

L.S. Bernstein, X. Jin, B. Gregor and S.M. Adler-Golden, *The Quick Atmospheric Correction (QUAC) Code: Algorithm Description and Recent Upgrades*, SPIE Optical Engineering, Vol. 51(11), 111719 (2012).

Copyright 2012 Society of Photo-Optical Instrumentation Engineers. One print or electronic copy may be made for personal use only. Systematic reproduction and distribution, duplication of any material in this paper for a fee or for commercial purposes, or modification of the content of the paper are prohibited.

[http://dx.doi.org/DOI# http://dx.doi.org/10.1117/1.OE.51.11.111719](http://dx.doi.org/DOI#http://dx.doi.org/10.1117/1.OE.51.11.111719)

See the next page.

Optical Engineering

SPIEDigitalLibrary.org/oe

Quick atmospheric correction code: algorithm description and recent upgrades

Lawrence S. Bernstein
Xuemin Jin
Brian Gregor
Steven M. Adler-Golden

Quick atmospheric correction code: algorithm description and recent upgrades

Lawrence S. Bernstein

Xuemin Jin

Brian Gregor

Steven M. Adler-Golden

Spectral Sciences, Inc.

4 Fourth Avenue

Burlington, Massachusetts 01803-3304

E-mail: larry@spectral.com

Abstract. The quick atmospheric correction (QUAC) code performs atmospheric correction on multi- and hyperspectral imagery spanning all or part of the visible and near infrared–short wave infrared spectral range, ~400 – 2500 nm. It utilizes an in-scene approach, requiring only approximate specification of sensor band locations (i.e., central wavelengths) and their radiometric calibration; no additional metadata is required. Because QUAC does not involve first principles radiative-transfer calculations, it is significantly faster than physics-based methods; however, it is also more approximate. We present a detailed description of the QUAC algorithm, highlighting recent accuracy improvements. Example results for several multi- and hyperspectral data sets are presented, and comparisons are made to more rigorous correction approaches. © 2012 Society of Photo-Optical Instrumentation Engineers (SPIE). [DOI: [10.1117/1.OE.51.11.111719](https://doi.org/10.1117/1.OE.51.11.111719)]

Subject terms: atmospheric correction; quick atmospheric correction; remote sensing; hyperspectral; sensor; fast line-of-sight atmospheric analysis of spectral hypercubes; reflectance; algorithm.

Paper 120309SS received Mar. 1, 2012; revised manuscript received Jun. 27, 2012; accepted for publication Jul. 10, 2012; published online Jul. 31, 2012.

1 Introduction

Views of Earth's surface from aircraft and spacecraft are substantially degraded by the presence of the atmosphere. This degradation includes attenuation of reflected light as well as loss of contrast due to sunlight scattering by atmospheric aerosols and molecules. Remote-sensing applications routinely require that these atmospheric effects be removed from the imagery, to retrieve the inherent spectral reflectance of the surface materials. The process for removing atmospheric effects is referred to as atmospheric correction or compensation. An example of the application of atmospheric correction to hyperspectral imagery is shown in Fig. 1.

Many atmospheric correction methods and algorithms exist, including those based on first-principles radiation transport (RT) calculations,^{1–11} and empirical approaches such as the empirical line method (ELM),^{12,13} which relies on two or more known reflectances in the image. More approximate in-scene methods, such as the IAR (internal average reflectance)¹⁴ and FF (flat field)¹⁵ approaches, are often used for quick-look purposes. However, none of these methods provide the ideal combination of high accuracy, high computational speed, and independence from prior knowledge (ground truth, sensor calibration, measurement geometry, etc.). In 2004, we introduced a new atmospheric correction algorithm and code called quick atmospheric correction (QUAC®),^{16–19} which comes close to satisfying these attributes.

QUAC is an in-scene approach, requiring only approximate specification of sensor band locations (i.e., central wavelengths) and their radiometric calibration; no additional metadata is required. Because QUAC does not involve first principles RT calculations, it is significantly faster than physics-based methods; however, it is also more approximate.

Previous comparisons to the most widely used physics-based code, fast line-of-sight atmospheric analysis of spectral hypercubes (FLAASH),¹¹ have shown that the absolute accuracy of QUAC is $\sim \pm 15\%$ with respect to FLAASH-retrieved reflectances for a few well-calibrated data sets with well-characterized measurement conditions.²⁰ A comparison of QUAC and FLAASH is presented in Fig. 2, which exhibits the general trend that spectral shapes agree well and the largest differences are in the absolute normalization of the QUAC results. Whereas the accuracy of the physics-based methods is directly tied to the accuracy of the sensor calibration and measurement geometry, QUAC performance will not significantly degrade as sensor and measurement uncertainties increase. Finally, in contrast to physics-based methods, which require the presence of specific bands to correct for water absorption and aerosol scattering, QUAC works with any collection of visible and near infrared–short wave infrared (VNIR–SWIR) (e.g., VNIR only, SWIR only) bands for both multispectral and hyperspectral sensors.

In recent years, there has been a rapid rise in the number of fielded HSI systems that utilize QUAC. This has provided a unique opportunity to evaluate QUAC performance over a wide range of measurement conditions (different sensors, different types of scenes, etc.). QUAC usually provides good results; however, it can fail badly on occasion. A primary focus of our latest work is to understand the reasons for these failures and upgrade QUAC to yield good results for these and similar cases. We have also improved the computational speed by a factor of approximately 2. The CPU time required to correct a typical AVIRIS data cube (500 × 500 pixels, 220 bands) is now ~ 3 s for a PC (64-bit, 3.2-GHz processor, 8 Gb ram).

The original IDL/ENVI version of QUAC is distributed as part of the Exelis ITT Visualization Information Systems IDL/ENVI imaging processing software package. We

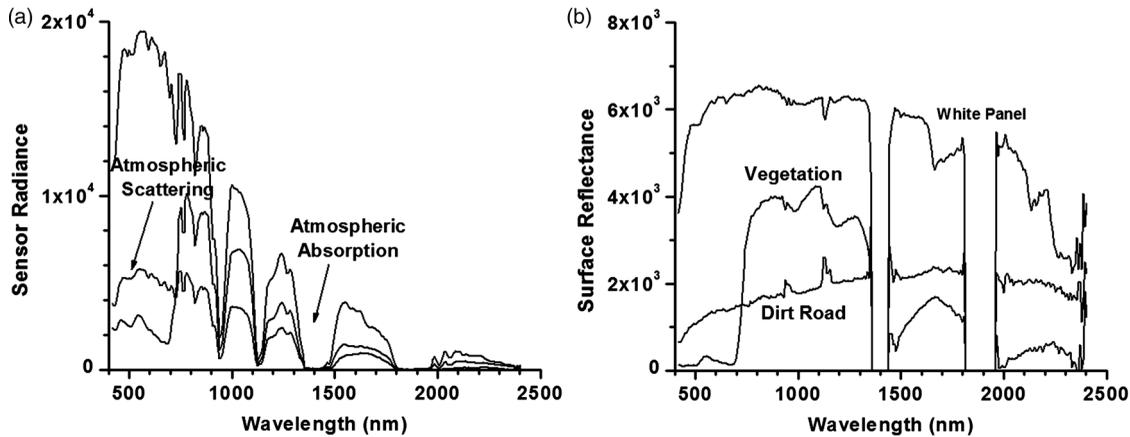


Fig. 1 An example of atmospheric correction for hyperspectral imaging (HSI) sensors, showing several at-sensor pixel spectra (a) and the corresponding atmospherically corrected surface reflectance spectra (b). Because of very strong atmospheric attenuation, due to water vapor, in the 1400 and 1900 nm regions, it is not possible to correct these regions, hence the presence of the zeroed-out bands in the corrected data. In spectral regions of moderately strong atmospheric attenuation, such as the 940 and 1130 nm water bands, residual atmospheric features are occasionally present in some of the corrected pixel spectra (e.g., Dirt Road). This arises because QUAC defines a scene-average gain function and thus does not account for the variability of humidity within a scene or spectral smile across the focal plane.

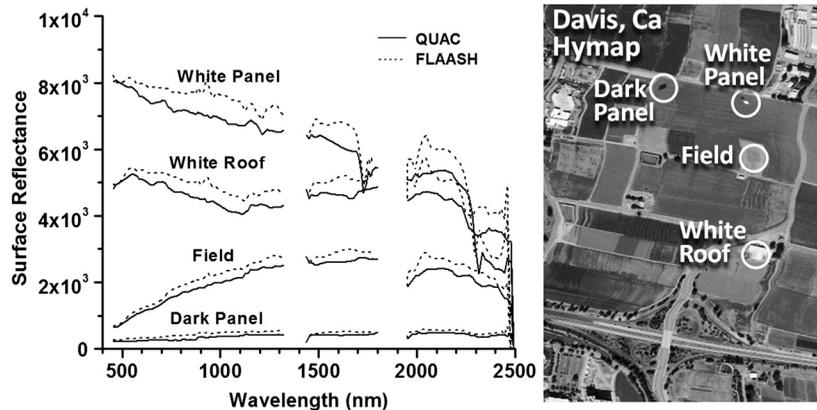


Fig. 2 Comparison of QUAC and FLAASH results for a well-ground-truthed HSI data collection over Davis, CA.²¹ The grayscale image was rendered at 600 nm, where vegetation (trees, crops, etc.) are the darkest materials in the image.

anticipate that the newest version of QUAC will be available in IDL/ENVI within the next year. Because of the rapidly growing need for performing fast atmospheric correction on board the sensor platform, we have also developed a C++ version of QUAC suitable for integration into automated data processing systems.

The organization of this paper is as follows. The next section describes the QUAC algorithm, highlighting recent improvements. This includes a more in-depth description of the underlying algorithm than previously reported. The section following presents example results highlighting the algorithm upgrades and includes comparisons to more rigorous correction approaches. Next, we assess the accuracy of QUAC and its sensitivity to several algorithm and sensor parameters. The final section provides a summary of key points and briefly describes anticipated future research activities.

2 Algorithm Description

The basic physics behind atmospheric correction is depicted in Fig. 3. The observed spectral radiance, L_{obs} , for a pixel

with surface reflectance, ρ_{sur} , is the sum of the three paths in Fig. 3:

$$L_{\text{obs}} = (A + C\rho_{\text{ave}}) + B\rho_{\text{sur}}. \quad (1)$$

The components in the term $(A + C\rho_{\text{ave}})$ are grouped together because they tend to be approximately constant over an image and thus can be considered as an offset common to all the image pixels. This simple linear relationship can be rearranged to express the retrieved surface reflectance in terms of the observed signal and derived atmospheric parameters,

$$\rho_{\text{sur}} = \text{Gain}(L_{\text{obs}} - \text{Offset}), \quad (2)$$

where $\text{Gain} = 1/B$ and $\text{Offset} = (A + C\rho_{\text{ave}})$. For a physics-based approach, A , B , and C are retrieved by comparison of certain spectral features to those predicted by RT calculations. For QUAC, we determine these parameters directly from the in-scene spectral data and a key underlying assumption.

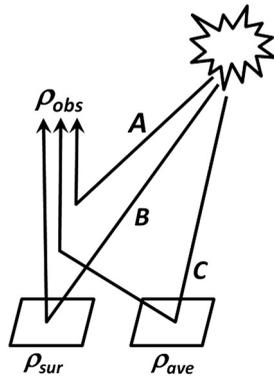


Fig. 3 The three types of paths, A, B, and C, that solar photons can travel on their way to a remotely located observer, where ρ_{sur} is the intrinsic reflectance of the observed surface pixel, ρ_{ave} denotes the spatially averaged reflectance of the surrounding pixels, and L_{obs} is the at-sensor radiance corresponding to the observed surface pixel.

For QUAC, the Gain and Offset are given by

$$\text{Gain} = \frac{\langle \rho_{end} \rangle_{lib}}{\langle (L_{obs} - C\rho_{ave})_{end} \rangle}, \quad \text{and}$$

$$\text{Offset} = \min(\text{pixel value for each band}), \quad (3)$$

where $\langle \rho_{end} \rangle_{lib}$ is the average of the endmember spectra derived from representing a reference library of material reflectance spectra and $\langle (L_{obs} - C\rho_{ave})_{end} \rangle$ is the average of a collection of endmembers retrieved from the observed, in-scene pixel spectra. An endmember represents a unique spectrum from a collection of spectra. In most cases, linear combinations of a small number of endmember spectra ($\sim 10 - 100$) can accurately represent a large number of spectra ($> 10,000$) associated with a spectral library or image. In QUAC, we use the SMACC (sequential maximum angle convex cone) code to find endmembers.^{22,23}

The key QUAC assumption, which empirically holds for most scenes, is that the average of diverse endmember reflectance spectra, excluding highly structured materials (e.g., vegetation, shallow water, mud), is always the same. More specifically, every image is assumed to contain at least a handful (~ 10 or more) of spectrally diverse materials whose average reflectance spectrum can be taken as a “universal” reference. The materials may include both natural and manmade materials, such as a dirt field, a water body, rocks, cars, roofs, or roads. It is unusual that this material diversity condition is not met, but it can occur; for example, in some all-water or all-desert scenes. However, such imagery is typically of much less interest for remote sensing.

2.1 Reference Material Spectral Library

The “universal” reference spectrum is derived by finding and averaging endmember spectra from a diverse collection of natural and manmade library reflectance spectra. We compiled a reference library (see Fig. 4) from the spectral libraries provided with ENVI.²⁴ Since the endmember representation of this library weeds out nearly degenerate spectra, it wasn't necessary to put a great deal of effort into selecting the library entries. However, the issue of whether a different selection process could produce improved QUAC results has been addressed, as discussed below.

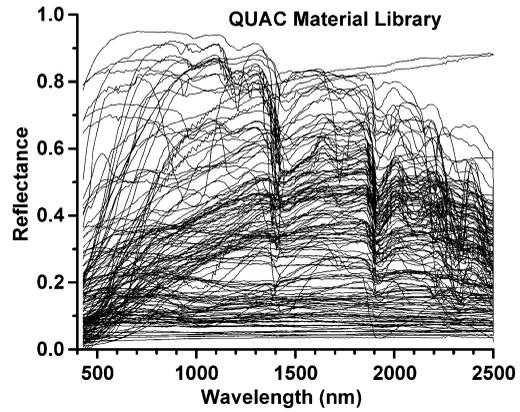


Fig. 4 Reflectance spectra for the 168 natural and manmade materials selected for the QUAC reference material library.

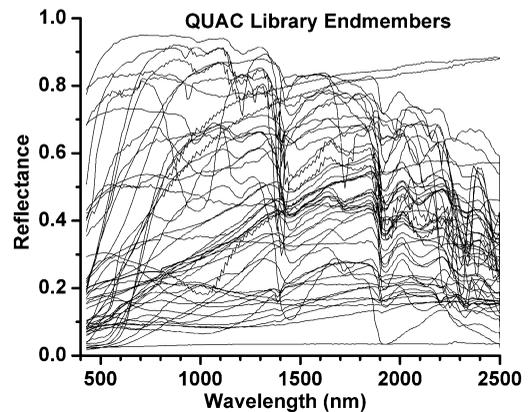


Fig. 5 Reflectance spectra for the 50 SMACC endmembers selected to represent the QUAC reference material library.

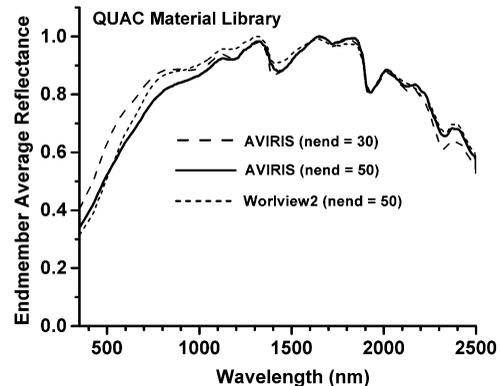


Fig. 6 Sensitivity of the average of the selected library endmembers to differences in the subset of bands used for endmember selection (e.g., AVIRIS vs. WorldView2) and on the number of endmembers used in the average. The peak value of each curve has been normalized to 1.0.

Typically, 50 scene endmembers are used in the correction process, and the reference spectrum is also based on the same number of endmembers. The representation of the reference library in terms of 50 endmembers is shown in Fig. 5. The average of these spectra provides the reference correction spectrum displayed in Fig. 6. As discussed later, both the library and in-scene endmembers are selected based on a small number, typically around four, of all the available

sensor bands. The curves in Figs. 5 and 6 are based on this subset of bands; however, for display purposes, all the spectral points for the library endmember and average spectra are plotted. For each run, QUAC uses both the same number of endmembers and also the same selection of sensor-specific spectral channels for determining the image and reference library endmembers. While this results in relatively small variations in the reference curves, its impact is even smaller on the retrieved image reflectances, because the variations are common to the library and image endmembers and they ratio out in the Gain determination [Eq. (3)].

The general shape of the reference reflectance spectrum has a simple physical origin. The decrease toward the long-wavelength edge arises because the molecular constituents of materials have relatively strong NIR vibrational absorption features that increase in strength with increasing wavelength. The decrease toward the short-wavelength edge arises because the molecular constituents have strong electronic absorption features that increase in strength with decreasing wavelength. Although we normalize the peak of this curve to unity for reasons discussed later, it is important to note that the peak average reflectance is ~ 0.4 , as can be inferred visually from Fig. 5.

Why does QUAC work, or equivalently, why should there be a “universal” reference spectrum? There is no first-principles reason for this, only a qualitative explanation of why an approximately “universal” curve is plausible. As can be seen by inspection of Fig. 4, a collection of a large number of different material reflectance spectra nonuniformly span the $\sim 0 - 0.8$ range of reflectance values. As evident in Fig. 5, the library endmembers span the space in a more uniform manner. We can expect most scenes to exhibit the same general behavior, that is, to contain a diversity of materials whose endmembers approximately spans the 0–0.8 reflectance range. Near-spectral duplicates and differences in abundances do not matter, since these issues are negated by finding the endmember representation of the image pixels.

There are additional issues for real scenes and real sensor data that need to be treated in the endmember selection process. For example, real scenes can exhibit solar glints, whose reflectances can far exceed 0.8. These excessively bright pixels need to be filtered out before endmember selection. Real sensors can have “bad” pixels with unphysical, highly structured spectra that can introduce spurious features into the

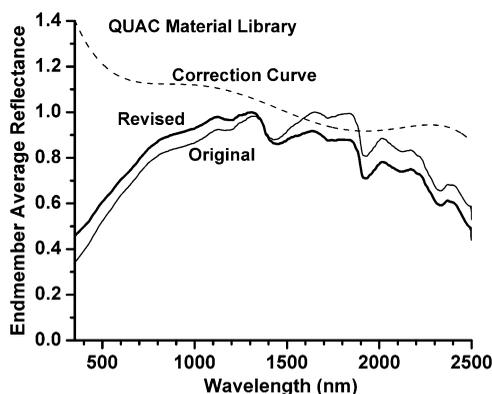


Fig. 7 Revision of the original endmember average spectral reflectance based on a best fit of the QUAC reflectance to several HSI data sets with high-quality ELM or FLAASH atmospheric corrections.

Gain curve. The methods used in QUAC to filter out these and other types of spurious spectra are discussed below.

QUAC is not a rigorous, first-principles algorithm, and there is no reason to expect that there is a single universal reference curve that will give exact results for all images. Therefore, the issue at hand is whether, on average, we can improve on the legacy reference curve. We addressed this issue by first finding a selection of hyperspectral images that were well corrected using either the ELM or FLAASH approach. We identified six corrected AVIRIS and COMPASS data cubes that were deemed of sufficient quality. We then compared the QUAC gain curves to the “truth gain” curves and determined a modest correction to the QUAC reference curve that resulted in better results, averaged over the six truth cubes. The legacy and revised reference spectra are presented in Fig. 7.

2.2 QUAC Process Flow

The implementation of the QUAC algorithm into software is summarized in the flow chart displayed in Fig. 8. This represents the major steps in the process and closely, but not identically, corresponds to modules in the code. We will highlight some of the key features associated with each process. Based on the number and location of the wavelength bands in the image, QUAC can identify the associated sensor, or a sufficiently close surrogate for the correction process. While new sensors can easily be added to QUAC, this is generally not needed, since the correction process depends only loosely on the sensor class properties.

The Offset calculation involves a number of data-conditioning steps to ensure that a valid baseline is determined. This includes (1) removal of border pixels, (2) averaging of adjacent pixels in a line, (3) rejecting values ≤ 0 , and (4) median filtering of adjacent pixels to remove spike artifacts. This module also finds a maximum value in a single spectral channel near the peak of the reference spectrum. As discussed below, this maximum value is used to convert the data to approximate reflectance units before endmember selection.

Before endmember selection the data are converted to an approximate 0–1 reflectance scale by dividing by the

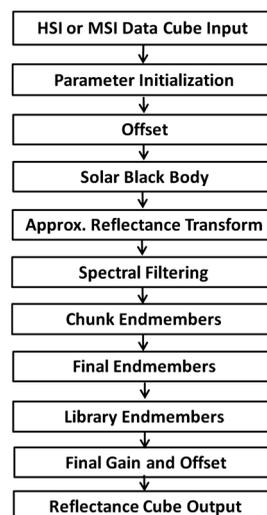


Fig. 8 Overview of the key steps in the QUAC software process flow.

aforementioned maximum value and an approximate solar blackbody curve normalized to unity at its peak value. It is not necessary, nor is it desirable, to use the exact solar blackbody curve. We adjust the effective solar temperature to yield approximate reflectance curves that span about the same range of values throughout the spectral domain of the sensor. The effective solar temperature varies between 4000 and 4500 K, depending on the sensor type, as contrasted to the actual solar temperature of 5700K. There are two important reasons for transforming to a reflectance scale: (1) it simplifies the process of setting spectrally dependent filter thresholds, since all sensors are put on a common scale, and (2) it ensures that all spectral regions are comparably weighted in the endmember selection process. The latter is particularly important because it maintains consistency with the selection of library endmembers, which is based on reflectance values.

It is important to filter out spectra that can introduce undesirable features and biases into the Gain curve. The most common example is vegetation, which has a rapidly rising red edge around 700 nm (see Fig. 1). Vegetation is often present and exhibits a lot of spectral variability, which means that many vegetation spectra would be selected as endmembers. This would produce a strong edge feature in the Gain curve around 700 nm and would result in imbalance in the Gain to either side of the red edge. However, the red edge makes it easy to implement a vegetation filter based on the normalized difference vegetation index (NDVI) for each pixel,

$$\text{NDVI} = \frac{R(850 \text{ nm}) - R(650 \text{ nm})}{R(850 \text{ nm}) + R(650 \text{ nm})}, \quad (4)$$

where $R(850 \text{ nm})$ denotes the approximate reflectance at a particular wavelength. In QUAC we discard pixels (by filling with zeros) with $\text{NDVI} > 0.17$. As shown later, muddy or very shallow water can also cause a similar problem and is filtered out based on a corresponding index we developed called NDMI (normalized difference mud index),

$$\text{NDMI} = \frac{R(795 \text{ nm}) - R(990 \text{ nm})}{R(795 \text{ nm}) + R(6990 \text{ nm})}, \quad (5)$$

where pixels are discarded for $\text{NDMI} > 0.3 - 0.5$ (depending on the sensor). Because the location of the mud blue edge at $\sim 900 \text{ nm}$ can vary, depending on the particulate concentration, we implemented two mud filters, with the second one defined at 1000 and 1600 nm. In the future, if other commonly occurring and abundant materials are found to present a problem, it is simple to construct and add a new filter. It is easy to recognize when this problem exists, since many instances of the undesired material will be present in the corrected endmember spectra.

Before endmember processing, the image is divided into smaller chunks. There are several reasons for this: (1) the full images may be too large to fit in the available memory, which will result in disk thrashing and a large slowdown in processing speed, and (2) we have found empirically that first finding an excess number of chunk endmembers and then down-selecting to the final number of endmembers (i.e., finding the endmembers of the chunk endmembers) generally yields better results than processing the entire image all at once. After

trying a number of chunking strategies, we have settled on dividing each data cube into 50 chunks and finding 20 endmembers for each chunk. Thus, a total of 1000 chunk endmembers is available for the selection of the final subset of endmembers used for the Gain curve. We currently use 50 endmembers, but the results are insensitive to the exact number.

For computational efficiency, the endmembers are found based on a small subset of the available channels for most sensors. For VNIR–SWIR HSI sensors, we typically use four channels in the 1000- to 2500-nm spectral region, which are approximately centered in atmospheric transmission windows. We do not presently select channels in the visible (VIS) spectral region; this is an important outcome of the recent work. We determined that this significantly improved the accuracy and robustness of the results over the entire spectral domain, and often led to large improvements in the VIS region. Although it seems counterintuitive that ignoring the VIS channels should produce better results, it can be understood as follows. As discussed earlier, the shapes of the reflectance curves in the VIS and NIR–SWIR spectral regions arise from different physical processes. Consider a scene containing a number of cars painted different colors. In the VIS region, the car spectra will be quite different, because (by design) the cars have different colors. However, in the NIR–SWIR spectral domain, the car spectra will be nearly identical, because the spectrum is dominated by the paint binder and not the coloring agents. If a VIS channel were included in the endmember selection, then all the car spectra would be selected. However, the degeneracy of the car spectra in the NIR–SWIR would introduce a bias into the endmember average in this spectral region. This should not be construed to mean that VIS channels should never be used for endmember selection. It only applies when a sensor spans both the VIS and NIR–SWIR spectral regions.

Before the final scene endmembers are determined, we filter out excessively bright pixels due to either glints or channel saturation. This is accomplished through a median filter applied to each channel. If any channel value of an endmember spectrum is > 2.25 times its median value, then that endmember is excluded from the final selection process.

As mentioned earlier, the library endmembers are selected using the same selection channels and the same number of endmembers as for the scene endmembers. Once the library and image endmembers are found, a preliminary Gain curve is determined via Eq. (3). There are a few refinements applied to this curve before it is used for the final atmospheric correction. One concerns removal of residual vegetation effects due to the presence of too much vegetation in the baseline spectrum. In highly vegetated scenes, with moderate to low visibility, aerosol scattering introduces a large vegetation component into the adjacency contribution. This is easy to spot in the corrected endmember spectra, as it appears as an inverted vegetation spectral residual superimposed on the endmember spectra. This problem is remedied by subtracting a small amount of the vegetation spectrum (typically a few percent) from the baseline; the amount subtracted is determined by the depth of the red edge in the endmember spectra. A second refinement relates to further improvement in the Gain curve in the VIS region. Even after the improved endmember channel selection discussed above, there was



Fig. 9 Grayscale rendering at 600 nm of the AVIRIS data for the Harrisburg, PA, airport. Vegetation (trees and grassy fields) correspond to the darkest materials for this display band.

still a residual tendency for the VIS region reflectances to be too large and increase toward the short-wavelength edge of the sensor. We found empirically that, in many instances, a better correction result could be obtained by constraining the Gain curve to be constant below ~ 650 nm. Based on this observation, we arrived at the following approximation for the modified Gain < 650 nm,

$$G(\lambda)_{\text{mod}} = \{G(\lambda)G(650 \text{ nm})\}^{1/2}. \quad (6)$$

As alluded to above, we need to set the absolute normalization for the Gain curve. We define two potential normalization values, based on different assumptions, and then select the preferred value. One method is defined by the absolute scale for the reference spectrum (i.e., before the peak value is normalized to unity). The other is based on an assumed typical peak value for an average “pure” green vegetation spectrum, ~ 0.4 at 850 nm. Empirically, we have found that the preferred method is the one that yields the smaller reflectances value for the specific data cube. If a data cube does not contain sufficient pure vegetation, typically ~ 100 pixels, then we use the other method.

For any QUAC run, the user can request plots that display some of the key steps in the correction process. We consider the example of an AVIRIS, 220 bands from $\sim 400 - 2500$ nm,²⁵ over flight of the Harrisburg, PA, airport. A grayscale rendering of the image is shown in Fig. 9 and the process plots are displayed in Fig. 10. Plot (a) shows that the baseline, in radiance units, rise steeply toward the visible and then drops off steeply in the near-UV close to the short-wavelength edge. The rapid rise is due to a combination of aerosol and molecular scattering, which increases with decreasing wavelength, and the solar black body function, which peaks in the VIS. The decline toward the edge is due to increased attenuation and the decline of the solar blackbody function. Plot (b) shows the 1000 chunk endmembers, which have been normalized for solar and maximum values. They display a comparable span of approximate reflectances for all channels, spanning the $\sim 0 - 1$ reflectance range in a complete manner (i.e., no large gaps). Plot (c) shows that the final 50 endmembers span the space in a more uniform manner, like the library endmembers shown earlier. Plot (d) shows an averaged pure vegetation spectrum which is used for the vegetation baseline correction. This correction

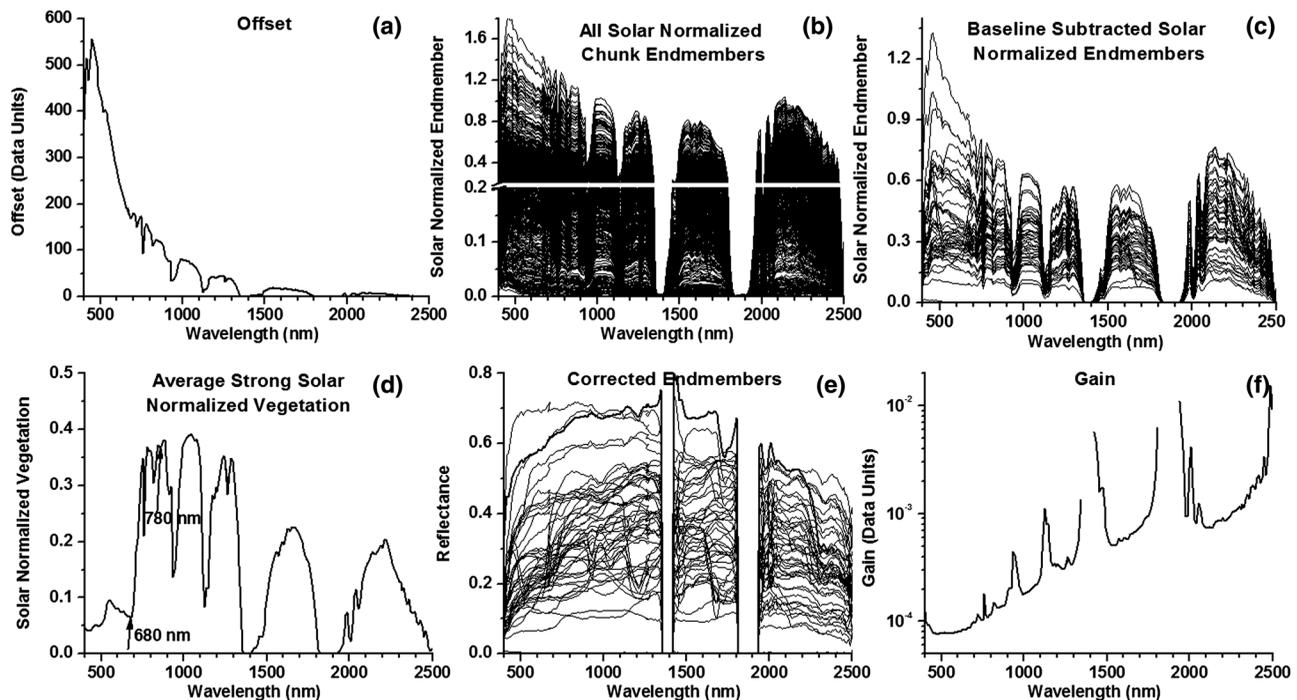


Fig. 10 QUAC output plots illustrating key steps in the process flow, including the offset spectrum determined by the lowest valid data value for each wavelength channel (a), the 1000 initial endmembers determined for the 50 chunks (note y-scale break at 0.2 for display purposes) (b), the final 50 endmembers selected from the 1000 initial endmembers (c), the average pure vegetation spectrum, also showing the blue and red wavelengths used for the NDVI-based vegetation filtering (d), the final 50 atmospherically corrected endmembers (e), and the Gain spectrum (f).

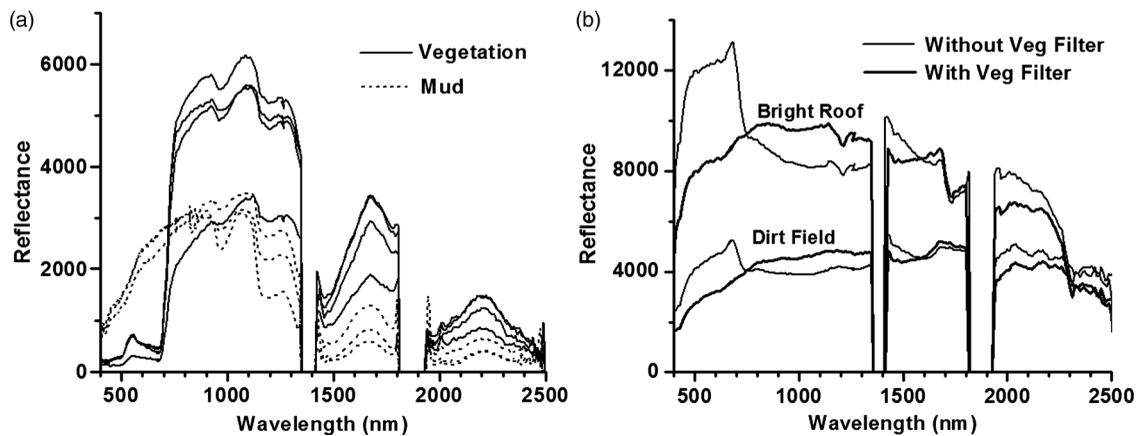


Fig. 11 Vegetation and mud can be abundant materials in a scene and exhibit highly variable spectra, as highlighted in (a). If these spectra are not filtered before selection of the in-scene endmembers, then spectral artifacts will be present in the corrected reflectances, as exemplified by the strong visible “finger” shown in (b) for the AVIRIS imagery in Fig. 9.

was not required for this data cube, as there is no significant vegetation contribution to the baseline spectrum in (a). Plot (e) shows the corrected final endmembers. One indication of a good correction is that the corrected endmembers approximately span the 0–0.8 reflectance range in a uniform manner and without a lot of highly structured endmembers. The zeroed out regions around 1400 and 1900 nm correspond to strong atmospheric absorption regions [see (b)], where the signals are too small to yield acceptable corrections. Panel (f) shows the final Gain curve. The spikes correspond to the edges of strong absorption regions, which contain sufficient signal to enable a good correction.

3 Example Results from Algorithm Upgrades

We highlight the results of some of the key algorithm upgrades discussed above, such as the mud filtering and improved performance in the VIS spectral region. Figure 11 illustrates the need for removing abundant and highly variable spectra with strong, localized features, such as the sharp dark-to-bright transitions that occur for mud and vegetation. Without filtering, a significant number of mud and vegetation endmembers are selected, and with filtering, none are selected. As seen in Fig 11, when vegetation endmembers are included in the correction process, the reflectances are quite distorted, exhibiting un-physically high values in the VIS spectral region. If the mud pixels are not filtered, the reflectances in the VNIR–SWIR spectral region are too

high. It is not necessary to devise a filter for every possible, strongly featured spectra; it is needed only for those that are both abundant and highly variable within a scene. A few strongly featured endmembers will not significantly skew the correction because their effect will be diluted by the much larger number of smoother endmembers.

The improved results in the VIS region arise both from using only NIR–SWIR channels for the endmember selection and also from using the Gain curve modification <650 nm. Previous to these upgrades, it was occasionally found that the retrieved reflectances exhibited a spike-like feature in the VIS region, as can be seen in Fig. 12 for the Bright Roof spectrum.

The final example illustrates how the mud filters can be used in a more general way to filter out sensor-related spectral artifacts. The WorldView-2 multispectral sensor has 8 bands covering the 400- to 1000-nm spectral region.²⁶ The location and widths of these bands mean that the sensor is not sensitive to the effects of muddy water. Therefore, we can use the two mud filters for other purposes. Using the full complement of bands without filtering to find the endmembers results in the selection of many spurious endmembers, as is evident in the lower-right plot in Fig. 13. This multitude of “sawtooth” endmembers, which result from slight misregistrations between the focal planes, is common to other WorldView-2 images. The vast majority of pixels in WorldView-2 images do not display this anomaly, but it is the

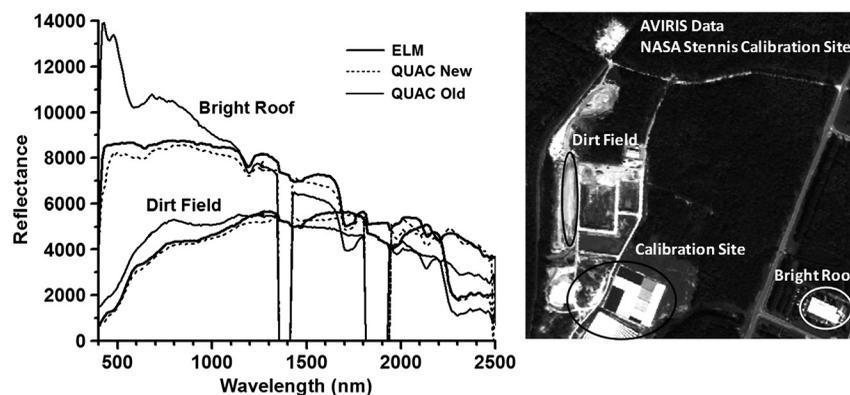


Fig. 12 Illustration of the UV-VIS spike problem (thin solid line) often encountered with the original version of QUAC (QUAC Old, thick solid line). This problem has been corrected (dotted line) by modifying the channels selected for finding the endmembers.

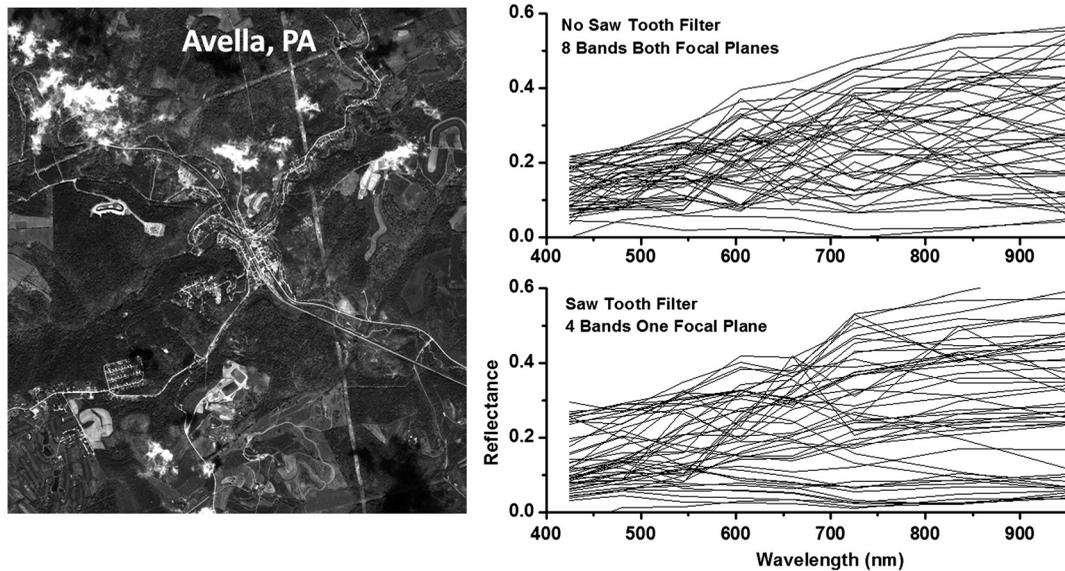


Fig. 13 Application of QUAC to the WorldView-2 MSI VNIR sensor showing the use of the MDMI filters to greatly reduce the influence of undesired sawtooth endmembers. The bottom right plot shows that without filtering, a large number of sawtooth endmembers are selected. The top right plot shows that the sawtooth endmembers have, for the most part, been removed and those that remain have significantly reduced spectral variations.

nature of an endmember algorithm that it will preferentially seek out anomalous spectra. By appropriately redefining the NDMI bands and threshold, we can filter out the sawtooth pattern. For WorldView-2, the proper selection for the filter pair is (478 nm, 546 nm) and (659 nm, 831 nm) with a threshold of 0.2. The result is shown in the upper-right plot of Fig. 13. The sawtooth endmembers have been greatly reduced in number, but not completely eliminated. While not fully optimized, these filter settings provide useful atmospheric corrections for WorldView-2 data cubes, as can be seen in Fig. 14 via comparison to a FLAASH correction. There are also issues with the FLAASH correction for WorldView-2, since the standard aerosol correction method cannot be applied (i.e., no SWIR band is available). Images containing well-characterized materials would be very helpful to establish the accuracies of the QUAC and FLAASH corrections for WorldView-2.

4 Discussion

Here, we assess the current accuracy of QUAC and the sensitivity of its correction to several algorithm and sensor parameters. Figure 15 shows comparison of QUAC

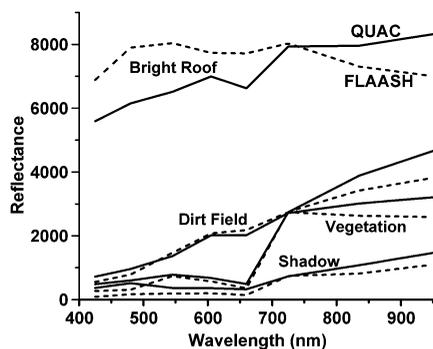


Fig. 14 Comparison of QUAC and FLAASH results for the WorldView-2 data in Fig. 13.

reflectances to those obtained for a variety of HSI and multispectral imaging (MSI) sensors, in which ground truth reflectance spectra were available. For the HYDICE and COMPASS comparisons, field-measured panel spectra were used. For the AVIRIS and HYMAP comparisons, FLAASH corrections, which agreed well with ground truth data, were used. For the LANDSAT7 comparison, a FLAASH correction was used, which was consistent with the FLAASH results for the nearly coincident HYMAP collect over Davis, CA. For each comparison, a ratio of the QUAC spectral reflectance to the “truth” spectrum was obtained, based on a single image pixel with an approximately flat spectral reflectance in the $\sim 0.2 - 0.4$ reflectance range. The ratio is not sensitive to the selection of image pixel used for the comparison.

Excluding the visible region below ~ 550 nm, these comparisons indicate that the overall normalization of the QUAC reflectances are within the range of $\sim \pm 10\%$. However, there can sometimes be a significant tilt to the spectrum, as exhibited by the HYDICE comparison, which shows a $\sim 30\%$ rise from the SWIR to the VIS spectral region. In contrast, the local spectral structure is well preserved, with relative variations in the range of $\sim \pm 3\%$. While we have made

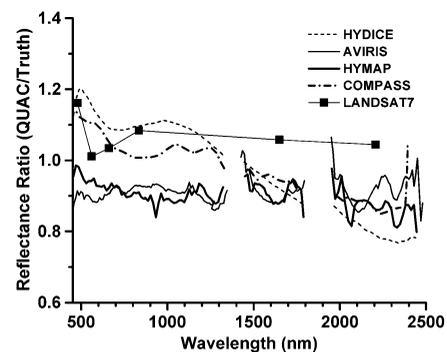


Fig. 15 Comparison of QUAC to well-ground-truthed, corrected data.

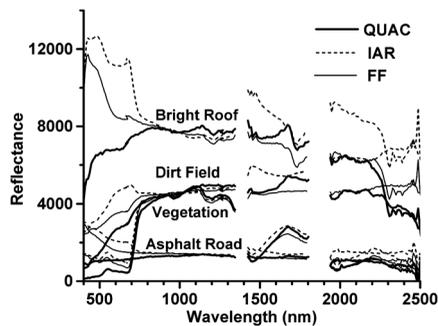


Fig. 16 Comparison of different in-scene-based atmospheric correction approaches for the AVIRIS Harrisburg airport imagery shown in Fig. 9.

some improvements in the VIS region, particularly the tendency to produce large reflectance spikes, further refinement is needed to match the performance in the NIR–SWIR spectral region.

We place QUAC in historical perspective by comparing to previously developed in-scene atmospheric correction approaches, as displayed in Fig. 16. These include the internal average reflectance (IAR)¹⁴ and flat field (FF)¹⁵ approaches, typically used for quick-look, data-survey purposes. In the IAR approach, each pixel spectrum is divided by the average spectrum for the entire image. This yields a relative reflectance spectrum; for comparison to QUAC, we normalized the IAR and FF spectra to the QUAC result for the Bright Roof at 1000 nm. The accuracy of the IAR approach can vary wildly, since it is tied to the most abundant material in a scene. If the dominant scene material is spectrally flat, then IAR will yield good relative spectral reflectances. In the example shown, the scene contains a significant fraction of very nonspectrally flat vegetation, which results in the large spikes in the VIS region. The IAR and FF methods do not include a baseline subtraction and thus will magnify the errors for the darker pixels. In the FF approach, each pixel spectrum is divided by a user-selected pixel spectrum. If the user can identify a spectrally flat material, then good results can be obtained with the FF method. In the example shown, we selected a Tarmac pixel from the Harrisburg airport image. As can be seen, this selection produced useful results in the VNIR–SWIR region, but poor results in the VIS region.

The default number of endmembers used in QUAC for the correction is set at 50. However, as is evident from Fig. 17,

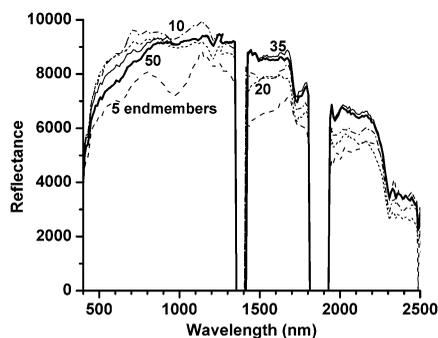


Fig. 17 Sensitivity of the corrected reflectance to the number of endmembers selected for the correction. The example spectra are for a bright roof in the AVIRIS Harrisburg airport imagery shown in Fig. 9.

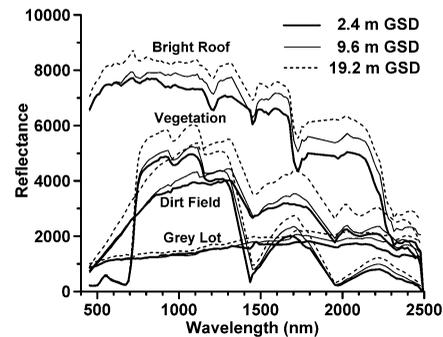


Fig. 18 Sensitivity of the corrected reflectances to sensor GSD, based on spatially resampling the HYMAP imagery shown in Fig. 2.

the results are not overly sensitive to this choice. As few as 10 endmembers (i.e., diverse image spectra) can yield a useful correction.

Finally, we consider the effect of sensor GSD (ground sampling distance) on the QUAC correction, as highlighted in Fig. 18. The dominant trend is for the absolute value of the reflectances to increase with increasing GSD. There is also a noticeable, but modest, change in the overall spectral shape (i.e., a tilt). While not shown in Fig. 18, we also considered a 30-m GSD, which applies to the LANDSAT7 and ALI MSI, space-based sensors. The results for this case follow the same trends, with the absolute values increasing by $\sim 25\%$ relative to the 2.4-m GSD spectra. This is consistent with the comparisons in Fig. 15, in which the LANDSAT7 results are about 25% higher than those for HYMAP for similar imagery. Similarly high reflectances were found when comparing QUAC and FLAASH results for the ALI MSI sensor. Although further evaluation of this GSD dependence is required, it appears that a simple, GSD-based normalization correction will lead to a significant improvement in the QUAC normalization approach.

We attribute the GSD dependence to a change in the amplitude of the average strong vegetation spectrum used in the normalization procedure. As the GSD increases, a larger proportion of vegetation shadows, cast by trees, enters into the average vegetation spectrum. This lowers the peak reflectance of the average vegetation, which, in turn, leads to an increase in the vegetation-based normalization factor.

5 Summary and Future Plans

We have presented a detailed description of the QUAC atmospheric correction algorithm, highlighting recent accuracy and computational speed improvements. QUAC can be applied to any multi- or hyperspectral sensor covering all or part of the VNIR–SWIR (~ 400 – 2500 nm) spectral range. However, the spectral filtering parameters may require some tuning, depending on the sensor spectral coverage. For example, a sensor based on only an InGaAs focal plane covers the ~ 900 -to- 1700 -nm spectral region. For this class of sensors, the standard NDVI-based vegetation filter, which uses bands at ~ 700 and 800 nm, will not work. However, an alternate pair of bands at ~ 1300 and 1500 nm provides good vegetation filtering for InGaAs sensors.

Our future plans for QUAC are driven by the rapidly increasing rate at which spectral sensors can generate enormous quantities of data. This is creating a need for faster

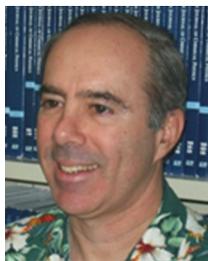
atmospheric correction performed on board the sensor platform. QUAC is well suited for a parallel computing environment because it breaks a data cube into many smaller chunks, each of which can be independently processed. This should enable the development of a real-time, on-board version of QUAC. Another area of interest concerns atmospheric correction for the frequently encountered case of nonideal viewing conditions. This primarily includes the effects of clouds and cloud shadows. QUAC performs best when the imagery is uniformly illuminated, which includes both clear sky and flying under complete cloud cover. For partly cloudy imagery, one may first parse the image into solar illuminated and shadowed regions, and then apply QUAC separately to each region. A key challenge is to perform this parsing in a fully automated manner.

Acknowledgments

We appreciate the technical and administrative contributions of many individuals, including T. Perkins, M. Matthew, M. Fox, and R. Sundberg of Spectral Sciences, Inc.; P. Villeneuve, A. Stocker, and R. O'Connor of Space Computer Corp.; S. Allen of The Analytical Sciences Corp.; J. Magarick, T. Smith, J. Nettles, G. Boer, and A. Sanders of the National Geospatial-Intelligence Agency; B. Thompson, M. Peitersen, and A. Ifarraguerri of Science Applications International Corp. (SAIC); and M. West, J. Grossmann, and M. Halper of MITRE. This work was funded under SAIC subcontract No. P010080877 and Spectral Sciences, Inc. IR&D.

References

- B.-C. Gao, K. B. Heidebrecht, and A. F. H. Goetz, "Derivation of scaled reflectances from AVIRIS data," *Remote Sens. Environ.*, **44**(2-3), 165-178 (1993).
- B.-C. Gao et al., "Atmospheric correction algorithm for hyperspectral remote sensing of ocean color from space," *Appl. Opt.* **39**(6), 887-896 (2000).
- M. J. Montes, B.-C. Gao, and C. O. Davis, "A new algorithm for atmospheric correction of hyperspectral remote sensing data," *Proc. SPIE* **4383**, 23-30 (2001).
- R. O. Green, D. A. Roberts, and J. E. Conel, "Characterization and compensation of the atmosphere for inversion of AVIRIS calibrated radiance to apparent surface reflectance," in *Summaries of the Sixth Annual JPL Earth Science Workshop*, JPL Publication 96-4, Vol. 1, pp. 135-146, Jet Propulsion Laboratory, Pasadena, CA (1996).
- C. J. Miller, "Performance assessment of ACORN atmospheric correction algorithm," *Proc. SPIE* **4725**, 438-449 (2002).
- S. M. Adler-Golden et al., "Atmospheric correction for short-wave spectral imagery based on MODTRAN4," *Summaries of the Eighth Annual JPL Earth Science Workshop*, JPL Publication 99-17, pp. 12-23, Jet Propulsion Laboratory, Pasadena, CA (1999).
- M. W. Matthew et al., "Status of atmospheric correction using a MODTRAN4-based algorithm," *Proc. SPIE* **4049**, 199-207 (2000).
- M. W. Matthew et al., "Atmospheric correction of spectral imagery: evaluation of the FLAASH algorithm with AVIRIS data," *Proc. SPIE* **5093**, 474-482 (2003).
- Z. Qu, A. F. H. Goetz, and B. Kindel, "High-accuracy atmospheric correction for hyperspectral data (HATCH)," in *Proc. of the Ninth AVIRIS Earth Sciences and Applications Workshop*, JPL Publication 00-18, pp. 373-380, Jet Propulsion Laboratory, Pasadena, CA (2000).
- R. Richter and D. Schlaepfer, "Geo-atmospheric processing of airborne imaging spectrometry data part 2: atmospheric/topographic correction," *Int. J. Remote Sensing* **23**(13), 2631-2649 (2002).
- T. Perkins et al., "Speed and accuracy improvements in FLAASH atmospheric correction of hyperspectral imagery," *Opt. Eng.* **51**(11), 111707 (2012).
- D. A. Roberts, Y. Yamaguchi, and R. J. P. Lyon, "Calibration of airborne imaging spectrometer data to percent reflectance using field spectral measurements," in *19th International Symp. on Remote Sensing of Environment*, Environmental Research Institute of Michigan, Ann Arbor, MI (1985).
- F. A. Kruse, K. S. Kierein-Young, and J. W. Boardman, "Mineral mapping at Cuprite, Nevada with a 63 channel imaging spectrometer," *Photogramm. Eng. Remote Sensing* **56**(1), 83-92 (1990).
- F. A. Kruse, "Use of airborne imaging spectrometer data to map minerals associated with hydrothermally altered rocks in the northern Grapevine Mountains, Nevada and California," *Remote Sens. Environ.* **24**(1), 31-51 (1988).
- D. A. Roberts, Y. Yamaguchi, and R. Lyon, "Comparison of various techniques for calibration of AIS data," in *Proc. of the 2nd Airborne Imaging Spectrometer Data Analysis Workshop*, JPL Publication 86-35, pp. 21-30, Jet Propulsion Laboratory, Pasadena, CA (1986).
- L. S. Bernstein et al., "A new method for atmospheric correction and aerosol optical property retrieval for Vis-SWIR multi- and hyperspectral imaging sensors: QUAC (QUick Atmospheric Correction)," in *Proc. of the 13th JPL Airborne Earth Science Workshop*, Jet Propulsion Laboratory, Pasadena, CA (2004).
- L. S. Bernstein et al., "Method for performing automated in-scene based atmospheric compensation for multi- and hyperspectral imaging sensors in the solar reflective spectral region," U. S. Patent No. 6909,815 B2 (2005).
- L. S. Bernstein et al., "Methods for determining a measure of atmospheric aerosol optical properties using a multi- or hyperspectral, multi-pixel image," U. S. Patent No. 7,046,859 (2006).
- L. S. Bernstein et al., "In-scene-based atmospheric correction of uncalibrated VISible-SWIR (VIS-SWIR) hyper- and multispectral imagery," *Proc. SPIE* **7107**, 710706 (2008).
- L. S. Bernstein et al., "Validation of the QUick atmospheric correction (QUAC) algorithm for VNIR-SWIR multi- and hyperspectral imagery," *Proc. SPIE* **5806**, 668-678 (2005).
- P. A. Rochford et al., "Validation and refinement of hyperspectral/multispectral atmospheric compensation using shadowband radiometers," *IEEE Transact. Geosci. Remote Sens.* **43**(12), 2898-2907 (2005).
- J. H. Gruninger, M. J. Fox, and R. L. Sundberg, "Hyperspectral mixture analysis using constrained projections onto material subspaces," in *Proc. of the International Symp. on Spectral Sensing Research*, Québec, Canada, pp. 162-170 (2001).
- J. Gruninger, J. Lee, and R. L. Sundberg, "The application of convex cone analysis to hyperspectral and multispectral scenes," *Proc. SPIE* **4888**, 188-198 (2003).
- A. M. Baldridge et al., "The ASTER Spectral Library Version 2.0," *Remote Sens. Environ.* **113**(4), 711-715 (2009).
- R. O. Green et al., "Imaging spectroscopy and the airborne visible/infrared imaging spectrometer (AVIRIS)," *Rem. Sens. Environ.* **65**(3), 227-248 (1998).
- T. Updike and C. Comp, "Radiometric Use of Worldview-2 Imagery," DigitalGlobe, Longmont, CO (2010).



Lawrence S. Bernstein received his BS degree in chemistry from the University of Delaware in 1971, and a PhD in physical chemistry from the University of California at Berkeley in 1974. This was followed by postdoctoral fellowships at Bell Laboratories and Harvard University involving the spectroscopy of Van der Waals molecules. From 1976 to 1981, he was employed as a research scientist at Aerodyne Research, Inc. In 1981, he co-founded Spectral Sciences, Inc., where he currently serves as the chief scientist. His current research activities include the spectroscopy and radiation transfer of high temperature flows and planetary atmospheres, and the simulation and atmospheric correction of hyperspectral imagery.



Xuemin Jin received his BS degree in physics from Shanxi University, Taiyuan, China, in 1982, the MS degree in physics from China Institute of Atomic Energy, Beijing, China, in 1985, and the PhD degree in physics from the University of Maryland, College Park, MD, in 1993. From 1994 to 1996, he was a postdoctoral research associate at Canada's national laboratory for particle and nuclear physics (TRIUMF), Vancouver, Canada, working on analytical models for strongly interacting multi-nucleon systems. From 1996 to 1999, he was a senior postdoctoral research associate at the Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, MA, working on fundamental spin structures of nucleon. He joined Spectral Sciences, Inc., in 2002. His current research interests include active combustion control systems based on optical sensing using emission tomography, tomographic spatial profiling of turbine engines, and tomographic-based, laser radar reconstruction of remotely located objects.



Brian Gregor received BS degrees in physics and philosophy from Northeastern University in 1998 and MS and PhD degrees from Boston University in 2002 and 2004, respectively. Since joining Spectral Sciences, Inc. (SSI) in 2004 he has been involved in the development of simulation codes for high-altitude rocket plumes and hyperspectral image processing. In the SSI laboratory, he has been active in the development of novel optical instruments applied to a

number of problems including turbine combustion, rocket plume characterization, detection of surface contaminants, and detection of blood coagulopathy.



Steven M. Adler-Golden joined Spectral Sciences, Inc. (SSI) in 1981 and currently leads their Passive Sensing Group. His experience at both technical and management levels includes atmospheric aeronomy, infrared/visible/ultraviolet radiation modeling, and prototype trace gas sensor development. He is currently active in development of atmospheric correction and data analysis algorithms for spectral imagery from visible to thermal infrared wavelengths. He received

his BS degree in chemistry from Yale University in 1974 and his PhD in physical chemistry from Cornell University in 1979. His PhD research under professor John Wiesenfeld and postdoctoral research under professor Jeffrey Steinfeld at MIT was in the area of small-molecule photochemistry and spectroscopy.