JPL XXXXXXXXXXX
EOS AIRS DRL XXX
ATBD-AIRS-01

Earth Observing System (EOS)

Atmospheric Infrared Sounder (AIRS)

AIRS Level 1C Algorithm Theoretical Basis

Evan Manning, Hartmut H. Aumann, Yibo Jiang, and Denis A. Elliott

JPL

Margaret H. Weiler

???

Larrabee Strow and Scott Hannon

UMBC

Version 3.0
August ??, 2014

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California 91109-8099

Release Record

Version	date released	comments
1.0	March 1, 2011	Initial Release
2.0	March 10, 2012	Final Release
3.0	Aug ??, 2014	Release with first software release. Clean and Fill only.

TABLE OF CONTENTS

1. Introduction and Review

1.1 AIRS Instrumentation

The Atmospheric Infrared Sounder (AIRS), launched aboard NASA's EOS Aqua spacecraft on May 4, 2002, is a grating array spectrometer having 2378 channels sensitive in the range 3.7 to 15.4 microns with some gaps. The spectral resolution ($\lambda/\Delta\lambda$) is about 1200. A combination of a design philosophy having radiometric accuracy as a foremost goal, cooled and temperature-controlled spectrometer hardware (including most of the optics), and thorough pre-flight calibration have made the AIRS a superb instrument that produces high quality radiance data [Strow *et al*., 2003]. AIRS has completed twelve years in routine operations at the end of August 2014. The instrument remains healthy and it is hoped that it will continue to operate and produce high quality data for several more years. This long data record makes AIRS an excellent candidate for producing useful records for climate trend analyses. To maximize its utility as a source of climate data, any instrument must be capable of being accurately cross-calibrated with other instruments so that lengthy multi-instrument records may be generated.

Because AIRS is a grating array spectrometer, not a Fourier transform spectrometer, each of the 2378 channels is independent of the others. Each channel is separately calibrated, and has the potential for noise behavior that is different from the other channels—even neighboring ones. In fact, the noise behavior of a channel can change (or a channel could even stop functioning) suddenly, perhaps due to a radiation hit, without any other channel being affected. Channel variability in noise characteristics, if unaccounted for, complicates error estimates, noise estimates, and cross-instrument calibration.

In this document, we will describe the methodology and procedure that will produce cleaned-up AIRS spectra that can be conveniently used for instrument cross-calibration and studies of climate trends. The final output will be a prescription for producing Level 1C radiances that will become publicly available in a future version of the AIRS science product generation software (currently in Version 6.0) operating at the Goddard Earth Sciences Data and Information Services Center (GES DISC).

Level 1A products contain raw detector counts. Level 1B consists of radiometrically calibrated radiances [Gaiser *et al*., 2003]. Both Level 1A and Level 1B have been produced since the start of routine instrument operations on September 1, 2002. Level 1C consists of cleaned up spectra, in which dead or unusually noisy channels have been marked and their radiances corrected or replaced using well-behaved correlated AIRS channels. Furthermore, the small gaps in the spectrum are filled in with reasonable values calculated from the same spectrum, and the discrepancies of overlapped channels between adjacent modules are corrected and the overlapped channels are removed.

1.2 Motivation for the production of Level 1C

An ideal spectrum should have the following properties:

1) All frequencies are included (at the design resolution) with no gaps and no overlapped spectral regions.
2) All channels have only Gaussian noise, so that far outliers do not interfere with integration over the pass band or interpolation during resampling.
3) Adjacent channels are truly independent measurements, with adjacent-channel correlations minimized.
4) The channel frequencies are fixed in time.

The AIRS L1B spectrum deviates by small amounts from the ideal spectrum described above. Although the instrument's in-flight performance has far exceeded its specifications, which were developed with temperature and humidity profile measurement in mind, the originally envisioned use of AIRS was for non-climate-related studies of atmospheric phenomena and for improved weather prediction. Because of its exceptional performance the AIRS data is now being used for climate studies. AIRS climate data records can be improved if the deviations from ideal performance are accounted for.

AIRS, a grating array instrument, has seventeen detector modules spread across the focal plane. The hardware design was simplified by permitting small gaps in the frequency coverage between some modules, and small overlaps in coverage between others. These gaps and overlaps complicate the task of integrating the spectrum over the pass band to enable cross-instrument comparisons and calibrations.

Because of the large frequency range of the entire instrument, each detector module is made from different material having different properties, and unique read-out integrated circuits (ROIC). The shortwave modules include circuits to remove spikes resulting from radiation hits, while these spikes are obvious in the longwave modules. Some channels exhibit non-Gaussian noise including "pops" (temporary changes in output level) and cold scene noise (scene-dependent noise larger at lower signal levels) [Weiler *et al*., 2005]. The AIRS detectors and their ROIC's have different susceptibilities to radiation hits and to the slow build up of total radiation dosage throughout the mission. Thus, a channel that has exhibited very low noise for years can suddenly undergo a noise increase. It is also possible for a channel that has been affected by a radiation hit to recover after a period of hours or days. So noise can vary independently over time from channel to channel.

Each channel occupies its own physical space on the focal plane. That implies that slight changes in channel frequency can occur, due primarily to changes in temperature gradients within the spectrometer optical train. The AIRS spectrometer's temperature is tightly controlled at one location by interactions between a radiator (cooling) and a heater. But small changes in internal gradients can occur and result in very small channel frequency shifts with time.

Grating array spectrometers, by their very nature, perform extremely well as far as point #3 above is concerned. However, the independence of the channels means that the other goals are not met automatically, and so must be handled explicitly. In this document, we describe each step towards a final Level 1C product that will fulfill the ideal. At first, we just handle points #1 and #2 above (spectral gaps and overlaps and variable and/or non-Gaussian noise). Point #4 will be handled last. Note that, because of the fact that the AIRS spectrometer is temperature controlled, point #4 (channel frequencies vs. time) has only a minor effect on AIRS data quality. The observed instrument frequency stability far exceeds the specification. But when the more stringent requirements for climate records are considered, it becomes clear that frequency shifts do need to be handled eventually.

Once "ideal" spectra have been produced, the use of AIRS data for many studies will be simplified and errors more easily estimated. Most importantly, it will be possible to cross-calibrate AIRS with other climate instruments more easily and reliably than that can be done today.

Figure 1 shows a typical L1B spectrum, where noisy or dead channels are indicated by the red dots at the bottom of the plot. The 17 modules and their location in the spectrum are also labeled in this plot. Of the 2378 AIRS spectral channels, about 58 have no response (dead), and an additional ~35 channels (varying with time of the measurements) have more than 2 K noise (compared to the nominal noise of 0.07-0.6 K). The presence of these ~95 channels is overemphasized in Figure 1. Table 1 lists the gaps and overlapped regions of the spectrum that will be discussed in detail in section 4.

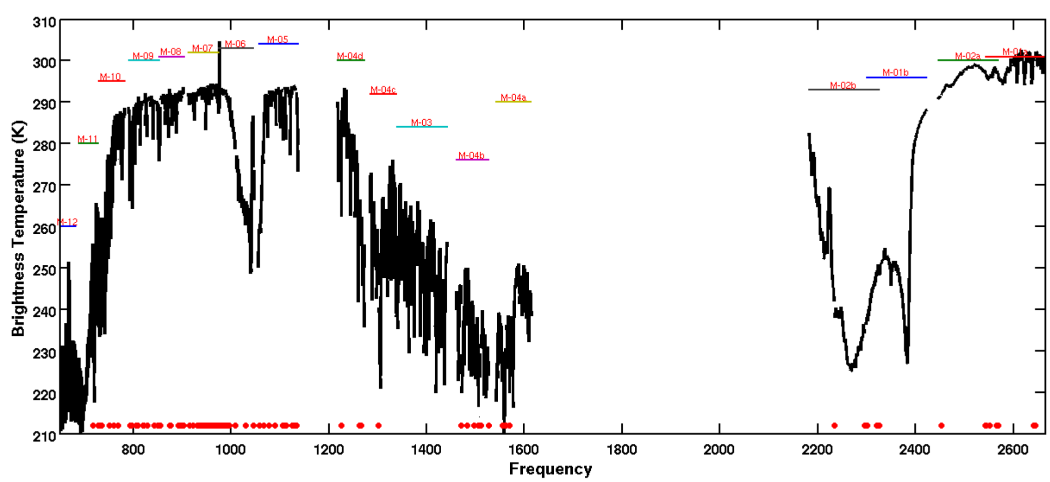

Figure 1. AIRS nighttime spectrum from January 1, 2006. The red dots at the bottom of the plot are indicators of dead or bad channels. M-* mark the spectral ranges of the detector arrays.

Table 1. AIRS Spectrum Gaps and Overlapped Regions

Type	Frequency Range (cm^{-1})	Comments
Gap	681.99-687.60	
	781.88-789.26	
	903.78-911.23	
	1046.20-1056.07	
	1136.63-1216.97	
	1272.59-1284.35	
	1443.07-1460.27	
	1527.00-1541.10	
	1613.86-2181.49	Large gap (will not be filled)
	2557.41-2558.53	
Overlap	727.83-728.36	M-11/M-12
	851.20-852.72	M-08/M-09
	973.48-974.63	M-06/M-07
	1337.64-1339.18	M-03/M-04c
	2301.72-2320.79	M-01b/M-02b
	2545.19-2565.21	M-01a/M-02a

The L1B spectra have four quality issues that can be addressed with the L1C software:

1. The radiances in L1B are correct, but the current frequencies may vary by up to 5 ppmf from the nominal frequencies due to diurnal and seasonal effects (Figure 2). The true frequencies are known to within 1 ppmf. The difference between the nominal and the actual frequencies causes errors in the interpretation of the radiances, which are very small compared to the random noise, but if ignored can produces artificial trends of as much as 10 mK/year in some channels.
2. AIRS has 2378 channels, but ~100 of the channels are "bad". We define "bad" as having NEdT > 2K at the reference temperature of 250 K. The median NEdT at 250 K is 0.2 K. ~60 of these ~100 channels were already dead (with -9999 flag value in radiance) during prelaunch testing.
3. Each channel has slightly different spatial response, so in highly inhomogeneous scenes, like the edges of cold clouds, there can be significant distortions.
4. There are small gaps between the modules. This does not include the large gap from 1614 cm^{-1} to 2181 cm^{-1} that covers the shortwave part of the water band.
5. The L1B channel sequence is not increasing monotonically in the overlapped regions.

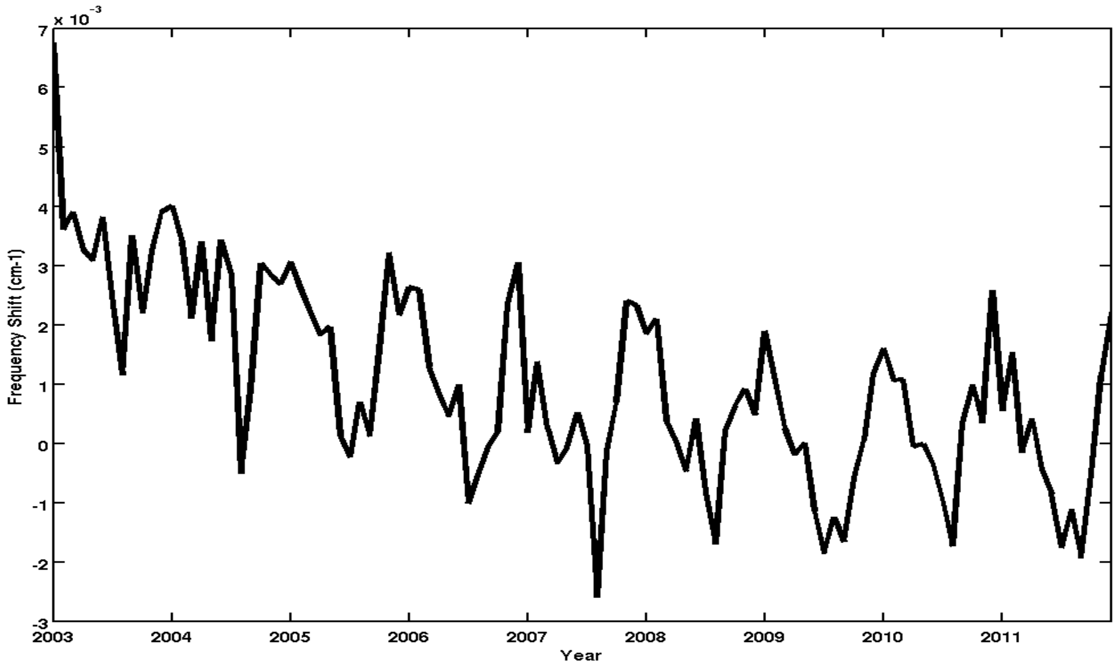

Figure 2. Monthly averaged AIRS instrument nighttime frequency shift in cm^{-1} from 2003 to 2011 in latitude range from -10 degree to 10 degree.

The issues in items 2 through 5 make it more difficult to use AIRS L1B data for band-integrated comparisons with lower resolution sounders. The L1C product will address these four items as listed below.

1. Provide "clean" substitute radiances for all of the 2378 instrument channels with significant problems. First substitute values for the worst channels are calculated using a buddy system, then a PCA reconstruction is used to create a reconstructed spectrum. The reconstructed values are used as the final L1C values for the channels that were previously filled using the buddy system and is also used to dynamically identify and replace other channels with problems.
2. Extend the clean spectrum to a new monotonic 2645-channel L1C AIRS frequency set. This involves eliminating some overlap channels and creating synthetic values for some new channels. The synthetic values are calculated using a simple buddy system.

1.3 Level-1C Approach

Level-1C is conceptually divided into 2 stages: cleaning and filling. The cleaning stage is concerned only with the 2378 channels present in the L1B product. Cleaning determines which L1B values are of lower quality and provides replacement values, producing a "clean" 2378-channel spectrum. Filling then removes overlap channels and provides values for synthetic channels in the instrument gaps, producing the final 2645-channel spectrum.

2. Spectrum Cleaning

6. PGE Execution Flow

This section describes the spectrum cleaning algorithm and its implementation. Cleaning depends on the high degree of redundancy among the channels. For any given channel there are any number of highly-correlated channels, channels which are sensitive to the temperature at roughly the same atmospheric altitudes and/or are sensitive to the same gases. The information from these correlated channels can be extracted and applied in 2 general approaches: buddy systems or non-local (Better term?) systems. In a buddy system, a small number of highly-correlated channels are selected, and the synthetic radiance for the channel to be replaced is a simple function of these. Non-local systems look at an entire spectrum or a large part of it and characterize the input spectrum using a number of parameters significantly less than the number of input channels. These parameters specify a reconstructed spectrum, the individual channels of which can be used as synthetic values for each channel.

AIRS L1C cleaning uses a combination of these approaches. First a buddy system fills in values for bad or suspect channels using a few good channels. This preliminary clean spectrum is then passed through a 100-PC process to produce a reconstructed spectrum. The reconstructed spectrum is used both to provide final replacement radiances for the known bad channels and as a reference to dynamically identify and replace radiances for other channels with transient problems.

The major steps in the spectrum cleaning process are:

- Identification of problem channels.
- First-order replacement of these bad channels using the buddy system. This step is required or else the PCA diverges.
- Measure Cij, the scene-specific impact from scene inhomogeneity.
- PCA reconstruction of the spectrum and replacement of the bad channels.
- Scan the spectrum for outliers based on PCA reconstruction and replace them with PCA values.

We can derive rough performance requirements for each level by working backwards through the steps above. In our final product we wish to avoid any artifacts or outliers bigger than ~2 K. This level varies with conditions as detailed below but we can use 2 K as the nominal target accuracy in this discussion. For the shortwave and longwave ends of the instrument the nominal noise level is near 0.6 K, so 2.0 K is a bit over 3 sigma, so only the extreme tails of the noise distribution will be distorted by this cut-off.

Since we identify outliers in the L1B spectrum as 2 K mismatches between the reconstructed and observed spectra, the accuracy of the reconstruction should be better than ~0.5 K over all channels and over all geophysical cases. This in turn requires a very good global training set, a relatively high number of PCs used in reconstruction, and a fairly good spectrum input to the PCA step.

We can assume that the PCA step spreads any distortion in its input over the output, with the most-impacted channel being the same channel with the input distortion, and a magnitude of ~20% of the input distortion. So we can tolerate input problems up to about 10 K if only one input channel in any region is imperfect but should hope for <~3 K to allow for multiple problem channels.

3 K is a relatively easy accuracy level to obtain with a buddy system, but is a relatively tight requirement on the initial step, identification of problem channels. Problem channels must be identified both so they can be replaced and so their bad values are not used as buddies in filling other bad channels. It is easy to identify most problem channels with L1B measured NEdT levels and a static list of problem channels, but extra checks are needed so almost all problems are caught. For the primary NEdT-level check, a threshold of 0.85 K means that 3 K outliers are 3.5-

sigma events. A check for measured NEdT level relative to a baseline for neighboring high-quality channels helps weed out unusual channels that are more likely to have non-Gaussian noise with more outliers and even biases.

2.1 Static Detection of Dead or Noisy Channels

Channels can be determined to need "cleaning" either statically on input or dynamically based on the mismatch between the observed spectrum and the PCA reconstructed one. This section details the initial static determination. This step is important because PCA output becomes distorted if there are too many bad observations in its input.

For each observation, typically about ~200 of 2378 channels are marked for replacement and others are excluded from use in replacing other channels. This process generally relies on noise levels, error flags, and radiances from the Level-1B product, so in principle different channels could be selected for replacement in each spectrum, but in operations we see mostly the same channels, including ~50 permanently dead or noisy from launch through the first 10 years of operations. A configuration file flags additional channels for exclusion when they would not be caught automatically. Ten of these channels are marked permanently bad because they have been unusable since the launch of the AIRS. These include 5 channels that are cross-wired, so they actually have low noise, but observe a spectral region broader than they were intended to.

Channel replacement is required whenever:

1) NEdT at 250 K scene temperature exceeds a threshold (currently 0.85 K). This is the primary quality check. Channels with more than 0.85 K of noise are likely to have >3 K outliers just from Gaussian distribution statistics, and these channels are also somewhat likelier than lower-noise channels to have non-Gaussian noise distributions, with the possibility of still higher outlier rates and even biases.
2) NEdT at 250 K exceeds a threshold (currently 3.0) times the baseline NEdT for the module and channel A/B state. Baseline NEdT for each channel assuming A+B channel state are contained in an ancillary file. The baseline is multiplied by an additional factor of sqrt(2) For channels in A-Only or B-Only modes. As with check #1, those channels with higher NEdT compared to a baseline are more likely to be non-Gaussian or biased.

Insert figure of baseline NEdT

3) NEdT is negative (indicating noise could not be characterized). Level-1B determines noise from the jitter among the downlinked observations of space and of the blackbody. Negative NEdT will be a -9999 flag value indicating that there were no valid observations for one or both of these calibration sources. Usually this means that the detector is completely unresponsive or is operating out of its sensitive range and so is saturated for at least one calibration source. These are the most impaired channels, sometimes referred to as "dead" channels.
4) L1B provides no calibrated radiance value, i.e. radiance value in L1B is the flag value of -9999.0. In addition to the same cases caught by the negative NEdT test, this test will flag cases where the individual readings have saturated, caused either by cosmic ray hits or by bright glints.
5) L1B radiance is out of a configurable range of geophysically expected values. Currently this range is [170, 420] K, with an additional margin of 5 times the channel's noise level. This test is also designed to catch extreme values from cosmic ray hits.
6) Channels are on a list of permanently or temporarily bad detectors. These include the 5 cross-wired channels and others that have been determined in testing to have undesirable characteristics but are not flagged according to the criteria above.

Channels not meeting the above conditions can instead be marked as suspect if they meet any of these conditions. Suspect channels are not replaced but are excluded from use in replacing other channels and have a lower threshold for dynamic replacement:

1) NEdT at 250 K exceeds a threshold (currently 0.70 K).
2) NEdT at 250 K exceeds a threshold (currently 1.75) times the baseline NEdT for the module and channel A/B state.
3) Observed radiance is negative. Negative radiances are expected in the shortwave region for very cold scenes, but are still suspect.
4) L1B CalFlag indicates a problem with gain or offset calculations, bad telemetry, or a "pop" event. Only pop events are common. When a pop occurs the zero level from the detector has changed between the observation of space before the current scan and the one after. There is therefore an additional uncertainty in the calibration.
5) AB_State from the current channel properties file marks the channel lower quality with state > 2. These states have been judged to be low quality by the AIRS calibration team.
6) Cij in channel properties file is less than a specified value (currently 0.92) indicating the channel is not well aligned spatially with the nominal boresight.

The median NEdT at the reference temperature of 250 K is ~0.2 K. The further detection of "bad" or "noisy" channels with the help of principal component analysis will be discussed later.

When channels are determined to need replacement, the reason is marked in the L1C product in a field named "L1cSynthReason". Channels that are only suspect are not flagged in the output.

2.2 Buddy-System Algorithm

In order to fill in the bad channels with reasonable values ("buddies") from the same spectrum, the most correlated channels (minimum standard deviation) are found based on the training set. The training set used is detailed in Appendix A. These channels will be sensitive to most nearly the same combination of geophysical parameters as the channel needing replacement.

The brightness temperature (BT) of the bad channel is replaced by a BT calculated from the BTs of most correlated channels. The correlated channel replacement list is calculated based on the following formula (1), which calculates the averaged deviation from the channel to be filled. Note that the "buddy" channels always come from the same detector module (Figure 1) as the channel that is being replaced.

$$\delta T(k,j) = \sqrt{\frac{1}{n}\sum_{i}^{n}(T_{ij} - T_{ik})^2} \qquad (1)$$

n: number of spectra in the training set
i: spectrum index, range from 1 to n;
j: channel number range from 1 to 2378;
k: channel to be filled range from 1 to 2378

T represents the brightness temperature in formula (1), *i* is the spectrum index in the training set, *j* is the replacement channel number and *k* is the channel number to be filled. $\delta T(k,j)$ represents the averaged deviation of each individual channel *j* from the filled channel *k*. For each channel *k*, $\delta T(k,j)$ values are sorted in the ascending order and the first 100 *j*s (with the least deviation from channel *k*) are selected and will be used to replace or fill the bad channels. This process is repeated for 10 15-K scene brightness temperate ranges from [220, 235) to [355 ,370) K. The resulting 100 *j*s (integer array size 2378x100) and the associated deviations (double array size 2378x100) and biases are calculated offline and are stored in an ancillary file, read at runtime.

The brightness temperature of the channel to be filled or replaced then is the average of the four best correlated of the useable channels weighted by the deviations from the filled or replaced channel. In level 1C implementation, the four most correlated channels are used in the final calculation of the brightness temperature T_k of the replaced channels (Formula 2).

$$T_k = \frac{\sum_j (T_j + fB_r(k,j)) \frac{1}{\delta T_r(k,j)}}{\sum_j \frac{1}{\delta T_r(k,j)}} \quad (2)$$

j: channel number range from 1 to 2378;
k: channel to be filled range from 1 to 2378
r: brightness temperature range from 1 to 10
B: brightness temperature bias
f: bias scale factor

The use of a bias allows us to find much better matches than otherwise, but there is a catch. The degree of bias between neighboring channels is highly scene dependent: the more clouds, the less spectral contrast. Therefore, for each channel to be replaced and in each spectrum, we first determine a bias scale factor *f*. *f* is selected from among 9 possible values {0.00, 0.25, 0.50, 0.75, 1.00, 1.25, 1.50, 1.75, 2.00} by finding the value that minimizes the penalized standard deviation of the candidate fill values $T_j + fBr(k,j)$. The penalty function requires that the standard deviation be 4x smaller for the extreme bias scale factors of 0.00 and 2.00 than for the nominal value of 1.00, which makes the process favor using the nominal value. The penalty factors for each *f* are {4.00, 3.25, 2.50, 1.75, 1.0, 1.75, 2.50, 3.25, 4.00}.

Figure 6 shows a typical AIRS spectrum after the "buddy-system" replacement of dead or noisy channels. All visible outliers are suppressed and reasonably good values are provided for missing channels. Errors of 1-5 K are still important but not obvious on this scale, so we demonstrate the further improvements made by PCA in figure 8 below and then characterize the final cleaned product in section 2.6.

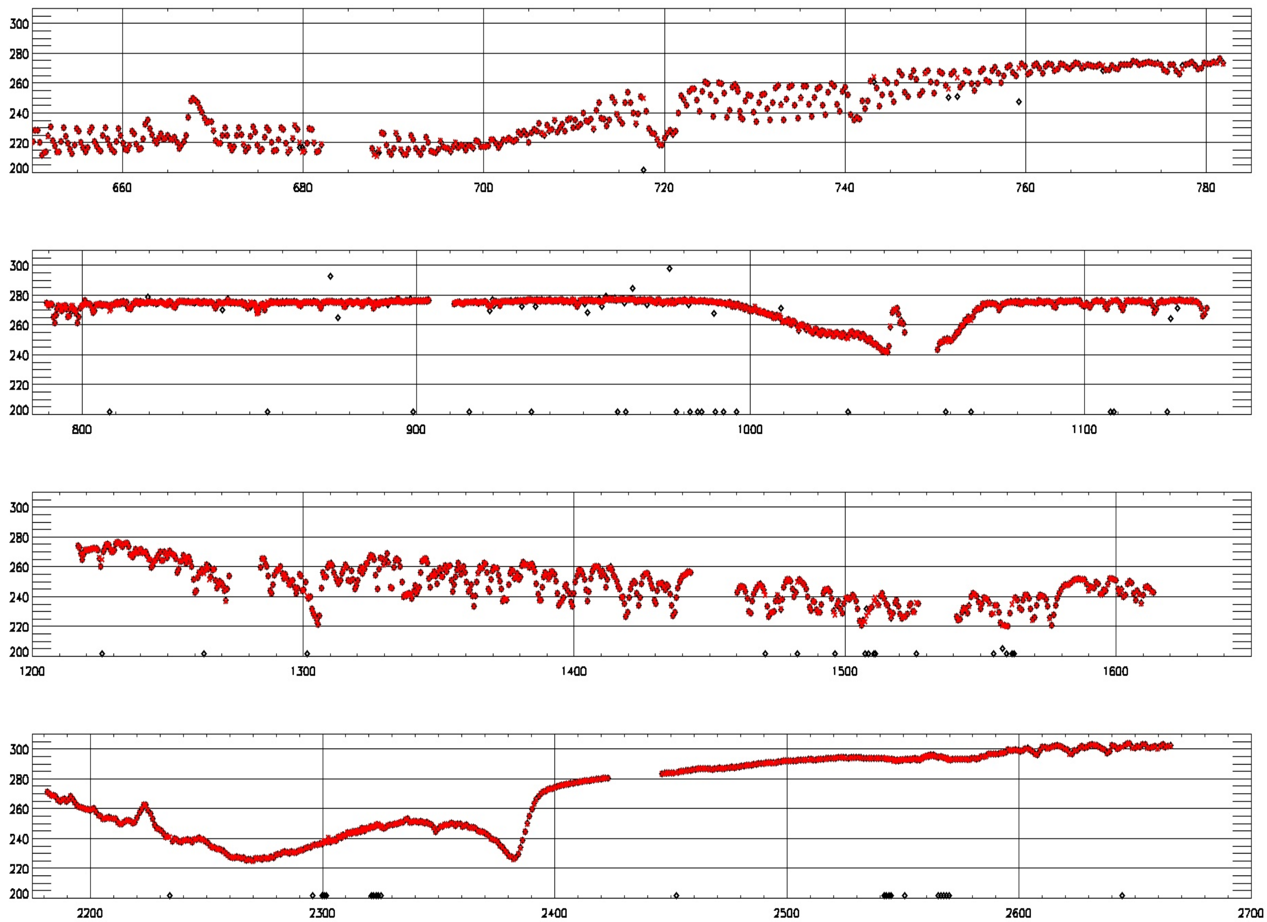

Figure 6. A typical AIRS spectrum after the first pass "buddy-system" replacement of dead or noisy channels. The black diamond symbol represents the original L1B spectrum and the red circle is the cleaned spectrum. Diamonds along the bottom of the spectrum are channels for which the L1B value was out of range or L1B did not provide a value.

2.3 Principal Component Analysis and AIRS Spectrum Reconstruction

Principal component analysis (PCA) involves a mathematical procedure that transforms a number of (possibly) correlated variables into a (smaller) number of uncorrelated variables called principal components. PCA is a simple, non-parametric method of extracting relevant information from complicated data sets. Any data can be expressed as linear combinations of eigenvectors, where the first principal component (associated with the first eigenvector) accounts for most of the variability in the data, and each succeeding component accounts for remaining variability. The real spectrum training set is used to create the fixed set of eigenvectors. As the number of components used to reconstruct the spectra increases, the reconstructed spectra will increasingly match the observed spectra, eventually including noise artifacts. We use 100 components in our PCA reconstruction.

The observed spectrum T can then be reconstructed from the PCA using formula (3) where T represents the reconstructed brightness temperature, E^T is the transposed eigenvector, and T_0 is the original spectrum. Note: all T are expressed relative to a spectral mean brightness temperature.

$$T = E^T T_0 \qquad (3)$$

Once a spectrum has been cleaned by the "buddy-system" replacement, PCA is used in order to further detect and correct the cleaned spectrum.

2.4 Cij Metric

The magnitude of Cij distortion of varies greatly from scene to scene, so it is necessary to have a metric of Cij in order to apply tight correction criteria only to Cij-impacted scenes. Cij effects are strongest for channels at the longwave end of the M-08 detector array, with frequencies near 850 cm-1, so this region is used for the metric. We define metric Inhomo850 as follows:

1. Determine the 10 "good" channels closest to the longwave end of M-08. Good channels are channels which are not bad according to the criteria in section 2.1 and also are not suspect, unless they are only suspect because of their Cij values.
2. Calculate the mean over these 10 channels of (BT1B – BTpca). Call this dBTm8.
3. Calculate the mean over these 10 channels of BTpca. Call this BTm8.
4. Determine the 10 "good" channels closest to the shortwave end of M-09.
5. Calculate the mean over these 10 channels of (BT1B – BTpca). Call this dBTm9.
6. Calculate the mean over these 10 channels of BTpca. Call this BTm9.
7. Calculate mean BT near 850: BT850 = (BTm8 + BTm9) / 2
8. RawInhomo850 is dBTm9- dBTm8.
9. In order to prevent Inhomo850 from getting too large for cold scenes, calculate:
   CijFactor = dBdT(BT850, 850) / dBdT(250, 850)
   Where dBdT is the slope of the Planck function for the given BT and frequency.
   If BT850 is > 250 K, set CijFactor to 1.0
10. Inhomo850 = RawInhomo850 * CijFactor

Inhomo850 is a mean brightness temperature difference in K, and represents the mean magnitude of the Cij effect over the most impacted channels. It is output in the L1C product files. Any scene with |Inhomo850| > ~0.84 K is considered Cij-impacted and any scene with |Inhomo850| > ~1.69 K is strongly impacted.

2.5 Dynamic Spectrum Cleaning

Using the reconstructed spectrum and the Inhomo850 Cij metric, it is possible to identify channels that were not flagged in the static checks of noise levels, etc. but need cleaning in this particular spectrum. Channels are flagged for cleaning if the difference between the input BT and the reconstructed spectrum exceeds the acceptable limit. These limits are complicated, depending on BT, Inhomo850, and whether the channel was earlier flagged as suspect.

The first dynamic check is for Cij. It is applied only to scenes where |Inhomo850| exceeds a threshold (currently 0.84 K), and only to channels known to be susceptible to Cij and that are currently marked with an "AB_State" that puts the channel in a vulnerable condition (Table 2).

Table 2 gives the Cij-sensitive frequency ranges. These are generally near the ends of detector modules where the optical path is furthest from the center. In some cases not all channels within a frequency range are vulnerable, only those with a particular set of 'AB_State's. For 745-785 cm^{-1}, there is little risk for channels in AB_State 0, indicating a good channel using both the A and B detectors, but channels that either use only one side or are of lower quality are vulnerable. For several other ranges we accept both AB_State 0 and AB_State 1, which indicates the channel is using only the A side but is otherwise in good condition. In these cases the B side is much more sensitive to scene inhomogeneity. In the remaining ranges all channels are sensitive no matter what AB_State is.

Figure 7 shows the data on which table 2 is based. One day of data was compared to PC reconstruction and the margin needed to ensure only 1 in 100,000 valid observations was flagged was calculated, separately for non-Cij scenes and for scenes with significant Cij. The difference is a measure of how important Cij is for each channel. Because channels can change between A-Only, B-Only, and A+B states as gain tables are uploaded over the mission, Cij sensitivity in each frequency range is keyed to current A/B state instead of static.

Table 2. Cij frequency ranges

Frequency range (cm^{-1})	AB_State conditions
745 - 785	AB_State not 0
800 - 845	AB_State >= 2
850 - 875	
880 - 900	AB_State not 0
910 - 925	
925 - 1000	AB_State not 0
1040 - 1150	AB_State not 0
1200 - 1220	
1340 - 1360	

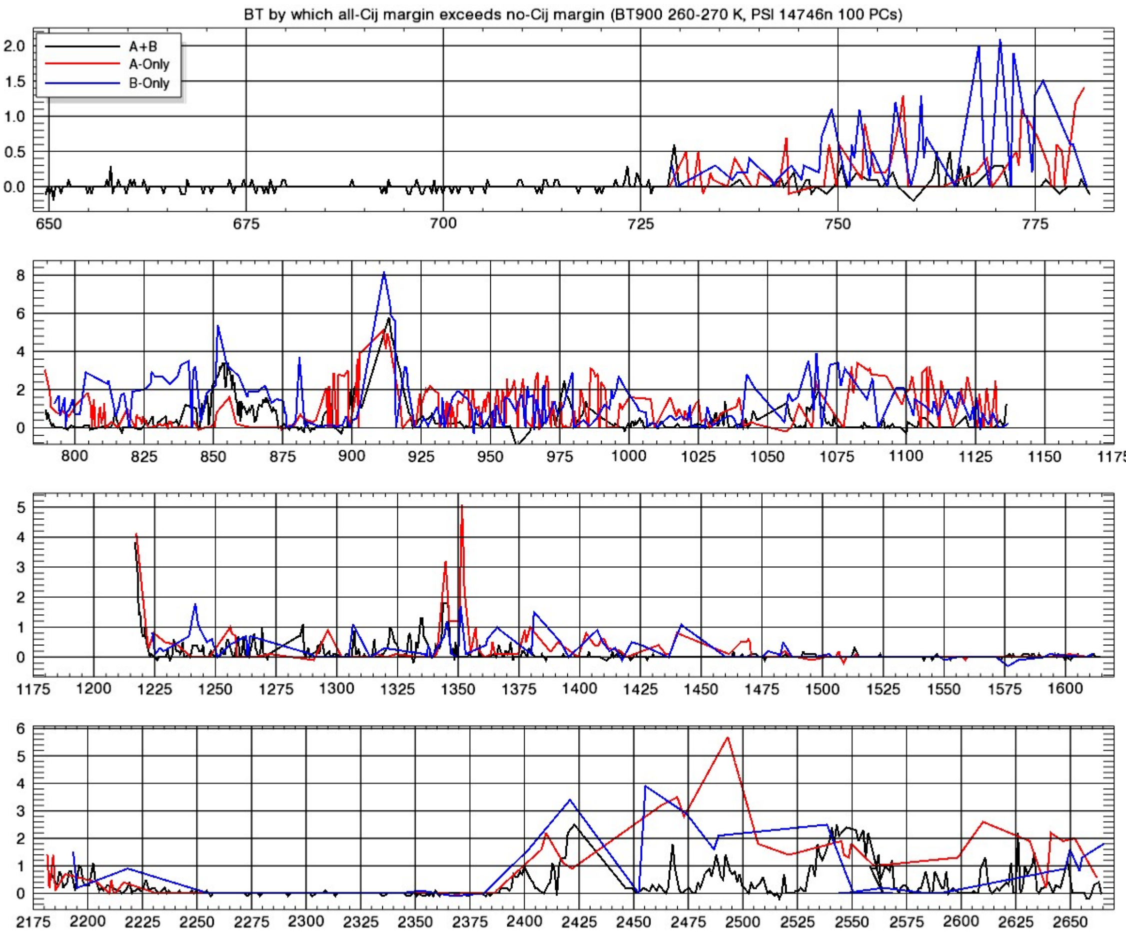

Figure 7. Estimated degradation from Cij effects depending on whether channels are A-Only, B-Only, or A+B

For scenes and channels that fit the criteria for Cij, we set the dBT threshold BT_thresh_Cij to 1.0 K if |Inhomo850| is in [0.84, 1.69) K or 0.7 K if |Inhomo850| is >1.69 K. Channels are flagged for replacement with reconstructed values if dBT > BT_thresh_Cij / CijFactor.

For all channels and scenes there is a check of dBT to eliminate outliers. An ancillary table provides thresholds for each channel and each 10K BT range, designed to be 1.25*the level at which 1 in 1,000 observations from a good channel would be flagged as an outlier just because it is in the extreme tail of the Gaussian noise distribution. The minimum dBT threshold is set to 2.0 K and several additional adjustments are made to compensate for spectral regions where PCA can perform poorly because of problems in its training set:

- Channels 2038 to 2048 (near 2324 cm-1) have the threshold raised to 2.75 K because the PCA has problems here over Siberia.
- Thresholds are raised by an additional factor of 1.5 for the photoconductive arrays M-12 and M-11 (650-728.4 cm-1) because these detectors are of high quality and training in this region is poor.
- For M-09, M-08, and M-07 (789-974 cm-1) threshold is set to 2.0 because there are many poor channels and noise dominates training errors in this region.

For channels marked suspect in section 2.1 the threshold is lowered 20%. All channels for which |dBT| exceeds the threshold, the reconstructed value is substituted and the reason for the substitution is marked in L1cSynthReason.

4. Spectrum Gap and Overlap Fill

UMBC has simulated the AIRS spectra without any gaps as shown in Figure 16 with a typical spectrum showing the filled gaps. A buddy-system algorithm simpler than the one discussed in section 2.2 is used here to get the best-correlated channels for the gap channels. Overlap channels are eliminated as part of the same process. In the overlap regions we select the better-behaved channels. 64 channels are eliminated in this way. There are 331 new gap channels, which are from the UMBC full spectra model. The new gap frequencies are listed in Table 3. The total number of channels increases to 2645 after the gaps are filled.

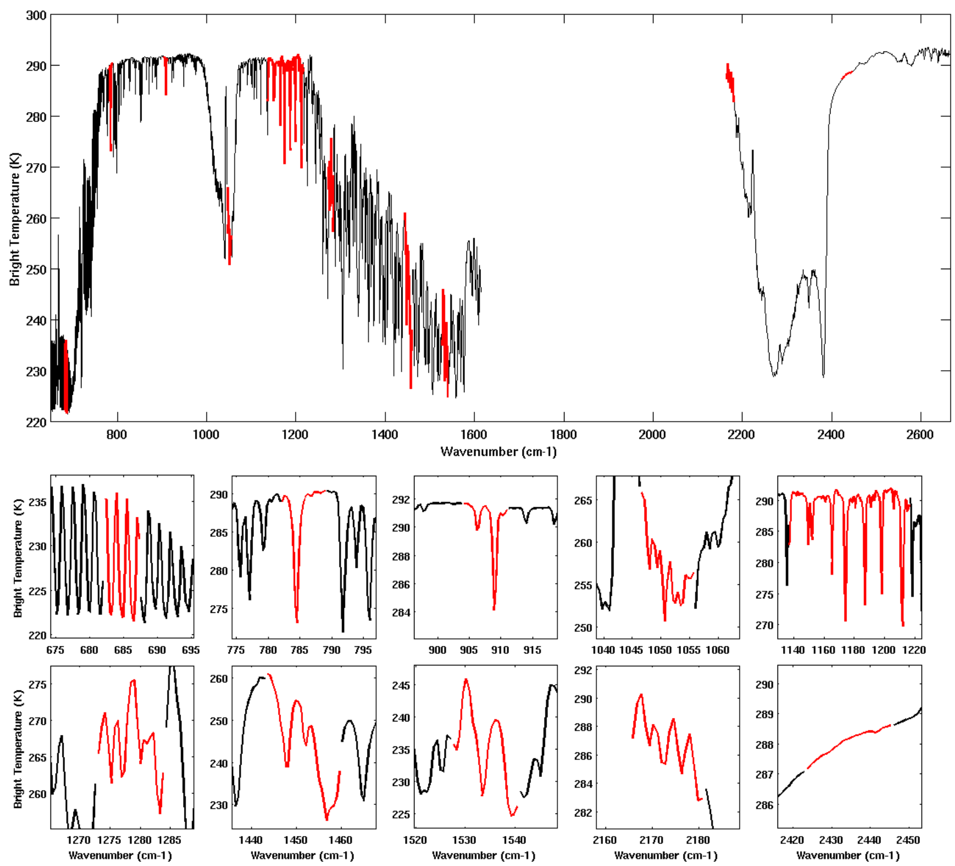

Figure 8. Sample spectrum from the full spectrum training set. The black lines represent the existing channels, and the red lines are the gap channels. Note that the channels with frequencies between 2160 and 2181 cm-1 are not filled in the L1C product.

Table 3. AIRS L1C Synthetic Channel Frequencies (cm^{-1})

Gap #1	Gap #2	Gap #3	Gap #4	Gap #5				Gap #6	Gap #7	Gap #8	Gap #9
682.25	782.22	904.12	1046.65	1137.16	1159.59	1182.03	1204.98	1273.13	1443.65	1527.65	2423.83
682.51	782.56	904.46	1047.10	1137.68	1160.11	1182.55	1205.50	1273.66	1444.23	1528.29	2424.80
682.76	782.89	904.80	1047.55	1138.20	1160.64	1183.07	1206.02	1274.20	1444.80	1528.93	2425.78

683.02	783.23	905.14	1048.00	1138.73	1161.16	1183.59	1206.55	1274.73	1445.37	1529.57	2426.75
683.27	783.56	905.48	1048.45	1139.25	1161.68	1184.11	1207.07	1275.27	1445.95	1530.21	2427.72
683.53	783.90	905.82	1048.90	1139.77	1162.20	1184.63	1207.59	1275.80	1446.52	1530.85	2428.70
683.78	784.23	906.15	1049.34	1140.29	1162.72	1185.16	1208.11	1276.34	1447.09	1531.49	2429.67
684.04	784.57	906.49	1049.79	1140.81	1163.25	1185.68	1208.63	1276.87	1447.66	1532.14	2430.64
684.29	784.90	906.83	1050.24	1141.33	1163.77	1186.20	1209.15	1277.40	1448.24	1532.78	2431.61
684.54	785.24	907.17	1050.69	1141.86	1164.29	1186.72	1209.68	1277.94	1448.81	1533.42	2432.59
684.80	785.58	907.51	1051.14	1142.38	1164.81	1187.24	1210.20	1278.47	1449.38	1534.06	2433.56
685.05	785.91	907.85	1051.59	1142.90	1165.33	1187.76	1210.72	1279.01	1449.96	1534.70	2434.53
685.31	786.25	908.19	1052.04	1143.42	1165.85	1188.29	1211.24	1279.54	1450.53	1535.34	2435.50
685.56	786.58	908.53	1052.49	1143.94	1166.38	1188.81	1211.76	1280.08	1451.10	1535.98	2436.48
685.82	786.92	908.87	1052.94	1144.46	1166.90	1189.33	1212.28	1280.61	1451.68	1536.62	2437.45
686.07	787.25	909.21	1053.39	1144.99	1167.42	1189.85	1212.81	1281.15	1452.25	1537.26	2438.42
686.33	787.59	909.54	1053.83	1145.51	1167.94	1190.37	1213.33	1281.68	1452.82	1537.90	2439.40
686.58	787.92	909.88	1054.28	1146.03	1168.46	1190.89	1213.85	1282.22	1453.39	1538.54	2440.37
686.84	788.26	910.22	1054.73	1146.55	1168.98	1191.42	1214.37	1282.75	1453.97	1539.18	2441.34
687.09	788.60	910.56	1055.18	1147.07	1169.51	1191.94	1214.89	1283.28	1454.54	1539.82	2442.31
687.35	788.93	910.90	1055.63	1147.59	1170.03	1192.46	1215.41	1283.82	1455.11	1540.46	2443.29
				1148.12	1170.55	1192.98	1215.94		1455.69		2444.26
				1148.64	1171.07	1193.50	1216.46		1456.26		2445.23
				1149.16	1171.59	1194.02	1273.13		1456.83		
				1149.68	1172.11	1194.55	1273.66		1457.41		
				1150.20	1172.64	1195.07	1274.20		1457.98		
				1150.72	1173.16	1195.59	1274.73		1458.55		
				1151.25	1173.68	1196.11	1275.27		1459.12		
				1151.77	1174.20	1196.63	1275.80		1459.70		
				1152.29	1174.72	1197.15	1276.34				
				1152.81	1175.24	1197.68	1276.87				
				1153.33	1175.77	1198.20	1277.40				
				1153.85	1176.29	1198.72	1277.94				
				1154.38	1176.81	1199.24	1278.47				
				1154.90	1177.33	1199.76	1279.01				
				1155.42	1177.85	1200.28	1279.54				
				1155.94	1178.37	1200.81	1280.08				
				1156.46	1178.90	1201.33	1280.61				
				1156.98	1179.42	1201.85	1281.15				
				1157.51	1179.94	1202.37	1281.68				
				1158.03	1180.46	1202.89	1282.22				
				1158.55	1180.98	1203.41	1282.75				
				1159.07	1181.50	1203.94	1283.28				
						1204.46	1283.82				

The buddy algorithm used for this gap filling takes advantage of the fact that it is operating on the reconstructed spectrum, and so doesn't have to dynamically adjust for bad channels. It is a simple 4-channel substitution. For each of the 331 gap channels we have 4 channel numbers {ch1, ch2, ch3, ch4}. There are 3 coefficients a1, a2, a3 giving weights for the first 3 channels. The weight for the 4^{th} channel is calculated as a4 = 1.0 – a1 – a2 – a3 so that the sum of the 4 coefficients is guaranteed to be 1.0. The BT for the gap channel then is simply:

BT = a1*BT(ch1) + a2*BT(ch2) + a3*BT(ch3) + a4 * BT(ch4)

7. Summary

An algorithm has been presented which can correct for the effect of time-variable non-Gaussian noise in AIRS spectra and fill these spectra out to a monotonic set with only one gap. This algorithm will be used as the basis for a new AIRS science data product—Level 1C—calibrated and cleaned radiances. That product will be made available to the public in a future version of the AIRS science software that runs at the GES DISC. Most channels have their calibrated radiances untouched from the basic product—Level 1B—calibrated radiances. The worst channels have their radiances replaced by PCA reconstructed values drawing information from strongly correlated channels. This replacement process is performed independently for every spectrum. Gap channels are filled with a simpler algorithm. The Level 1C product is not a substitute for Level 1B (which will remain the primary AIRS Level 1 product), but an enhancement for optional use by users.

7. Summary and Application

The process described here involves two passes. The first pass is computationally cheap and the other is expensive. The “buddy” system algorithm is used in the first pass in both cleaning and gap-fill process. The PCA is computational expensive but affordable in the second pass. It is required in order to detect bad channels which have been slipped through the NEΔT check. There are also two arbitrary thresholds used. One is the local granule average NEΔT cutoff used in the first pass. The other is the difference threshold between the observed brightness temperature and the PCA reconstruction used in the second pass.

We expect that use of the fully complete Level 1C product will significantly enhance the utility of the AIRS instrument for climate studies. AIRS Level 1C data will be more easily cross-calibrated with other instruments. Some subtle instrument effects (such as very small frequency shifts) will be fully accounted for, reducing the likelihood that effects of an instrument feature could be misinterpreted as a climate trend.

References

Chahine, M. T., T. S. Pagano, H. H. Aumann, R. Atlas, C. Barnet, L. Chen, M. Divakarla, E. J. Fetzer, M. Goldberg, C. Gautier, S. Granger, F. W. Irion, R. Kakar, E. Kalnay, B. H. Lambrigtsen, S. Y. Lee, J. Le Marshall, W. McMillan, L. McMillin, E. T. Olsen, H. Revercomb , P. Rosenkranz, W. L. Smith, D. Staelin, L. L. Strow, J. Susskind, D. Tobin and W. Wolf, The Atmospheric Infrared Sounder (AIRS): improving weather forecasting and providing new insights into climate, *Bulletin of the American Meteorological Society*, **87**, 891-894, DOI: 10.1175/BAMS-87-7-891 (2006)

Clerbaux, C., Boynard, A., Clarisse, L., George, M., Hadji-Lazaro, J., Herbin, H., Hurtmans, D., Pommier, M., Razavi, A., Turquety, S., Wespes, C., and Coheur, P.-F.: Monitoring of atmospheric composition using the thermal infrared IASI/MetOp sounder, *Atmos. Chem. Phys.*, **9**, 6041–6054, doi:10.5194/acp-9-6041-2009 (2009).

Gaiser, S. L., Hartmut H. Aumann, L. Larrabee Strow, Scott E. Hannon, and Margaret Weiler, In-Flight Spectral Calibration of the Atmospheric Infrared Sounder, IEEE TRANSACTIONS ON GEOSCIENCE AND REMOTE SENSING, VOL. 41, NO. 2 (2003).

Jolliffe, I. T.,Principal Component Analysis. Springer-Verlag. pp. 487. doi:10.1007/b98835. ISBN 978-0-387-95442-4 (1986).

Strow, L. L., Scott E. Hannon, Margaret Weiler, Kenneth Overoye, Steven L. Gaiser, and Hartmut H. Aumann, Prelaunch Spectral Calibration of the Atmospheric Infrared Sounder (AIRS), IEEE TRANSACTIONS ON GEOSCIENCE AND REMOTE SENSING, VOL. 41, NO. 2 (2003)

Weiler , M. H., Kenneth R. Overoye, James A. Stobie, Paul B. O'Sullivan, Steven L. Gaiser, Steven E. Broberg, and Denis A. Elliott, Performance of the Atmospheric Infrared Sounder (AIRS) in the Radiation Environment of Low-Earth Orbit, SPIE International Symposium on Optical Science and Technology, Vol. 5882, San Diego (2005).

Appendix A: Buddy Cleaning Training Set

A.1 The Night-Time Clear-Sky Buddy Training Set

While the AIRS spectrum consists of 2378 independent spectral channels, the information content of the spectrum as a whole is much less. For all AIRS channels, there exist other AIRS channels that are radiometrically similar. We can therefore replace the radiance of known bad or noisy channels with ones computed using these approximately equivalent channels or "buddies".

We prefer to select buddies for each channel by finding channels with similar statistics over a representative set of observed spectra. However, for our 7 permanently bad channels and about 100 others that are dynamically bad on the day used for this set, we can't do this. So these channels are first replaced using synthetic (simulated) spectra, then we can produce a training set with all channels good for training the final buddy replacement list and the principal component step.

The first training set is derived from AIRS night-time clear-sky spectra model simulation by UMBC radiation transfer model with 49 climatology spectra at satellite zenith angles 0, 10, 20, 30, 40, 50 degrees. Some of the spectra are shown in Figure 3.

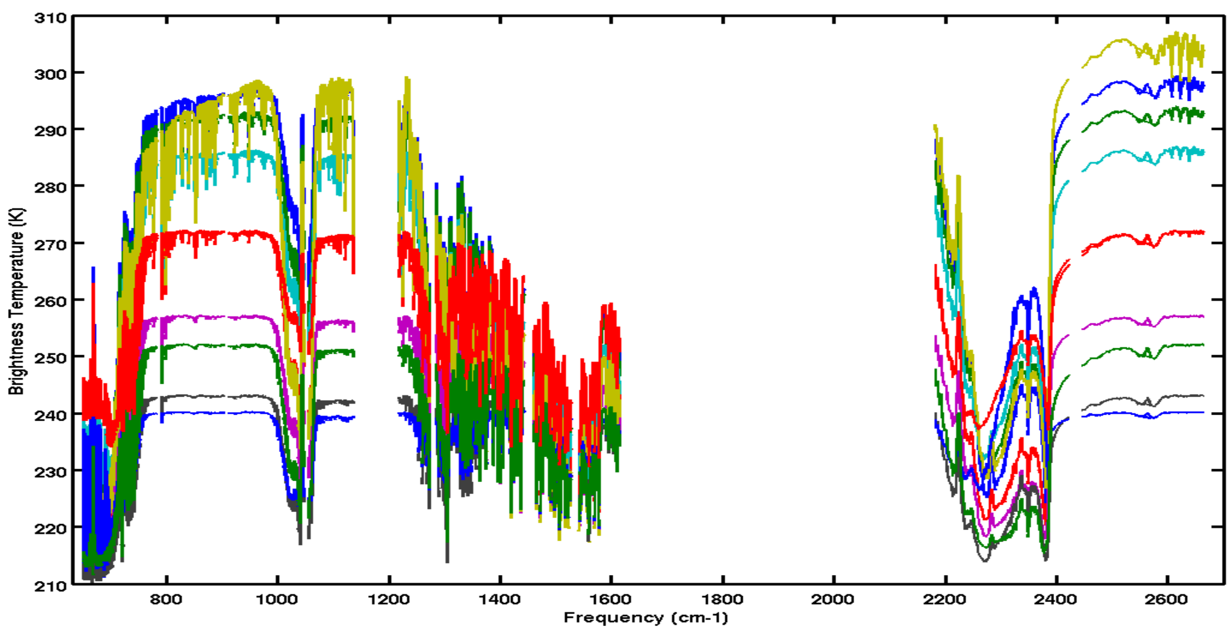

Figure 3. Examples of night-time clear-sky training sets.

A.2 The Global Buddy Training Set

The second training set (Figure 4) uses real AIRS spectra observed on January 1, 2008. There are 21,502 profiles in this training set, and they include day-time and night-time spectra, cloudy and dust cases. This was the largest number of samples in the training set that the PCA training algorithm could accommodate. The spectra were selected such as to represent all the factors that may affect the measurement. Each of the 21,502 spectra was cleaned and filled by the first set "buddy" replacement process. Some of the spectra are shown in Figure 4. This training set was used to find the new set of "buddies" which will be used in L1C process.

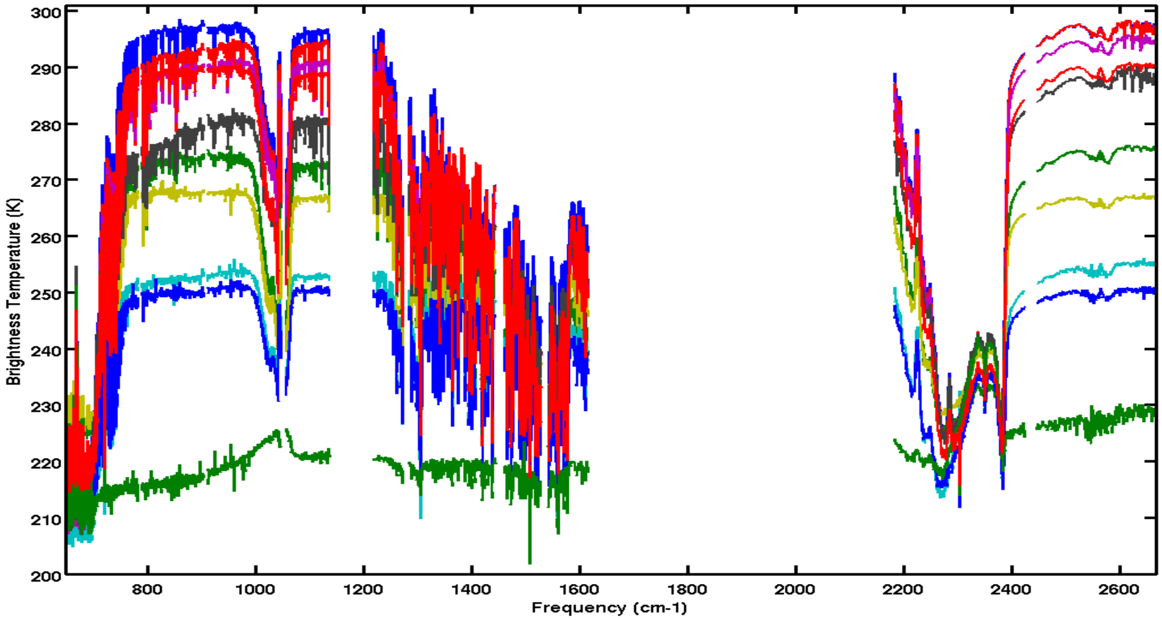

Figure 4. Example of second training sets.

The first step in selection of profiles for inclusion in the training set was to evaluate the median brightness temperature (BT) for each of the 17 detector module for each spectrum. See the Level-1B ATBD for more information on detector modules. Then the median of those medians was calculated as the overall median BT for the spectrum. The key parameters for distinguishing profiles were:

1) overall median BT
2) Delta BT between module M-12 (649-682 cm^{-1}) BT and overall median BT
3) Delta BT between module M-01b (2300-2422 cm^{-1}) BT and overall median BT
4) Delta BT between module M-01a (2546-2665 cm^{-1}) BT and overall median BT
5) Delta BT between module M-02a (2446-2564 cm^{-1}) BT and overall median BT
6) Delta BT between module M-04d (1217-1272 cm^{-1}) BT and overall median BT

6-dimensional bins were created by dividing each parameter into intervals:

1) overall median BT: 2K bins
2-5) delta M-12 through delta M-02a: range of values encountered is divided into 12 equal bins
7) delta M-04d: range of values encountered is divided into 8 equal bins

Then from each non-empty bin, one spectrum was selected at random. Assuming the overall median BT covers a span of about 100 K, this gives 50*12*12*12*12*8 ~= 8 million bins, but most combinations are empty.

Appendix B: PCA Cleaning Training Set

The training set used for PCA is a composite of 2 separate sets derived from simulated artifact-free spectra from 2 sources. The composite overcomes the shortcomings of earlier PCA sets trained on just one or the other. The UMBC set alone failed for desert cases, while the NASA Langley set was weak in the shortwave region for day cases.

Appendix B.1: UMBC PC training set

A training set was received from Larrabee Strow and Sergio Desouza of UMBC and placed at JPL in /asl/data/rtprod_airs/2009/03/01/JPL2834/ . This set represents an entire day of simulated data with 240 6-minute granule files in RTP format. In order to accommodate the use of this data for training a PC set which can be applied to AIRS data throughout the AIRS mission, spectra in the set have different CO_2 levels and spectral shifts over the entire expected range. These spectra have 2834 channels, representing the 2378 AIRS instrument channels and 456 synthetic channels in spectral gaps. For the generation of this training set we used only the original 2378 channels.

Evan Manning wrote the IDL program /home/evan/L1C/global_spectra2.pro to extract from this day of data a PC training set of <~10,000 representative spectra. This program traverses the spectra in time order and assigns to each spectrum a tag which is combined of the following elements:
1) Day vs Night (the divide is solar zenith angle = 90) (2 bins)
2) Land/Sea + Latitude band. 30-degree bins. non-polar bins are divided into land vs sea. Sea has landfrac < 1%. (10 bins)
3) BT1231 (brightness temperature at 1231 cm-1 in 10K increments from [170, 180) to [340, 360)) (19 bins)
4) Tsurf-BT1231 (cloud impact in 10K increments from [-40, -30) to [200, 210)) (26 bins)
5) for each of 5 bands, BTband-BT1231 in 10K increments from [-110, -100) to [100, 110) The bands are [650, 800], [800, 1200), [1200, 1700), [1700, 2400), and [2400, 2700) cm-1. (20 bins per band)

This allows for up to 2*10*19*26*20^5 ~= 30 billion bins but of course most bins are empty.

Each spectrum's tag is compared to the tags of the spectra previously collected. If the tag does not match any then the new spectrum is added to the training set. Note that this procedure makes no explicit provision to ensure that the training set represents the full range of scan angles, CO_2 levels, or spectral shifts but seems likely to contain a good mix as long as the input data does.

The resulting training set is train2.2009-03-01f.h5, with 7377 spectra in 100-spectrum objects.

Appendix B.2: Zhou PC training set

A training set was received from Dan Zhou of NASA Langley Research Center and placed at JPL in /home/hha/DanZhou.AIRS.simulation/ . It is organized in 84 files with names in the pattern "AIRS_RAD_xx_y1_y2.bin". xx is the SARTA satang for all spectra in that file and [y1, y2] is a range of sunangs. Each file contains 35043 spectra, representing the same 35043 representative geophysical states. The first 26600 spectra are under clear conditions, and the rest are under cloudy conditions. An ancillary file named "data_info.asc" contains lat. (col #1), long. (col #2) and cloud optical depth (col #3).

3. Spectra Shift

Evan Manning wrote the IDL program /home/evan/L1C/zhou_training2.pro to extract from this data a PC training set of <~10,000 representative spectra. This program traverses the spectra and assigns to each spectrum a tag which is combined of the following elements:

1) Day vs. Night (Night iff sunang is 95 degrees) (2 bins)

2) Latitude band. 30-degree bins. North and south tropics are combined to a single bin as are north and south midlat. (4 bins)

3) BT1231 (brightness temperature at 1231 cm-1 in 10K increments from [170, 180) to [340, 360)) (19 bins)

4) Cloud optical depth, {[0.0, 0.1), [0.1, 0.3), [0.3, 1.0), [1.0, 3.0), >=3.0} (5 bins)

5) for each of 5 bands, BTband-BT1231 in 10K increments from [-140, -130) to [130, 140) The bands are [650, 800], [800, 1200), [1200, 1700), [1700, 2400), and [2400, 2700) cm-1. (28 bins per band)

This allows for up to 2*4*19*5*28^5 ~= 13 billion bins but of course most bins are empty.

Each spectrum's tag is compared to the tags of the spectra previously collected. If the tag does not match any then the new spectrum is added to the training set. The order of file traversal is manually set to put the most extreme scan angles and solar zenith angles first:

```
'AIRS_RAD_00_00_30.bin',
'AIRS_RAD_60_75_90.bin',
'AIRS_RAD_00_75_90.bin',
'AIRS_RAD_60_00_30.bin',
'AIRS_RAD_40_75_90.bin',
'AIRS_RAD_60_45_60.bin',
'AIRS_RAD_00_45_60.bin',
'AIRS_RAD_20_00_30.bin',
'AIRS_RAD_20_75_90.bin',
'AIRS_RAD_40_00_30.bin',
'AIRS_RAD_40_45_60.bin',
'AIRS_RAD_20_45_60.bin',
'AIRS_RAD_00_95_95.bin',
'AIRS_RAD_60_95_95.bin',
'AIRS_RAD_40_95_95.bin',
'AIRS_RAD_20_95_95.bin'
```

The resulting training set is zhoutrain2.h5, with 13021 spectra in 100-spectrum objects.

Appendix C. Spectrum Shift

Spectrum shifting is not supported in the v6.0 release of AIRS L1C. This section details a candidate algorithm that is currently implemented in software but disabled operationally.

The overall position of the AIRS channels are set with the grating model y-offset, which is an offset distance (micrometers) in the y-axis (dispersed) direction on the focal plane. This y-offset is added to the nominal y-position for the 2378 channels and then run through the AIRS grating model to generate channel frequencies.

Analysis of the AIRS radiances showed the channels were at -13.5 µm in September 2002 and November 2003. The grating model y-offset is calibrated to channel frequency for a reference grating temperature = 155.1325 K. But for unknown reason, at least two of the AIRS modules have shifted a little after launch, and the overall shift after launch is plotted in Figure 2. It shows that the frequency shift is not only seasonal dependent, but also latitude dependent. By regression analysis of the measured frequency shift from 2002 to 2010, the empirical prediction of the frequency shift has been determined by Strow and Hannon from UMBC (ref).

The frequency shift is about 8ppm in ten years, this could introduce 0.3K bias in brightness temperature. The changing frequency in L1B radiance data sets may cause certain confusion for users. Since the frequency shift is minor, we can determine the corrected spectrum using formula (4)

$$R = R_0 + \frac{dR}{d\nu} \cdot \Delta\nu \qquad (4)$$

where R is the radiance at frequency ν after the frequency shift, R_0 is the spectrum before shift, $\Delta\nu$ represents the frequency shift amount, $\frac{dR}{d\nu}$ is the derivative of the spectrum at frequency ν. The $\frac{dR}{d\nu}$ is calculated from the grating model with different y-offset. The grating model provides two series of spectra with y-offset -14 µm and -15 µm respectively. Figure 13 shows the $\frac{dR}{d\nu}$ (y-axis) calculated for the shift from nominal -14 µm to -15 µm shift (black line), and for the shift from nominal -15 µm to -14 µm shift (blue line) verses the derivative calculated from the spline interpretation (x-axis). The linear fit to the data is expressed in formula (5), which shows the linear relationship between the calculated derivative and the derivative from spline interpretation. This relationship is critical to our frequency shift calculation since the slope or derivative is different for different spectrum, the spline derivative has to be calculated dynamically. Finally the shifted spectrum is shown in formula (6).

$$\frac{dR}{d\nu} = a\frac{dR^S}{d\nu} + b \qquad (5)$$

$$R = R_0 + (a\frac{dR^S}{d\nu} + b) \cdot \Delta\nu \qquad (6)$$

3. Spectra Shift

The shift parameters *a* and *b* are frequency dependent variables, their values for each channel are stored in the look up table.

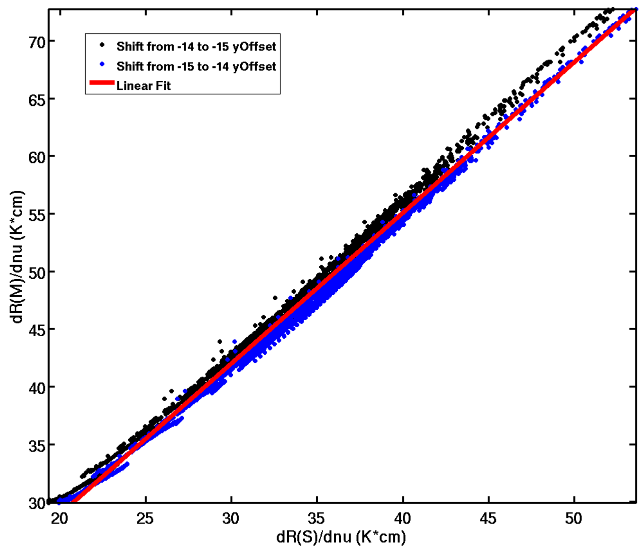

Figure 13. Scattering plot of the calculated spectrum derivative vs. the derivative from -14 μm to -15 μm shift (black line) and -15 μm to -14 μm shift (blue line). Red line is the linear fit to the data.

The model spectra are used to validate the algorithm. Figure 14 indicates that the averaged brightness temperature difference between shifted and true spectrum is much smaller around 5 mK as comparing to other algorithm with difference around 50 mK such as spline fit only. The distribution of the difference clearly shows that our algorithm is the best as comparing to spline re-sampling method and non-shift.

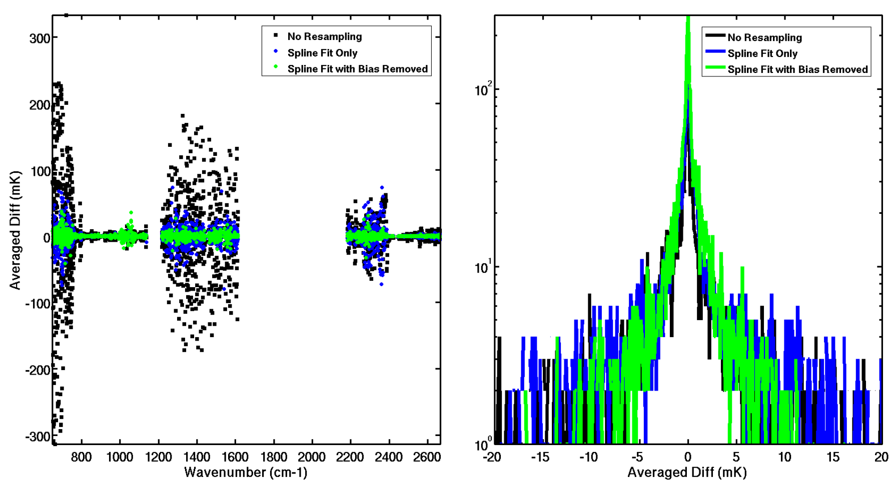

Figure 14. Averaged brightness temperature difference between shifted and original spectrum (left panel) and its distribution (right panel) of the brightness temperature vs. frequency from model spectra with no resampling (black dot or line), spline fit (blue dot or line), and spline fit with bias removed (green dot or line).

3. Spectra Shift

Figure 15 presents that the root mean square difference (RMS) of the brightness temperature between shifted and true spectrum is much smaller around 1 mK as comparing to the difference around 10 mK for spline fit. The distribution of the difference also shows that our algorithm improve the final spectrum.

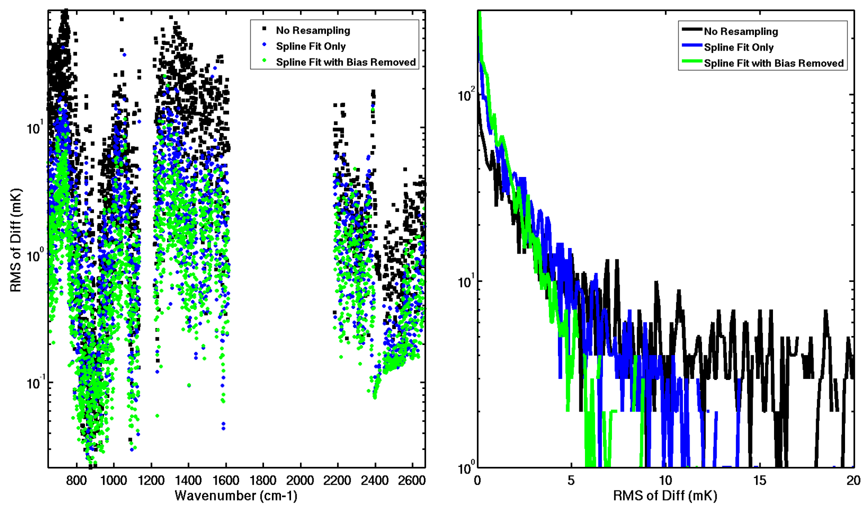

Figure 15. Root mean square difference (RMS) of the brightness temperature between shifted and original spectrum (left panel) and its distribution (right panel) of the brightness temperature vs. frequency from model spectra with no resampling (black dot or line), spline fit (blue dot or line), and spline fit with bias removed (green dot or line).