International Journal of Image and Data Fusion
2012, 1–23, iFirst

Taylor & Francis
Taylor & Francis Group

A radiometrically-accurate super-resolution approach to thermal infrared image data

Christopher G. Hughes[a]* and Michael S. Ramsey[b]

[a]*Caltech/NASA Jet Propulsion Laboratory, Pasadena, CA 91109, USA;*
[b]*Department of Geology and Planetary Science, University of Pittsburgh, Pittsburgh, PA 15260, USA*

(*Received 9 April 2012; final version received 6 July 2012*)

Super-resolution is an image processing and analysis technique used to improve the original (or native) spatial resolution of data. Super-resolution approaches have commonly sacrificed radiometric accuracy for visual appeal or vice versa. The results presented here are a significant modification and improvement of an algorithm originally applied to the thermal infrared (TIR) data from the Earth-orbiting Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) instrument. The algorithm has been programmed in the Interactive Data Language (IDL) scripting but could be easily coded in other programming languages. The primary focus of these modifications was to adapt the algorithm for use with visible and TIR data from the Mars-orbiting Thermal Emission Imaging System (THEMIS) instrument in addition to other improvements, such as a user-defined Point Spread Function (PSF) using an alpha notation. In addition, the previous requirement for an intermediate spatial/spectral resolution dataset has been removed after determining it to be unnecessary for accurate and visually pleasing results. This super-resolution approach is now more transparent to the user, and provides data from the intermediate steps, which allows for more accurate analysis of the results. The super-resolved TIR data from both the ASTER and the THEMIS are radiometrically accurate, interpretable, reproducible and maintain an excellent qualitative appearance.

Keywords: super-resolution; remote sensing; TIR; ASTER; THEMIS

1. Introduction

Super-resolution is a process for obtaining a spatial resolution greater than that of the original (or native) resolution of the data. A variety of techniques that fuse the original data with an additional higher-resolution data source have been used. In addition, super-resolution can be performed by calculating areas of overlap between different data. The most commonly applied technique is pan-sharpening, in which a single high-spatial resolution channel is used to enhance multi-spectral lower-resolution data (Pohl 1999, Zhukov *et al*. 1999, Aiazzi *et al*. 2002, Garzelli *et al*. 2004, Zhang 2004, Wang *et al*. 2005). The Multisensor Multi-resolution Technique (MMT) (Zhukov *et al*. 1999) works with multi-spectral data at both higher and lower spatial resolutions. However, in all of these techniques a trade-off has been noted between those that are the most visually appealing and the ones that are most radiometrically accurate (Zhukov *et al*. 1999). It is rare to find a

*Corresponding author. Email: topher@caltech.edu

ISSN 1947–9832 print/ISSN 1947–9824 online

http://dx.doi.org/10.1080/19479832.2012.711377
http://www.tandfonline.com

super-resolution methodology that is both quantitatively accurate and qualitatively acceptable. The algorithm (Tonooka 2005, Hughes and Ramsey 2010) modified and examined within this work was developed specifically for this purpose, and with these trade-offs in mind. It produces radiometrically accurate, interpretable and reproducible results while maintaining a good qualitative appearance.

The super-resolution technique presented here is a significant modification of an algorithm (Tonooka 2005) that was originally applied successfully to multi-spectral resolution data from the Earth-orbiting Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) instrument (Yamaguchi *et al*. 1998). The primary focus of these modifications was to enable the algorithm to process data from the Mars-orbiting Thermal Emission Imaging System (THEMIS) instrument (Christensen *et al*. 2004) as well as other improvements/testing. Other instruments, with higher spatial resolution in different spectral regions than the THEMIS instrument, have shown a need for higher-resolution Thermal Infrared (TIR) data than that provided natively by THEMIS, such as for investigating small-scale mineralogical deposits. However, it is unlikely that any TIR instrument with a higher spatial resolution than THEMIS will be sent to Mars in the near future. The modified super-resolution algorithm may help meet the need for this higher-resolution TIR data. The THEMIS instrument has five (four effective) VNIR bands between 0.425 and 0.86 μm with a spatial resolution of either 18 m/pixel; ASTER has three bands between 0.52 and 0.86 μm with a spatial resolution of 15 m/pixel. ASTER has five TIR bands between 8.125 and 11.65 μm and a spatial resolution of 90 m/pixel; THEMIS has 10 bands (eight effective for the surface) between 6.78 and 14.88 μm with a spatial resolution of typical 100 m/pixel. Older ASTER scenes have an intermediate set of data, with six bands between 1.6 and 2.43 μm and a spatial resolution of 30 m/pixel. This instrument degraded in quality, and then produced no good data after May 2008. As a result, ASTER is now similar to the THEMIS instrument in spatial and spectral resolution, although the data from the THEMIS instrument is organised and processed differently than ASTER. As a result of these changes in ASTER and the desire to super-resolve THEMIS data, significant modification of the algorithm was necessary. These modifications include changes to permit the algorithm to use data with only two different spatial resolutions, while retaining the ability to make use of intermediate spatial/spectral data if present, and to expand the range of acceptable input data for the instrument. Other modifications included performance improvements and the use of a different clustering method within the algorithm. Significant testing was performed to assess the impact of these modifications and ensure that the modified algorithm produced the same results on test data as the original method. The algorithm, as described, has been programmed in the Interactive Data Language (IDL) scripting, but could be easily implementable in other programming languages depending on the need/application. Previous work (Hughes and Ramsey 2010) describes the algorithm and provides a flowchart through various steps; this work focuses on changes from the original (Tonooka 2005).

2. Modification of the algorithm

The original technique (Tonooka 2005) is a 10-step process, in which the second five steps use the same process as the first five steps but different input data. In the first five steps, the algorithm uses ASTER visible/near infrared (VNIR) data to super-resolve the ASTER shortwave infrared (SWIR) data. Steps 6–10 repeat the process to super-resolve the

ASTER TIR data using VNIR and the newly super-resolved ASTER SWIR data. In adapting this algorithm to apply to the current ASTER configuration, which lacks SWIR data, and the THEMIS instruments' two-resolution data, it was necessary to simplify the algorithm to the first five steps. Super-resolution of the ASTER data with SWIRS bands using the old 10-step algorithm and the SWIR data, and the new five-step algorithm, while ignoring the SWIR data, produced identical results across multiple scenes. Changes from the original algorithm (Tonooka 2005) were made within each step. The steps and how they vary from the original are documented below.

2.1 *Convolution with the point spread function (Step 1)*

In the first step, higher resolution data are convolved with the instrument's Point Spread Function (PSF) to produce a dataset with the same pixel size as the lower resolution data. The PSF describes the amount of blurring, due to contributions from the adjacent area for each pixel, created by the instrument optical design. The PSF is instrument-specific and is independent of the surface being imaged. In an ideal instrument, 100% of the signal sensed within any given pixel would originate only within the area on the surface corresponding to that pixel. In actual instruments, some percentage of the sensed signal for any given pixel originates within the neighbouring pixels. In many instruments, the PSF is either symmetrical or assumed to be symmetrical in both the X and Y axes. The original algorithm defines the PSF of the ASTER instrument as a two-dimensional Gaussian, with a value determined from pre-flight measurements. In modifying the algorithm, the PSF specification has been converted to use the alpha notation described in Townshend *et al.* (2000). This provides a more generalised approach for PSF usage, and enables application to other instruments. A single value, α, is used to define the shape of PSF of the instrument, with different equations used to define the percent contribution of each neighbouring pixel; the total contribution of the pixel and the surrounding neighbour pixels is 100% (Townshend *et al.* 2000). The ASTER PSF's α is calculated to be 6.56×10^{-2}, based on the Gaussian equation and the specified standard deviation given in Tonooka (2005), equivalent to 75.5% of a pixel's signal originating from within that pixel (Table 1).

Step 1 is implemented as three actions. First, the high-resolution image is resized to match the number of pixels in the low-resolution image. This is accomplished using the IDL frebin function (available from http://idlastro.gsfc.nasa.gov/), which in turn makes use of the IDL rebin or congrid functions for image resizing. Next, the resized data are convolved with the PSF using an IDL convolve function, from the same site. Finally, as

Table 1. The ASTER PSF used in this algorithm, based on the Gaussian equation in (Tonooka 2005).

0.0043	0.0570	0.0043
0.0570	0.7546	0.0570
0.0043	0.0570	0.0043

Notes: This PSF is based on an alpha value of 0.065, and shows 75.5% of the signal recorded for a pixel originates within the spatial area of that pixel, and the rest originates within the surrounding pixels.

the data were in Digital Number (DN) format originally, the degraded resolution data are rounded.

2.2 *Identifying homogeneous pixels (Step 2)*

The second step creates a homogeneous pixel map. The original algorithm (Tonooka 2005) does this by comparing the degraded resolution pixels and their components (original, high-resolution) pixels. If the standard deviation of the component pixels (in each band) is less than a specified threshold, then degraded pixel is considered homogeneous. In Tonooka (2005), the threshold is defined as the average of the standard deviations over the whole image of each band. Based on this definition, the code calculates a standard deviation for each band. Next, these values are averaged together to create the threshold for homogeneity. The component pixel's standard deviation is then calculated for each band. If the greatest standard deviation is still less than the threshold value, that pixel is marked as homogeneous. It is not possible to convolve the edge pixels, as some of the neighbouring pixels used as input to the process are missing. As a result, all edge pixels are assumed to be non-homogeneous. The original algorithm (Tonooka 2005) did not explicitly define the edge pixel handling behaviour, so this assumption may be a difference between the two.

2.3 *The cluster tree (Step 3)*

The third step generates a spectral tree by clustering the homogeneous pixels, and then further clustering the co-located low-resolution data within each cluster. Clustering is a means of grouping together the data measured in multiple ways, such as spectral bands, such that each cluster contains members are more similar to each other than to other data. In both the original (Tonooka 2005) and this version of the algorithm, the distance measurement used for clustering is the Mahalanobis distance instead of the Euclidian distance. The MD measures the difference in variance and correlation between bands of data (Mimmack *et al*. 2001). The MD can be thought of as equivalent to the Euclidian distance for multi-dimensional data, except that it takes into account the differences of scale along each axis and discounts dimensions that are highly correlated. The original implementation uses the MD in conjunction with the K-Means clustering algorithm to create the cluster tree. Because the K-Means algorithm requires an *a priori* assumption as to the number of clusters in the data, we have implemented the ISODATA algorithm (Ball and Hall 1967) instead, modified to use the MD.

The ISODATA algorithm is similar to the K-Means algorithm, without requiring any assumptions about the final number of clusters. In both cases, a number of initial cluster centres are provided to the algorithm. For ISODATA, choosing too few cluster centres initially will lead to a large number of clusters being split; choosing too many centres initially will lead to a large number of clusters being joined. Either case increases the run-time of the algorithm, and in extreme cases may prevent the algorithm from converging to a solution in a reasonable amount of time.

In the ISODATA implementation within the modified algorithm, after the user has defined the number of initial clusters, this value is increased by 10%. This extra 10% improves the distribution of initial end-members throughout the entire data-space, rather than focusing along the edges, by permitting the elimination of clusters too

near each other. A random pixel is selected from the degraded data. This pixel becomes the centroid for the first cluster. The MD between this pixel and every other pixel is then calculated. The pixel with the largest MD is selected as the second centre. The third cluster centre is the pixel with the largest combined MD to the first and second clusters; the fourth cluster centre has the largest combined MD to the first three centres and so on. In the case of two or more pixels having the largest value in common, the first of these is chosen. If the MD to any existing cluster centre is 0 then that pixel is rejected as a new cluster centre, preventing the same spectral composition from being selected multiple times.

After all potential centres have been calculated, the distances between each are measured. The 10% additional cluster centres are removed by finding the cluster centre pairs that are closest to one another, and removing the first member of those pairs. This de-emphasises data variability near extreme data values, and leads to a more even distribution of the centres across the data space while not adding significantly to the algorithm's processing time. After the additional centres have been removed, all data are clustered within the ISODATA by seeding the algorithm using the remaining cluster centres. Subsequent iterations recalculate the cluster centres based upon the values of the data within each individual cluster.

2.4 *Assignment of super-resolved values (Step 4)*

During step 4, initial DN values are assigned to each super-resolved pixel. As in the original algorithm, these values are selected from both the data tree created in Step 3 and from nearby homogeneous pixels. The user is required to define a spatial radius, in the number of low-resolution pixels, to check for the best fit. The algorithm finds all homogeneous pixels within that radial distance and calculates the MD to each pixel, resulting in the selection of the spectrum that has the lowest MD within that area. After the best fit within the map is found, the data are then compared to the cluster tree. This is done by finding the minimum MD between the co-located high-resolution pixel and the high-resolution cluster centres. The associated low-resolution sub-clusters are then compared to the original spectrum. If the minimum MD from these comparisons is less than the MD from the map, the spectrum of the sub-cluster centre is used.

In data where three spatial resolutions are available, the MD is calculated by assigning fractional values to the high- and middle-resolution MDs. Weighting between spectral regions is defined by a user input value between 0 and 1. This value, W, is multiplied by the higher-resolution MD, and 1-W is multiplied by the intermediate-resolution MD. The default recommended W value of 0.7 from Tonooka (2005) reflects the greater importance given to matching the higher spatial resolution data. The MD of the intermediate data measures differences within only the pixels co-located with the higher resolution cluster; small differences within the limited data are magnified beyond the MD that would be calculated over the entire data.

2.5 *Radiometric correction (Step 5)*

No significant changes from the original algorithm were necessary in the final step. After allocation of spectra in the previous step, obtained values are only best-fit data. In order to maintain radiometric accuracy, the newly super-resolved data are degraded back to the original resolution with the PSF for comparing with the original data. A per-pixel

correction factor is then calculated by finding the difference between the original and degraded super-resolved data. The correction factor is allocated according to the MD of the super-resolved pixels. Super-resolved pixels with high MD are spectrally less similar to their source than lower MD pixels, and are therefore more likely to be incorrect, while a super-resolved pixel with an MD of 0 is a perfect match. By allocating the correction according to the MD, pixels that are less similar to their sources receive greater correction, and a pixel with an MD of 0 would receive no correction as long as the adjacent pixels had an MD value. If the radiometric correction were allocated equally across the component super-resolved pixels, previously good pixels would be over-corrected while poorer-fitting pixels, as compared to their source, would receive an inadequate correction. The resulting correction is then added to the values created in Step 4.

3. Testing methodology

3.1 *Datasets*

Three different sets of data were used for the performance and the behaviour analysis of the algorithm. The first set consisted of ASTER L1B data, granule id AST_L1A.003:2006463023, acquired on 16 May 2000 over the Ishioka City, Japan (Figure 1). These data were the same as that used in Tonooka (2005), and were chosen for comparative purposes. The subset used for super-resolution covers an area of 9.9×8.1 km. Where pixel locations are specified, pixel 0, 0 in the upper left of the super-resolved area was located at pixel 2496, 2460 of the full scene VNIR data. One significant difference from Tonooka (2005) is that the ASTER SWIR bands were crosstalk corrected prior to the application of the modified algorithm. ASTER SWIR bands suffer from crosstalk contamination between detectors due to a stray light error in which incident light to band 4 propagates to other bands via multireflection (Iwasaki and Tonooka 2005). A correction for this problem was not available at the time of the original super-resolution implementation.

The second data-set consisted of THEMIS data I33902002 (TIR) and V33902003 (VIS), and provided coverage of a putative chloride deposit on Mars (Figure 2). This deposit is located along the southern portion of a crater near 180.5° E, −27.0° N, in northwestern Terra Sirenum, with the chloride unit at 180.46° E, −27.25° N. The unit covers an area of roughly 15×10 km. These data were converted to emissivity and atmospherically corrected using the previously described methods (Bandfield *et al.* 2004, Hughes and Ramsey 2010) prior to the application of the super-resolution algorithm.

The third set is an artificially-generated terrain, referred to as TerrGen, using an alpha value of 1.4645×10^{-1} (or 50% of the pixel energy originates from that pixel), two end-members and a two-fold difference in the spatial scale between the high- and low-resolution pixels. Both the high- and low-resolution data have four bands. End-member spectra were created and selected separately for each resolution, with arbitrary unit-less values between 0 and 1000. As the super-resolution algorithm looks only for spatially co-located patterns within data, what the data actua × lly measures is irrelevant. The algorithm works equally well on raw DN, radiance, emissivity or even non-spectral measurement values, with the caveat that the end-user must ensure no other factors can influence patterns in one data set over the other. As an example, the ASTER and THEMIS data are both atmospherically corrected to minimise contributions from the intervening distance.

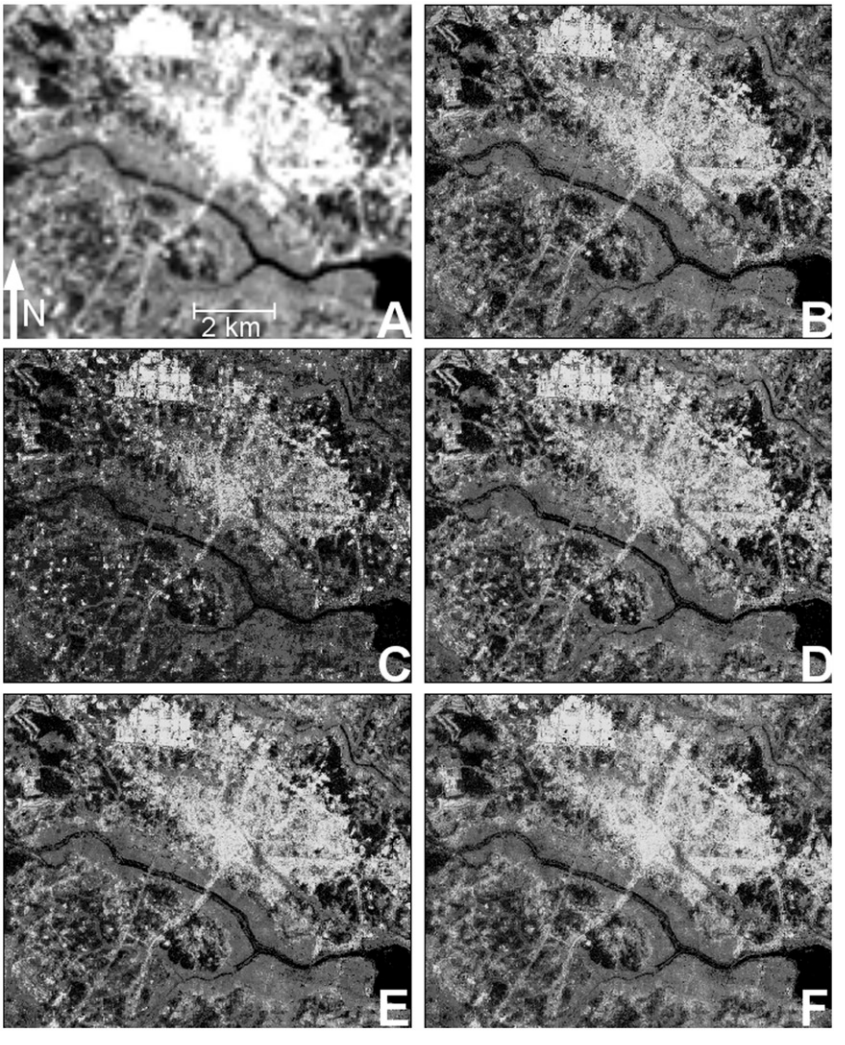

Figure 1. The original data compared to super-resolved data using different alpha values for calculation of the PSF. All images are ASTER Band 14, and are histogram matched to image A for easy comparison. (A) The original Band 14 data, clipped to the area of super-resolved data, with a linear 2% stretch. (B) The super-resolved product with an alpha of 0.06565 (75.5%), which is the correct alpha value to use for ASTER and produces both the most radiometrically accurate and clearest overall result. (C) The super-resolved product with an alpha of 0.5, or 0% of the pixel's energy originates from within its spatial area. (D) The super-resolved product with an alpha of 0.25 (25%). (E) The super-resolved product with an alpha of 0.14645 (50%). (F) The super-resolved product with an alpha of 0.0 (100%).

TerrGen pixels were allocated randomly at a low-spatial resolution, with a 30% chance of being purely end-member 1, a 30% chance of being purely end-member 2 and a 40% chance of being a randomly mixed pixel. This split was used to provide a sufficient number of each possible combination for mixed pixels. In mixed pixels, each sub-pixel had an equal chance of being either end-member. Mixed sub-pixels could have easily been created in this

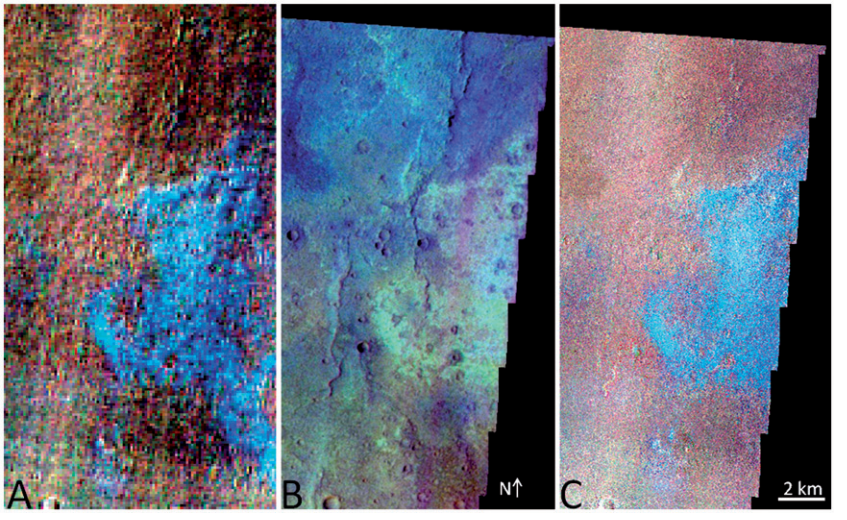

Figure 2. The original high- and low-spatial resolution THEMIS data compared to the super-resolved data. (A) The original resolution (108 m/pixel) THEMIS TIR bands 8/7/5 emissivity data shown in RGB. (B) The original resolution (36 m/pixel) THEMIS VNIR bands 4/3/2 radiance data shown in RGB. The black areas along the top and right of the figure are regions of no data, as the VNIR instrument has a smaller image footprint than the TIR data. (C) Super-resolved (36 m/pixel) THEMIS TIR bands 8/7/5 emissivity data shown in RGB. The black bars along the top and right side of the figure are regions that could not be super-resolved due to a lack of co-located high resolution data.

process as well, but only end-member pixels were used to simplify measurement and interpretation of results. During this process, a high-resolution map and a 'super-resolution' map were created. Pixels from the 'super-resolution' map were convolved with the PSF for the creation of the low-resolution map, leading to a large number of possible final spectra for any given low-resolution pixel. Five possible types of resultant spectra can be created just within the space of the low-resolution pixel: purely end-member 1, mostly end-member 1, mixed evenly, mostly end-member 2 and purely end-member 2. Spectra were then modified during convolution with the defined PSF, in which half of the pixel value is derived from the adjacent pixels. As an example of how many ways any individual pixel's spectrum could be modified by the surrounding pixels, ISODATA clustering of this dataset found 245 separate clusters in less than 100 iterations rather than two (end-members) or five (pixel composition types) that may be initially expected.

The performance of each of the five steps is determined by an user input. An initial run was performed for each data-set, using the default values (Table 2) for all variables. Default values were based on those provided by Tonooka (2005) for the ASTER data, and used as a guide for the THEMIS and TerrGen data. During testing of these variations for each step, one variable was systematically altered with all non-tested values held fixed at their default value. Output data were then compared with each other and with data from the initial super-resolution results of each data-set. This was done by examining the behaviour of the super-resolution algorithm during data processing, the quality of the output in terms of image appearance and the statistical distribution of the data as compared to the original non-super-resolved data.

Table 2. The default values used for variables within the super-resolution algorithm.

Variable	ASTER data	THEMIS data	Artificial data
Alpha	0.06565	0.1	0.14645
V_T_Distance	10	10	10
V_S_T_Distance	10	n/a	n/a
Weight	0.7	n/a	n/a
ISODATA Variables			
Initial VIS Clusters	50	50	50
Initial SWIR Clusters	10	n/a	n/a
Initial TIR Clusters	5	10	10
ChangeLimit	0.50%	0.50%	0.50%
MaxStdDev	4	4	4
MinDistance	2	2	2
MaxPairs	4	4	4
MinMembers	0.01%	0.01%	0.01%
VIS BandMax	255	0.012	1000
SWIR BandMax	255	n/a	n/a
TIR BandMax	4300	1.05	1000
Limit	100	100	100

Notes: During testing, one value was allowed to vary in a systematic fashion, whereas all other values were held fixed.

3.2 *Convolution with the PSF*

To test the effect of incorrect alpha values on the super-resolution process, data were super-resolved using five different alpha values. Each dataset was super-resolved with an alpha-equivalent value to 0%, 25%, 50%, 75% and 100% of a pixel's signal originating from within the area of that pixel. In cases where the correct alpha value is close to one of these values, the correct value was used instead: the ASTER data were processed with an alpha value equivalent to 75.5% instead of 75% and the THEMIS data were processed with an alpha value equivalent to 50.6% instead of 50%.

3.3 *Identifying homogeneous pixels*

Homogeneous pixels can be defined in a number of different ways. The original algorithm (Tonooka 2005) defines a single threshold value as the average of the standard deviations of each band for the entire image. To test the impact of other definitions, the code was modified to test the following threshold definitions: a single value is defined as the average of the standard deviations of the super-resolved area or area-of-interest, multiple values with each band being compared to the standard deviation of that band across the entire image and multiple values with each band being compared to the standard deviation of the super-resolved area or area-of-interest. For the ASTER data, the super-resolved area matched with that used in Tonooka (2005). For artificial data, the area-of-interest was defined as the middle 250×500 low-resolution pixels. For the THEMIS data, the super-resolved area was defined as the area immediately surrounding the putative chloride deposit.

3.4 *The cluster tree*

Step 3 represents the most significant change from the original algorithm (2005). There are several user-defined values within the ISODATA process, but the two with the largest impact on the final results are the initial number of clusters used and the maximum number of iterations. These variables define how far the system can start from a final, stable cluster map (not known a priori) and still reach that state. Other values, such as the minimum number of members per cluster or the maximum number of clusters to join in any iteration, can impact this process, but do not have the same controlling effect. As a result, testing of step 3 was done by varying these values one at a time. The initial number of clusters was varied between five and 1000, and the number of iterations was varied between 10 and 10,000. The super-resolved ASTER data results were also compared with Tonooka (2005).

3.5 *Assignment of super-resolved values*

The main variable impacting the result of step 4 is the radius used for searching the nearby homogeneous pixels. By default, the program uses a radius of 10 low-resolution pixels, which is the same value used in Tonooka (2005). Tests were performed by setting this value to 5, 15 and 20 pixels. Data were also super-resolved using a value of 0 (all values from the tree). This permitted the examination of the influence of the cluster-tree on the super-resolution process.

3.6 *Radiometric correction*

The correction factor calculated within this step is allocated among the super-resolved sub-pixels based upon their MD. This can be interpreted in two different ways, with both being tested. In the first method, the difference between the original and degraded super-resolved data was divided by the convolved MD of the super-resolved pixels, and this amount was multiplied by the MD of each super-resolved pixel to create the correction. In the second method, the difference between the original and degraded super-resolved data was divided by the sum of the MD of the associated super-resolved pixels, and this amount was multiplied by the MD of each super-resolved pixel to create the correction. An additional third method was tested, in which just the difference between the original and degraded super-resolved data was used as the correction.

The clearest way to illustrate the impact of the correction factor allocation can be seen by examining a single low-resolution pixel and its associated super-resolved sub-pixels. ASTER band 10 data from Earth is presented in DN, whereas the THEMIS band 4 data show emissivity. Checking both data types allows a comparison of the effects of the relative scale; ASTER DN are generally several orders of magnitude larger than their associated MD values, THEMIS emissivity are of the same general magnitude. The pixels whose upper left super-resolved sub-pixel was located at 240, 240 were examined in both cases. By examining a single pixel and a single band, it is feasible to show how the correction factor was developed, and how it was then allocated using different methods.

4. Results

4.1 *Convolution with the PSF*

The DN and the calculated calibrated radiance of the super-resolved ASTER data shows the impact of choosing an incorrect PSF (Figure 1). If the correct alpha value is used, the mean calibrated radiance of the super-resolved data is closest to that of the original data compared to the other datasets. As the alpha value moves away from the correct value in either direction, there is an increase in the mean calibrated radiance values. Similarly, the correct PSF produced the fewest number of VNIR clusters, with 100 clusters found in ISODATA. As the alpha value used for calculating the PSF moves away from the correct value, a greater number of clusters are found.

The impact of an incorrect PSF on THEMIS emissivity data is less clear. All five average image spectra are plotted on the original average image spectrum. The average spectral difference between the data super-resolved with any PSF and the original data was 4.0×10^{-5}. The largest difference between the original and the super-resolved average spectrum in any one band is 1.2×10^{-3}, with total spectral differences across all bands ranging between 9.0×10^{-5} and 4.0×10^{-3}. The TerrGen data shows similar results, with average data values similar to the original average image spectrum.

4.2 *Identifying homogeneous pixels*

The threshold used for determining whether a degraded high-resolution pixel is homogeneous or not is defined in Tonooka (2005) as the band average of the spatial standard deviation over the whole image of each VNIR band. This was calculated as a single value by averaging the standard deviation of each band together. For the ASTER data, the average standard deviation of the entire scene for bands 1, 2 and 3 N is 32 DN (rounded to the nearest integer value). For the sub-scene, it is less than half that amount (15 DN). As expected from this being an average, the threshold value is greater than the standard deviation of at least one band in both cases. Where the threshold values are determined by the band instead, the thresholds become 35, 25 and 37 (for the whole image) or 12, 13 and 20 (for the sub-image) for each of the ASTER VNIR bands. The difference between 12 DN (band 1, area of interest only) and 32 DN (the average of the standard deviations of the entire scene) is significant; 12 DN represents less than 5% of the total possible data range (0–255 DN), whereas 32 DN is 12.5%. The default threshold value (32 DN) identifies 82,442 pixels convolved to the intermediate spatial resolution (93% of the scene) and 6991 pixels convolved to the low-spatial resolution (70% of the scene) as homogeneous.

The THEMIS data had a default threshold of 1.43×10^{-4} emissivity units, or 2.4% of the data range, with the standard deviation of band 4 being one quarter and band 5 being one half of this value. If only the area of interest is considered, this single threshold decreases to 1.09×10^{-4} emissivity, with an equivalent drop in the standard deviation across all bands. The default value selects 74,393 pixels convolved to the lower resolution as homogeneous, or 11% of the scene.

The difference between per-band and the single value can be most clearly illustrated within the TerrGen artificial data. Where a single threshold value is used from the entire image, the calculated value of 380.8 represents over one-third of the data range. As the TerrGen data are randomly distributed, unlike a natural surface, there is no significant difference between the threshold values of the whole image and the image subset. Ignoring

the effect of the PSF, this does correctly select only those pixels that are purely one of the two end-members. For those pixels in which there was an even contribution from both end-members, three bands exceed this threshold; whereas the pixels in which the contribution is mostly from one end-member have two bands exceeding the threshold. However, as the impact of the PSF is considered, contributions from surrounding pixels work to homogenise the data. As a result, the super-resolution program found 322,648 homogeneous pixels instead of the 300,000 homogeneous pixels created by TerrGen. This represents 4.5% of the total scene being incorrectly selected.

4.3 *The cluster tree*

During processing of both the ASTER and the THEMIS data, a limitation within the IDL environment was encountered. Within some operating systems, IDL array creation of over 800 megabytes (MB) is problematic, with problems commonly occurring before this size is achieved. Within the version of IDL used for this work (32 bit), there is also a size limit of 2.0 gigabytes (GB) for an array for any operating system. There were tens of thousands of homogeneous pixels in both data sets, and cluster assignment is done by creating a floating point array (4 bytes per array cell) with dimensions equal to the number of clusters on one axis and the number of homogeneous pixels on the other in order to track MD between pixels and cluster centres. This array rapidly exceeded the memory allocation or the memory-addressing capability of IDL, and prevented testing of very large values of the initial number of clusters (i.e. all homogeneous pixels as their own cluster centre), and limited the number of iterations possible for moderately large starting values.

The artificially-generated data were not examined within this test due to these memory limitations. With over 322,000 homogeneous pixels, there was a low limit to the number of possible clusters. An initial 500 clusters reached this limitation through splitting before 10 iterations. For any larger number of initial clusters, this limitation was reached during array creation and before any initial cluster assignments could be made. This limitation is implementation dependant. The algorithm would not have this issue in a different language, or if steps were taken to allow writing of memory to disk during array creation.

Results from testing the ASTER and the THEMIS data are presented in Table 3. In both data sets, a metastable state exists around 10 clusters, as seen in the results for five initial end-members regardless of the number of iterations. Due to the small size, relative to the initial starting value of the other tests, this metastable state was not encountered at other initial values. Within 50 initial end-members and varying iterations, the default case of 100 iterations shows the greatest number of clusters in the ASTER data, and is bordered by fewer means on either side in the THEMIS data. In most cases, the initial number of means and the number of iterations are both positively correlated with the final number of clusters.

4.4 *Assignment of super-resolved values*

In Tonooka (2005), a circle with a radius of 10 low-resolution pixels is checked first for the best-fit homogeneous pixel spectrum before comparing with the cluster tree. In the data used by Tonooka (2005), this resulted in a mean MD of 1.14 and 71.8% of the super-resolved pixels sourced from the map (Figure 3). With increasing radii, a decrease is seen in the mean MD and the percentage sourced from the map appears to asymptotically

Table 3. The initial starting number of VNIR clusters and the maximum number of ISODATA iterations are compared to the final number of VNIR clusters.

THEMIS					
		Iterations			
		10	100	1000	10000
Initial clusters	5	10	10	10	11
	50	85	118	92	142
	500	534	1192	1190	
	1000	673	2117	2000	
ASTER					
		Iterations			
		10	100	1000	10000
Initial clusters	5	10	10	10	10
	50	64	100	95	98
	500	249	472	528	566
	1000	284	470	540	575

Notes: These two variables are the dominant factors in determining the final number of clusters. Both show positive correlation with the final number. Missing values in the THEMIS table are due to memory management issues.

approach ~80%. Decreasing the radius results in a significant decrease in both fit, as measured by average MD and percentage sourced from the map, with a radius of 0 showing 99.7% from the tree and an average MD of 2.64. The 0.3% not sourced from the tree were pixels whose MD exceeded the default value used to indicate no fit within the map; the same issue is present with the artificially-generated terrain but not with the THEMIS data. The overall appearance of the image, in terms of visual appeal, improves with an increasing radius. The artificial terrain and the THEMIS data show the same trend in mean MD, although the artificial terrain sources entirely from the image by a radius of 5, reflecting the relatively few true end-members present within this dataset.

4.5 *Radiometric correction*

Table 4 shows the distance, initial value and correction factors of each of the three different radiometric correction methods for the ASTER data. The original low-resolution pixel associated with these data had a value of 1205 DN, whereas the uncorrected DN values range from 1129 to 1234. After correction, the equal allocation method has the value closest to the original data, with a convolved value of 1204.82. As the DN values are integers, this is rounded to 1205 DN. The weighted method using the convolved distance of 0.5187 as the divisor to calculate the correction factor produces the next best results, with a convolved value of 1204.71; also rounded to 1205 DN. The weighted method using the summed distance of 18.2893 produces non-radiometrically accurate results. The convolved and rounded value from this method produces 1190 DN.

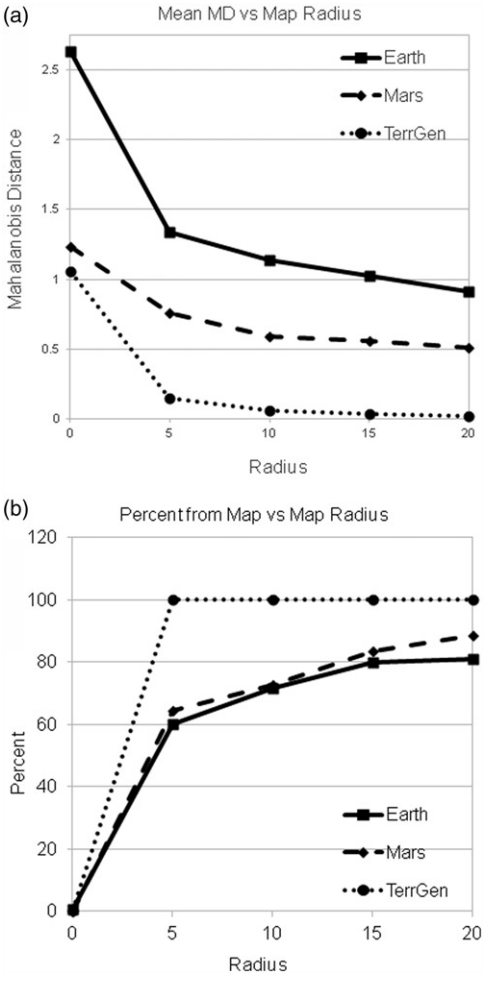

Figure 3. The effect of altering the radius searched within the image for an adjacent homogeneous pixel. (A) In all three data-sets, there is a decrease of mean MD with increase in radius, asymptotically approaching a lowest value for each data set. For the ASTER data, this appears to be 0.9, for the THEMIS data it appears to be 0.5, and for the artificial terrain it appears to be 0.02. Better data fits have lower MD values. (B) The percentage of super-resolved pixels sourced from the image asymptotically approaches a final value. In the natural data (Earth and Mars), this value is less than 100%. This shows that the cluster-tree is a necessary component of the super-resolution process, and improves the final product. However, as can be seen in a radius of 0, it is not sufficient on its own.

The results of the same tests applied to the THEMIS band 4 are shown in Table 5. The original low-resolution pixel associated with these data had an atmospherically corrected emissivity of 0.9766, whereas the non-corrected pixels had values ranging between 0.9689 and 0.9770. The same pattern of correction results was seen in the ASTER data. The equally allocated correction produced a convolved value of 0.9763, the weighted by convolved distance correction produced a convolved value of 0.9761 and the weighted by summed distance method produced a convolved value of 0.9737.

Table 4. Super-resolved pixel (240, 240) of ASTER band 10 is traced from the initially assigned super-resolved value through each of the three different radiometric correction methods.

ASTER distance map (M.D.)					
1.2657	0.4841	0.3093	0.3374	1.4847	0.6375
0.1921	0.3466	0.3236	0.6581	0.1890	0.6772
0.2679	0.8097	0.3748	0.5881	0.4190	0.2892
0.0314	0.4584	0.5856	0.2464	0.6438	0.7273
0.4084	0.3574	0.7225	1.2173	0.6196	0.5358
0.2605	0.5037	0.1964	0.1469	0.3966	0.5775
ASTER initial values (DN)					
1170	1167	1192	1195	1196	1184
1203	1192	1214	1195	1188	1173
1202	1196	1195	1213	1195	1194
1161	1196	1195	1209	1234	1192
1169	1183	1129	1205	1195	1195
1169	1170	1194	1140	1195	1129
ASTER convolution correction (DN)					
44.5812	17.0522	10.893	11.8831	52.2968	22.4544
6.76696	12.2077	11.3971	23.1811	6.65795	23.8537
9.43707	28.5205	13.2033	20.7135	14.7598	10.187
1.10428	16.1467	20.6272	8.67833	22.6755	25.6179
14.3847	12.5898	25.449	42.8773	21.8264	18.8732
9.17682	17.7412	6.91888	5.17361	13.9714	20.3404
ASTER sum correction (DN)					
1.2643	0.4836	0.3089	0.3370	1.4831	0.6368
0.1919	0.3462	0.3232	0.6574	0.1888	0.6765
0.2676	0.8088	0.3744	0.5874	0.4186	0.2889
0.0313	0.4579	0.5850	0.2461	0.6431	0.7265
0.4079	0.3570	0.7217	1.2159	0.6190	0.5352
0.2602	0.5031	0.1962	0.1467	0.3962	0.5768
ASTER equal correction (DN)					
18.2693	18.2693	18.2693	18.2693	18.2693	18.2693
18.2693	18.2693	18.2693	18.2693	18.2693	18.2693
18.2693	18.2693	18.2693	18.2693	18.2693	18.2693
18.2693	18.2693	18.2693	18.2693	18.2693	18.2693
18.2693	18.2693	18.2693	18.2693	18.2693	18.2693
18.2693	18.2693	18.2693	18.2693	18.2693	18.2693
ASTER convolution corrected (DN)					
1215	1184	1203	1207	1248	1206
1210	1204	1225	1218	1195	1197
1211	1225	1208	1234	1210	1204
1162	1212	1216	1218	1257	1218
1183	1196	1154	1248	1217	1214
1178	1188	1201	1145	1209	1149
ASTER sum corrected (DN)					
1171	1167	1192	1195	1197	1185
1203	1192	1214	1196	1188	1174
1202	1197	1195	1214	1195	1194
1161	1196	1196	1209	1235	1193
1169	1183	1130	1206	1196	1196
1169	1171	1194	1140	1195	1130

(Continued)

Table 4. Continued.

ASTER equal corrected (DN)					
1188	1185	1210	1213	1214	1202
1221	1210	1232	1213	1206	1191
1220	1214	1213	1231	1213	1212
1179	1214	1213	1227	1252	1210
1187	1201	1147	1223	1213	1213
1187	1188	1212	1158	1213	1147

Notes: The associated original resolution pixel has a value of 1205 DN.

Table 5. Super-resolved pixel (240, 240) of THEMIS band 4 is traced from the initially assigned super-resolved value through each of the three different radiometric correction methods.

THEMIS distance map (M.D)		
0.0623	0.0835	0.0890
0.1742	0.2030	0.1805
0.0096	0.5068	0.0704
THEMIS initial values (emissivity)		
0.9696	0.9769	0.9697
0.9735	0.9770	0.9704
0.9689	0.9708	0.9731
THEMIS convolution correction (emissivity)		
0.0012	0.0016	0.0017
0.0034	0.0039	0.0035
0.0002	0.0098	0.0014
THEMIS sum correction (emissivity)		
0.0002	0.0002	0.0003
0.0005	0.0006	0.0005
0.0000	0.0015	0.0002
THEMIS equal correction (emissivity)		
0.0041	0.0041	0.0041
0.0041	0.0041	0.0041
0.0041	0.0041	0.0041
THEMIS convolution corrected emissivity		
0.9708	0.9785	0.9715
0.9769	0.9809	0.9739
0.9691	0.9807	0.9744
THEMIS sum corrected emissivity		
0.9697	0.9771	0.9700
0.9740	0.9776	0.9709
0.9690	0.9723	0.9733
THEMIS equal corrected emissivity		
0.9737	0.9810	0.9739
0.9776	0.9811	0.9745
0.9731	0.9750	0.9772

Notes: The original resolution pixel associated with these data has an atmospherically corrected emissivity value of 0.9766.

5. Discussion

5.1 *Convolution with the PSF*

The choice of an incorrect PSF can lead to the creation of radiometrically inaccurate products during the super-resolution process, as evident in the ASTER data. The PSF is frequently measured for instruments before launch, either directly or indirectly through measurement of the Modulation Transfer Function (MTF). For Earth-orbiting satellites, the PSF can also be calculated by examining images that contain long linear features, such as walls, parallel to both image dimensions. As the calculated contribution moves away from the correct value, an increasing DN value is seen. At larger alpha values, this was expected; the larger alpha value indicates a decrease in the contribution from the pixel, with a greater signal coming from the surrounding pixels. However, the radiometric correction is applied only and entirely to the pixel, leading to an overcorrection. At smaller than actual alpha values, the opposite process leads to the same result. During the calculation of the amount of correction necessary to make the resultant super-resolved product radiometrically accurate, the image is convolved with the PSF. If too little contribution is calculated as coming from the surrounding region, the correction factor that is calculated is too large. However, this correction factor is still split only within the super-resolved sub-pixels. The correct alpha value is therefore critical, and occupies an optimum for the radiometric correction process.

With the two end-member artificial data and the THEMIS data, the impact of an incorrect PSF is negligible. In these datasets, there is a trend towards better results with an increasingly centred PSF, or a higher percentage originating from within that pixel. The THEMIS data has a limited dynamic range and, arguably, represents very few end-members. This commonality between the two datasets may be a result of few end-members. The choice of PSF may become important only with greater data diversity, as more can be contributed from the surrounding pixels. Additional bands, or intermediate spatial / spectral bands as in ASTER SWIR, can be used to increase data diversity only if the surface exhibits greater spectral diversity within these added bands than in the original data.

5.2 *Identifying homogeneous pixels*

The ASTER data show an extremely high percentage (93%) of homogeneous pixels at the default value. This may be accurate, given that the scene has a significant urban component; these pixels, although not homogeneous in the sense of being a single pure end-member, may contain the same end-members in roughly the same percentages, resulting in spectrally similar pixels. However, it may lead to too many pixels, relative to the total number, being used to create the cluster-tree. The THEMIS data presents a perhaps more accurate case, in selecting significantly fewer pixels (11%). Both scenes select roughly the same absolute number of pixels as homogeneous; this value falls within the range that creates good cluster-trees.

Where the ASTER and THEMIS data are compared as percentages of homogeneity, non-intuitive values are observed. Mars is generally spectrally bland, with a surface heavily dominated by basalt and dust (Bandfield *et al*. 2000, Ruff and Christensen 2002). The Earth, in contrast, has numerous spectral end-members to cause heterogeneity. The difference is the comparison to the average standard deviation. ASTER data, due to the numerous end-members, has a significantly higher standard deviation relative to the

THEMIS data. This difference is responsible for the lower percentage of homogeneous pixels on Mars. The spectral blandness present in most of the THEMIS data sets a higher threshold for homogeneity. A threshold choice other than standard deviation may result in better selection in such cases.

The artificially generated terrain with two end-members is an extreme example, created by selecting the two most dissimilar end-members from 1000 randomly created ones. As a result, it has a single threshold value for both the image and a spatial subset of over one-third the data range, and even the smallest threshold on a per-band basis is nearly one quarter of the data range. This led to the incorrect selection of 22,648 non-homogeneous pixels in the default case.

The impact of non-homogeneous pixels, selected by a threshold that is too low, reduces the utility of the cluster-tree created in the next step. Inclusion of data from an entire scene, rather than the subset being super-resolved, tends to produce too high a threshold value, permitting too many pixels to be marked as homogeneous. A balance is necessary between having too low a threshold and including non-homogeneous pixels within the cluster-tree creation, and having too few pixels for good clustering. The best results appear to arise by using either a per-band threshold, defined only from the area being super-resolved, or a single threshold defined by the minimum per-band standard deviation threshold. Another technique worth investigation in the future would be to cluster all data through some minimal number of loops, and then using the resulting cluster sizes to help specify homogeneous versus non-homogenous pixels. In this method, homogeneous will be near the centre of clusters, while non-homogenous pixels would be more distant. These clusters could then be constrained based on the cluster-size, and the results used to initialise the next step.

5.3 *The cluster tree*

In Tonooka (2005), K-means clustering was used to build the cluster-tree, with 50 VNIR clusters, 10 SWIR clusters per VNIR cluster (or 500 clusters total) and five TIR clusters per VNIR/SWIR cluster (or 2500 clusters total). As these clusters were derived within K-means, these numbers were invariant, and required significant knowledge of the scene prior to clustering. ISODATA clustering used in this algorithm had the same values as the initial number of clusters, but allows the number of clusters to vary, alternatively splitting if they grow too large or joining of their centres become too close. As seen in Table 3, the final number of clusters shows a strong positive correlation to both the initial number of clusters and to the number of iterations used for clustering.

In several cases, a relatively low number of final clusters are found, with values near or less than the initial number of clusters. This is driven by the initial end-member selection method. Initial means are selected by picking a random starting homogeneous pixel as the first means and then selecting those homogeneous pixels that are furthest in total MD from the previously selected means. Because the resulting initial mean pixels will generally be chosen from the periphery of the data space, there is a tendency towards consolidation of clusters within the first several iterations. It is not uncommon for the number of clusters to drop from 33% to 50% of the initial starting value. It is only after peripheral clusters have migrated inward that splitting of clusters becomes dominant, and the number of cluster starts to grow. Finally, where metastable states exist, such as the case for 10 end-members in both datasets, clustering will halt or return repeatedly to this state.

This returning trend can be seen in the case of the THEMIS data with five initial end-members and 10,000 iterations, in which the 11th cluster was a small offshoot of one of the other clusters. Given several more iterations, it is likely that this cluster would be eliminated as its members migrated back to the originating cluster. The formation of the offshoot cluster was driven by a few extreme pixels, relative to the originating cluster. There were an insufficient number of these pixels, however, to meet the minimum cluster size, and so the cluster was not stable across multiple iterations.

This step of the algorithm is the one of the most CPU-time intensive, and can cause significant issues for memory management. Future work will focus upon improvements to the code in this step. These can be incorporated by making use of implementations available in the libraries of other languages, such as C or python.

5.4 *Assignment of super-resolved values*

The ASTER data show a definite limit to the percentage of data sourced from the map versus the cluster-tree. A radius of 10 produces acceptable results; an increasing radius results in a lower mean MD, indicative of a lower fit of the model data to the original. A radius of 10 should take less time to compute than the larger radii, and profiling of the algorithm shows this to be the single most time-intensive step within the program, including the data clustering. However, the amount of time added by increasing the radius is not linear. The time to complete processing with double the radius is less than double the original time required to complete this step. The majority of the time cost for this step is up-front, with some additional time needed for increasing the radius.

Super-resolution of some environments may also benefit from a larger radius. The areas in which there are dispersed homogeneous deposits, such as some putative chloride units on Mars, a radius of 10 may not be sufficient to find those units within the map. These units will be clustered together within the tree, and so an initial spectrum can be selected from the tree. However, the strength and weakness of the cluster-tree is that it averages together a large number of pixels; smaller scale spectral variability is lost. This is helpful where trying to determine what an image end-member might be, but is less useful where trying to provide an accurate spectrum of a specific pixel. A larger correction factor in step 5 would be necessary as a result of the loss of small-scale spectral variability while sourcing from the tree.

The artificial TerrGen data were created with only two end-members. Due to the relatively large threshold value, over 33% the data range, a significant percentage of the image was marked homogeneous and clustering resulted in 25–40 low-resolution clusters per run. As a result, the use of the tree falls off immediately. Original super-resolved values derive entirely from the map by a search radius of five low-resolution pixels. Where the large numbers of 'homogeneous' pixels were averaged together into the clusters, they produced centres close to the five core pixel values. However, as there were so many pixels with different values, a closer fit to the associated high-resolution pixel was always found within the immediate area. As the search area increases, the fit of this match improves (i.e. mean MD decreases). The same rate of decreasing mean MD was observed in the ASTER data. In particular, the maximum MD for a given radius had a trend of being roughly $2 \times$ the average MD of the 5 pixel smaller radius in both sets of data. This test was meant to examine the effect of altering the radius with which to search for adjacent homogeneous

pixels in step 4, but served as a good illustration of the problems associated with choosing too large a threshold value in step 2.

The suggested radial value for super-resolution has been increased from 10 (proposed in Tonooka (2005)) to 20 (based on these results). This value balances the cost of computational time to the improvements in spectral fit, particularly for natural data. Testing shows that the calls to the tree are necessary, as the mean MD for a purely image-based sourcing is more than twice that of the radius-10 case within the ASTER data. The artificial terrain data shows a similar decrease in the MD with an increase in the radius.

5.5 *Radiometric correction*

Although the equally weighted correction produced the most radiometrically accurate results where the data were convolved back to the original resolution, this method is not recommended. Ignoring the fit of the data, as measured by the MD of their original assignment, and treating them as all equally good or bad is too simplistic an approach. Pixels which were exact matches to their source (MD of 0) would receive just as much of a correction as pixels that were far from their source. As a result, the previously correct pixels would be over-corrected, and the poorer-fitting super-resolved pixels would continue to be overcorrected. Allocating the correction factor strictly according to the summed weight fails due to under-correction. Larger spatial scale differences between data-sets will lead to larger under-corrections, with the only redeeming quality of this method being that it will not over-correct. The method first used in Tonooka (2005), based on equations 6, 7, 13, and 14 of that work, provides the best results by allocating the correction according to the MD of the super-resolved pixels.

The correction allocations of all the three methods examined ignore the impact of convolution with the PSF. In most cases, this leads to pixels that are radiometrically accurate where not convolved with their surroundings, but become slightly inaccurate where convolved. The magnitude of the inaccuracy is inversely proportional to the alpha value used for the PSF generation. The calculation of the radiometric correction for any given low-resolution pixel includes the contribution of uncorrected values from neighbouring pixels to the PSF, but this value is then allocated solely among the sub-pixels of the pixel of interest. Methods of radiometrically correcting data that take this into account might improve the ultimate super-resolution results. Similarly, a method such as the summed distance method, that consistently under-corrects, may be applied multiple times. This may lead to more accurate results without the loss of the improved spatial clarity or the risk of over-correction in other pixels at the cost of increased computation time.

6. Conclusion

A super-resolution technique was developed for the Earth-orbiting ASTER instrument, and presented in Tonooka (2005). Although this technique was originally developed for just the ASTER instrument, it was modified to work with data from the Mars-orbiting THEMIS instrument (Hughes and Ramsey 2010). In the process of implementing and adapting the algorithm, significant changes were necessary to process the THEMIS data. These changes have also been applied to new ASTER scenes, which no longer have SWIR data, and found to produce equal or better results.

The super-resolution method presented here is a five-step process, and the impact of making alterations to each step was examined. During the first step, the impact of choosing an incorrect PSF was shown. More accurate estimates of the PSF produced more radiometrically accurate final products, though results from an inaccurate PSF still have qualitative value. In step 2, the choice of threshold is seen to influence clustering results. A threshold defined on a per-band basis or as a band-average based only on the area of interest produces better results than the original threshold definition. The third step clusters the data and a new approach to clustering has been implemented. ISODATA clustering was found to be adaptable enough to recover from many poor user choices of initial starting conditions. The most critical choices were the starting number of clusters and the number of iterations before forcing a halt. With too few initial clusters, ISODATA may prematurely halt or return repeatedly to a metastable state. Due to the method of end-member selection, the ISODATA algorithm will undergo an initial reduction in the number of clusters; a low number of permitted iterations will cause the clustering to end during this stage. Performance issues associated with the algorithm were noted in this step, particularly during the ISODATA assignment of homogeneous pixels to new clusters. The fourth step is also CPU-intensive. Increasing the search radius for good matches improves the final results at a small cost in time. The majority of the time-cost in this step is up-front instead of scaling linearly or exponentially with the radius. Different methods of allocating the radiometric correction in the final step were also examined; the method proposed in Tonooka (2005) provides the best results and accounts for radiometric differences at the super-resolved scale.

During the super-resolution processing, some performance trade-offs are necessary. The algorithm can run more rapidly at the cost of greater memory usage, but will encounter limits imposed both by the operating system and the programming environment. The current implementation is also limited to five or fewer bands for high-resolution data, and 10 or fewer bands for low-resolution data due to these same memory restrictions. Workarounds for such arbitrary limits should be examined in future work. Similarly, the problem of determining how the radiometric correction should account for the PSF needs to be addressed. The results produced by the current algorithm are good and interpretable, but improvement is possible.

The modified algorithm is an improvement of the original implementation. The algorithm has been extended to run on multiple data types, including both the THEMIS data and atmospherically-corrected ASTER data. As part of extending the algorithm to other data, a user-defined PSF has also been implemented, using an alpha notation. The requirement for intermediate spatial/ spectral resolution data, provided by the ASTER SWIR instrument in the original implementation (Tonooka 2005), has been removed after it was determined unnecessary for radiometrically correct results. Multiple data have been used to examine the suggested values for user input, and in some cases better default values have been found. These new default values were found to perform better with both the ASTER and THEMIS data. ISODATA has been used in the place of K-means to build the cluster tree in the modified algorithm examined in this work. This replacement requires less *a priori* knowledge by the user, and provides greater flexibility to the knowledgeable user. The modified algorithm is more transparent to the user, and provides data from intermediate steps as well as runtime output. This permits users to determine how their final product was created and aids in tracing the propagation of incorrect input through the algorithm. In all steps, the CPU cost of the functions has been optimised for performance.

The algorithm relies on a number of choices by the end user, which determine the quality of the super-resolved data. Testing of the modifications has led to a new set of default recommendations for these values. This default configuration produces radiometrically-accurate, 15 m/pixel and 36 m/pixel TIR data for ASTER and THEMIS, respectively, and does so at a reasonable time-cost. Users with no insight into the algorithm can super-resolve ASTER and THEMIS data with good results by making use of the default values. The modified algorithm will accept other data sources for super-resolution as well, permitting this process to be extended to other Earth and Mars data, as well as data from other planetary surfaces. These other data sets need not be limited to the spectral range examined within this work. Users should ensure they understand the statistical distribution of their data prior to application of the algorithm, although the new default recommendations will provide a good starting point for most data.

Acknowledgements

This work was supported by the M.S. Ramsey's funding from the NASA THEMIS participating scientist program (NMO7106s0) and the NASA ASTER Science Project (NNX08AJ91G). This research was completed as part of the C.G. Hughes' dissertation at the University of Pittsburgh, and was not done in his capacity as an employee of the Jet Propulsion Laboratory, California Institute of Technology. This work was improved through pertinent comments made by the dissertation committee and the external reviewers.

References

Aiazzi, B., *et al*., 2002. Context-driven fusion of high spatial and spectral resolution images based on oversampled multiresolution analysis. *IEEE Transactions on Geoscience and Remote Sensing*, 40 (10), 2300–2312.

Ball, G.H. and Hall, D.J., 1967. A clustering technique for summarizing multivariate data. *Behavioral Science*, 12 (2), 153–155.

Bandfield, J.L., *et al*., 2004. Atmospheric correction and surface spectral unit mapping using Thermal Emission Imaging System data. *Journal of Geophysical Research-Planets*, 109, E10008, doi:10.1029/2004JE002289.

Bandfield, J.L., Hamilton, V.E., and Christensen, P.R., 2000. A global view of Martian surface compositions from MGS-TES. *Science*, 287 (5458), 1626–1630.

Christensen, P.R., *et al*., 2004. The Thermal Emission Imaging System (THEMIS) for the Mars 2001 Odyssey Mission. *Space Science Reviews*, 110 (1–2), 85–130.

Garzelli, A., *et al*., 2004. Pan-sharpening of multispectral images: a critical review and comparison. *In: Geoscience and remote sensing symposium, 2004. IGARSS '04 proceedings. 2004 IEEE international*, Anchorage, AK, USA, 81–84.

Hughes, C.G. and Ramsey, M.S., 2010. Super-resolution of THEMIS thermal infrared data: compositional relationships of surface units below the 100 meter scale on Mars. *Icarus*, 208 (2), 704–720.

Iwasaki, A. and Tonooka, H., 2005. Validation of a crosstalk correction algorithm for ASTER/SWIR. *IEEE Transactions on Geoscience and Remote Sensing*, 43 (12), 2747–2751.

Mimmack, G.M., Mason, S.J., and Galpin, J.S., 2001. Choice of distance matrices in cluster analysis: defining regions. *Journal of Climate*, 14 (12), 2790–2797.

Pohl, C., 1999. Tools and methods for fusion of images of different spatial resolution. *International Archives of Photogrammetry and Remote Sensing*, 32, 76–82.

Ruff, S.W. and Christensen, P.R., 2002. Bright and dark regions on Mars: particle size and mineralogical characteristics based on Thermal Emission Spectrometer data. *Journal of Geophysics Research*, 107 (E12), 5127.

Tonooka, H., 2005. Resolution enhancement of ASTER shortwave and thermal infrared bands based on spectral similarity. *Proceedings of SPIE*, 5657, 9–19.

Townshend, J.R.G., *et al*., 2000. Beware of per-pixel characterization of land cover. *International Journal of Remote Sensing*, 21 (4), 839–843.

Wang, Z.J., *et al*., 2005. A comparative analysis of image fusion methods. *IEEE Transactions on Geoscience and Remote Sensing*, 43 (6), 1391–1402.

Yamaguchi, Y., *et al*., 1998. Overview of Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER). *IEEE Transactions on Geoscience and Remote Sensing*, 36 (4), 1062–1071.

Zhang, Y., 2004. Understanding image fusion. *Photogrammetric Engineering and Remote Sensing*, 70 (6), 657–661.

Zhukov, B., *et al*., 1999. Unmixing-based multisensor multiresolution image fusion. *IEEE Transactions on Geoscience and Remote Sensing*, 37 (3), 1212–1226.