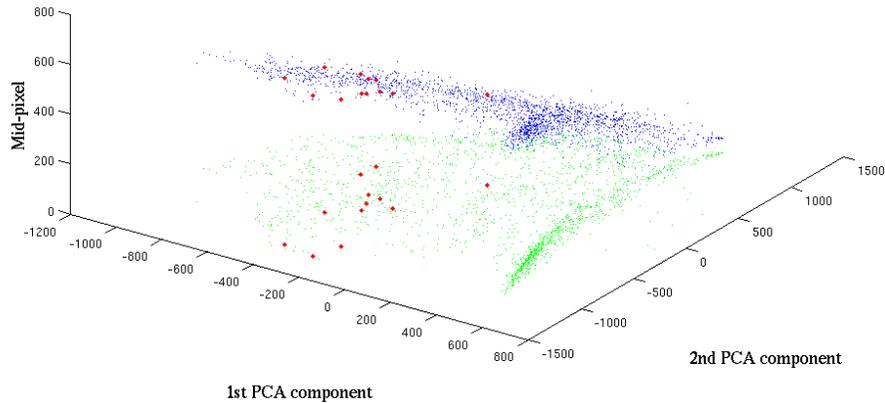


Figure 5. The above graph illustrates evidence that statistically based prediction can be effective for estimation of missing pixels. The Z axis corresponds to a value of a pixel of interest in an image where pixel loss is simulated from SEVIRI data. The X and Y axis are projections in the two largest principal components of pixel values in a neighborhood of the simulated lost pixel value. The point cloud along a surface indicates that given a training set modeling the surface, the missing value (Z) can be modeled by projection of the known X and Y values onto the surface.

Given a set $\{g_{i,j}, \mathbf{w}_{i,j}, i \in S(i^*, M) \setminus I(i^*, n), j \in \{J_0 : J_1\}\}$ of middle pixels $g_{i,j}$'s and corresponding windows $\mathbf{w}_{i,j}$'s over the area that omits the strip of the width $s = 2n + 1$ around the missing row, we fit a least squares best fit linear relationship between the value of the middle pixel g_{i_m, j_m} and the values of the window \mathbf{w}_{i_m, j_m} in the least square sense, i.e.

$$\arg \min_{\boldsymbol{\alpha}} \sum_{\substack{i_m \in S(i^*, M) \setminus I(i^*, n) \\ j_m \in \{J_0 : J_1\}}} |g_{i_m, j_m} - (\boldsymbol{\alpha}, \mathbf{w}_{i_m, j_m})|^2, \quad (3)$$

where (\cdot) is a dot product and $\boldsymbol{\alpha}$ is a block of the mixing parameters:

$$\boldsymbol{\alpha} = \{\boldsymbol{\alpha}_{i,j,k} | i \in S^0(n+1, n), j \in S(n+1 : n), 1 \leq k \leq K\} \quad (4)$$

that are arranged in the 3D block the same way as $\mathbf{w}(i, j)$'s. The mixing parameters α are the solution of the least squares problem that determine the prediction function. The values of the estimated parameters α depend on the window size, as well as the width M and length $J_1 - J_0$ of the patch from which training windows are chosen to do the least squares fit. The parameters we used were set experimentally. The restriction of windows to a patch around the missing pixel is simply a heuristic. It is clearly possible to attempt to restrict the windows based on a similarity function on the values surrounding the window. This will be considered in future work.

5. RESULTS

In this section we analyze the merits of the error mitigation algorithm that we have developed. In order to compare our algorithm's performance, we have run some standard interpolation schemes on the data by artificially introducing some errors. The results are shown in Figure 6. Of the four schemes (linear, cubic, bilinear, and spectral) we have tested, we found that the cubic interpolation scheme did the best in terms of the Root Mean Squared error (RMSE).

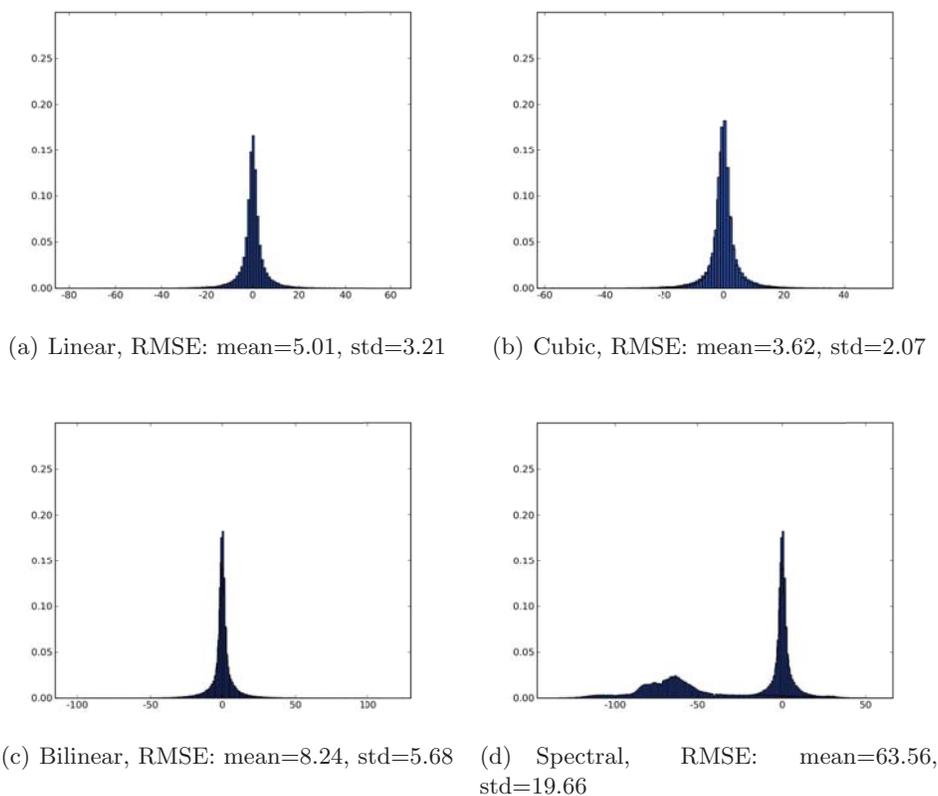
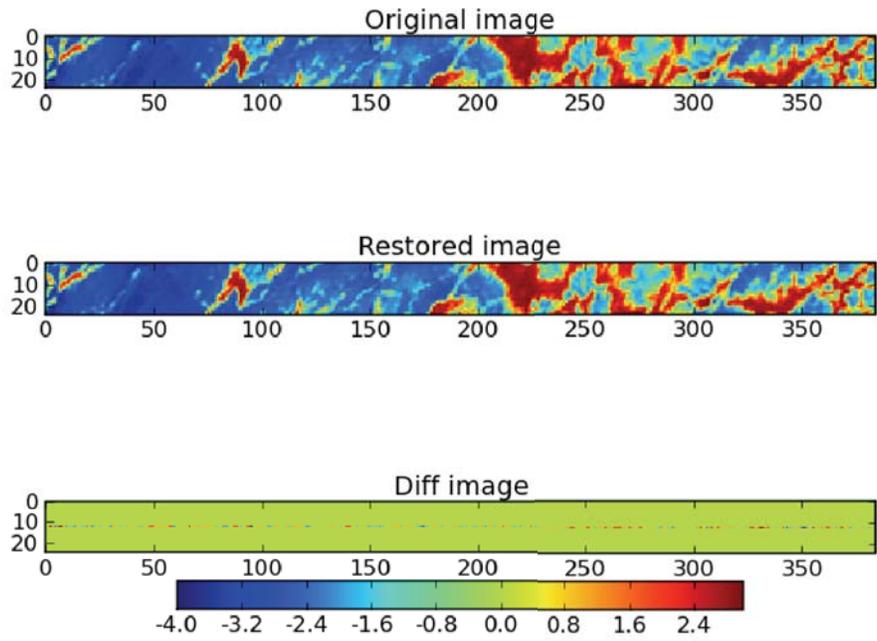
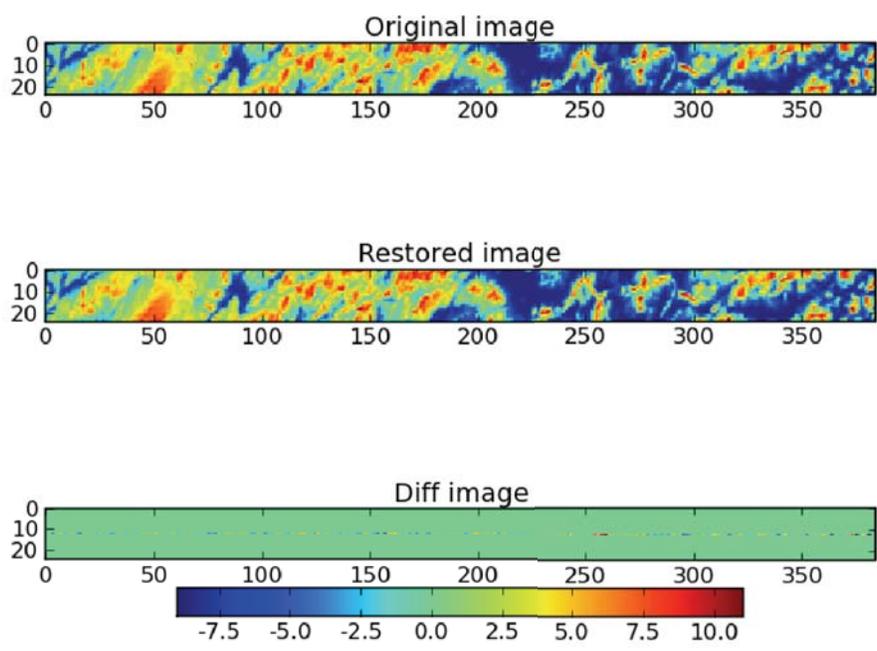


Figure 6. The above graphs show histograms of errors for various standard spatial and spectral interpolation schemes. The results are shown on SEVERI images with simulated scanline drop. The results for spatial and spectral interpolation alone show cubic interpolation has better results with respect to RMSE. It is important to note that in this case we only considered a single dropped scanline.



(a) Infra-red channel: IR108



(b) Visible channel: VIS06

Figure 7. The above figure shows the restored images using our algorithm assuming a missing scanline from the original SEVIRI images. The diff images show the magnitude of the errors upon restoration.

Figure 7 shows the restored image and the diff image for a subimage of a SEVIRI granule. The mid-row of this subimage is restored and compared with the original image. Though the restored image looks very similar to the original image, small errors exist. It will be interesting to see the magnitude of RMSE for the restored scanline for a large sample of images. Figure 8 shows histograms of errors for 4 different channels, each are picked at random from the Infra-red, near Infra-red, Visible, and water vapor channel groups, respectively. We have run our algorithm on 1000 sample subimages for each channel. The algorithm restores the water vapor channels to perfection. However, it does poorly on the near Infra-red channel which is expected, since it is poorly correlated with other channels.

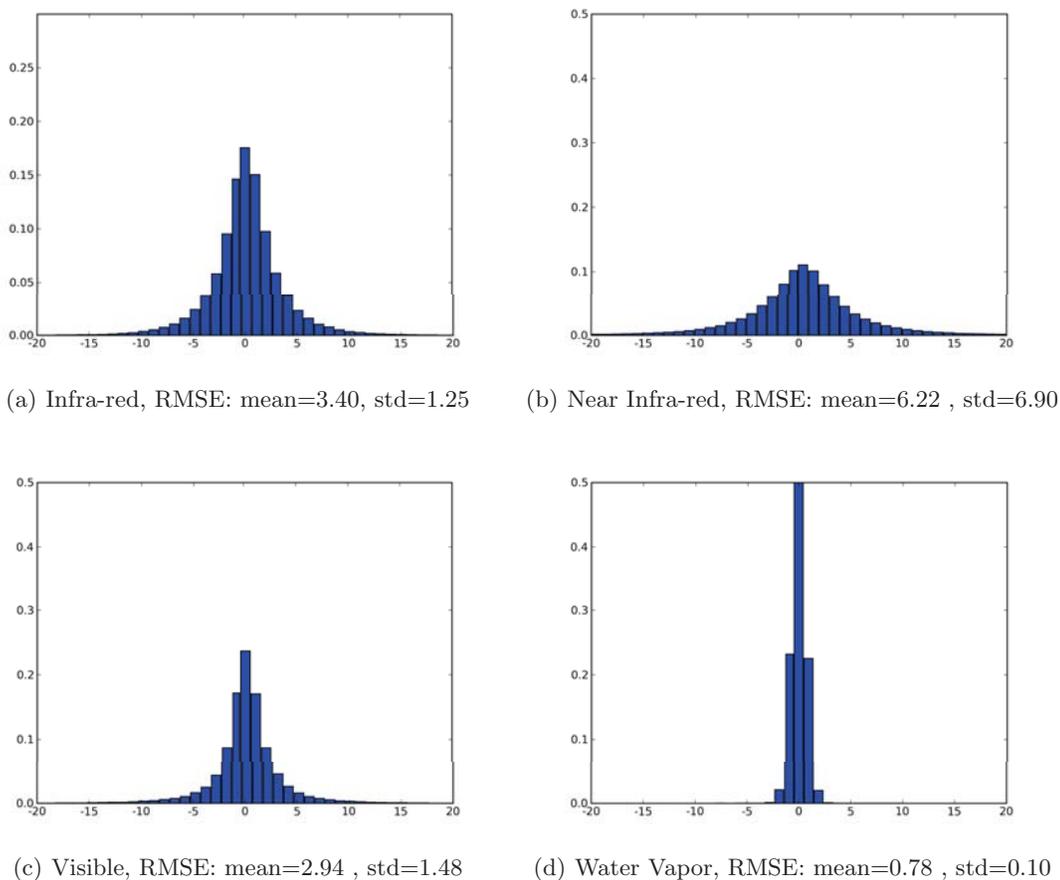
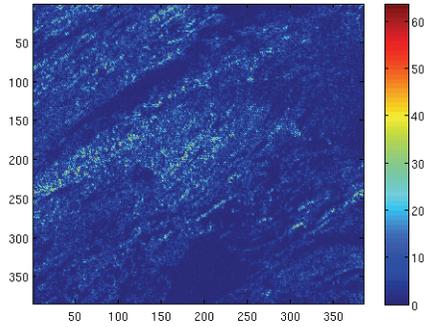
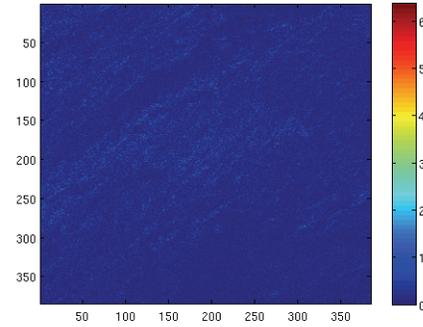


Figure 8. The above graphs show histograms of errors for four sample channels that belong to IR, NIR, VIS, and WV groups of channels using the linear regression algorithm. It can be seen that our algorithm fills the missing scanlines for the water vapor channels to perfection and does significantly well than the standard interpolation schemes on the other channels.

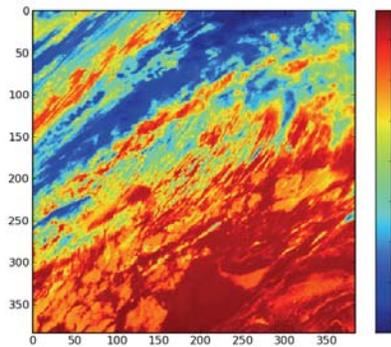
Furthermore, it should be interesting to see how the image structure affects restoration using our algorithm. Figure 9 shows the errors of an image restored where each scanline is assumed to be a missing scanline and restored sequentially. A sample SEVIRI image is restored using the cubic interpolation algorithm as well as the linear regression algorithm. Figure 9 shows that the diff image with cubic interpolated restoration still has a lot of structure left to be predicted, while, the diff image with linear regression restoration shows the algorithm has predicted much of the image structure. Figure 10 shows the performance of our algorithm with varying window size in terms of the RMSE measures. 7×7 window size seems to be the optimal window size for our restoration algorithm.



(a) Cubic: diff image



(b) Linear Regression: diff image



(c) Original image

Figure 9. The above figure shows the errors of an image restored where each scanline is assumed to be a missing scanline and restored sequentially. The mean RMSE and standard deviation across scanlines of the diff image produced by cubic interpolation method are 6.16 and 1.73, respectively, and while using linear regression method are 1.79 and 0.35.

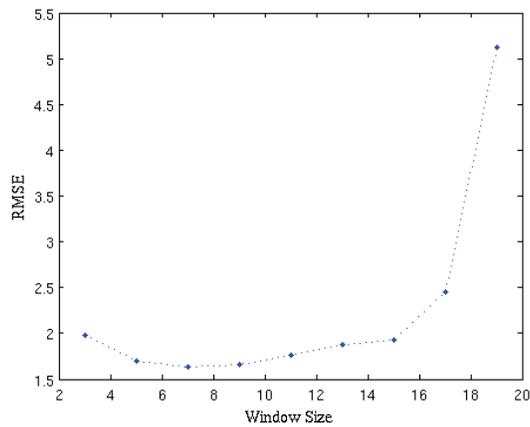


Figure 10. The figure shows the RMSE curve with respect to varying (square) window size. The figure shows that the optimal RMSE is achieved at window size 7×7

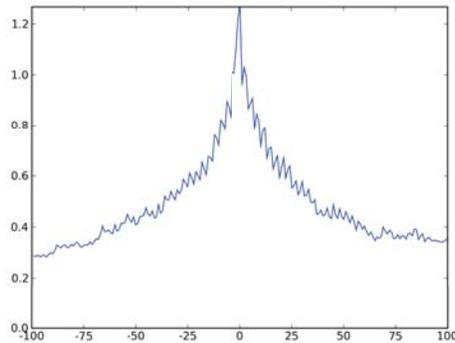


Figure 11. The above figure shows the histogram of nearest neighbors across the rows from windows whose mid-pixel belong to the center scanlines of a 1000 random 200×200 subimages of SEVIRI. The histogram looks zagged since SEVIRI has three detectors onboard for Infra-red channels.

A local neighborhood is taken as database for the purpose of efficiently computing good representatives for the missing patch in a probabilistic sense instead of using the global image. We need to estimate the right size of these local neighborhoods so that we find good representatives. Since we find our representatives using a statistical-based approach we would like to know the distribution of distances from a window with a missing pixel. Figure 11 shows the histogram of nearest neighbors across the rows from the center row of a given 200×200 subimage. One is least likely to find good representatives as one moves farther from the center scanline. Hence, it suffices to choose a local neighborhood around the missing scanline for prediction.

6. CONCLUSION

We have presented a novel but elegantly simple algorithm which provides a good estimate of missing or corrupted data. Our algorithm far outperforms the standard interpolation techniques currently used by NASA to fill in missing detector scanlines in multi-spectral imagers such as MODIS. Also, unlike currently used spatial estimation methods, the statistical regression method we proposed exploits the redundant information present in other spectral bands of the imager. Due to the non-parametric nature of the estimator, we avoid image blurring inherent in spatial interpolation methods, which have implicit smoothness priors. This method could be used as a general algorithm to predict missing data in multi-spectral images used in scientific studies. Furthermore, the computational efficiency of our algorithm is ideal for the purposes of error mitigation in large volume data transmission. Further study needs to be done to evaluate the performance of our algorithm in presence of more than one contiguous missing scanlines. In a future paper, we would like to evaluate our algorithm on other existing multi-spectral imager data.

REFERENCES

- [1] Yeh, P.-S., Xia-Serafino, W., Miles, L., Kobler, B., and Menasce, D., "Implementation of ccscs lossless data compression in hdf," in [*Earth Science Technology Conference*], (June 2002).
- [2] Ballester, C., Bertalmio, M., Caselles, V., Sapiro, G., and Verdera, J., "Filling-in by joint interpolation of vector fields and gray levels," *IEEE Trans. Image Proc.* **10**(8), 1200–1211 (2001).
- [3] Chan, T. and Shen, J., "Mathematical models for local nontexture inpainting," *SIAM J. Appl. Math* **62**(3), 1019–1043 (2002).
- [4] Chan, T., Kang, S., and Shen, J., "Euler's elastica and curvature-based inpainting," *SIAM J. Appl. Math* **63**(2), 564–592 (2002).
- [5] Bertalmio, M., Sapiro, G., Caselles, V., and Ballester, C., "Image inpainting," in [*Proceedings of SIGGRAPH*], *ACM*, 417–424 (2000).

- [6] Bertalmio, M., Bertozzi, A., and Sapiro, G., “Navier-stokes, fluid dynamics, and image and video inpainting,” in [*Conference on Computer Vision and Pattern Recognition*], *ACM* **1**, 355–362 (2001).
- [7] Gustavsson, D., Pedersen, K., and Nielsen, M., “Geometric and texture inpainting by gibbs sampling,” in [*Swedish Symposium on Image Analysis (SSBA '07)*], (2007).
- [8] Gustavsson, D., Pedersen, K., and Nielsen, M., “Image inpainting by cooling and heating,” in [*Scandinavian Conference on Image Analysis*], Ersboll, B. and Pedersen, K., eds., *Lecture Notes in Computer Science* **4522**, 591–600, Springer, Berlin (2007).
- [9] Zhu, S. and Mumford, D., “Prior learning and gibbs reaction-diffusion,” *IEEE Trans. Pattern Anal. Mach. Intell.* **19**(11), 1236–1250 (1997).
- [10] Zhu, S., Wu, Y., and Mumford, D., “Minimax entropy principle and its application to texture modelling,” *Neural Comput.* **9**(8), 1627–1660 (1997).
- [11] Zhu, S., Wu, Y., and Mumford, D., “Filters, random fields, and maximum entropy (frame): Towards a unified theory for texture modeling,” *Int. J. Comput. Vis.* **27**(2), 107–126 (1998).
- [12] Criminisi, A., Perez, P., and Toyama, K., “Object removal by exemplar-based inpainting,” in [*Conference on Computer Vision and Pattern Recognition*], *CVPR'03* **2**, 721–728 (2003).
- [13] Criminisi, A., Perez, P., and Toyama, K., “Region filling and object removal by exemplar-based image inpainting,” *IEEE Trans. Image Proc.* **13**(9), 1200–1212 (2004).
- [14] Efros, A. and Leung, T., “Texture synthesis by non-parametric sampling,” in [*Proceedings of 7th International Conference on Computer Vision*], 1033–1038 (1999).
- [15] Efros, A. and Freeman, W., “Image quilting for texture synthesis and transfer,” in [*Proceedings of SIGGRAPH*], (August 2001).