

Interactive comment on "The Atmospheric Infrared Sounder Version 6 cloud products" by B. H. Kahn et al.

B. H. Kahn et al.

brian.h.kahn@jpl.nasa.gov

Received and published: 9 September 2013

The authors appreciate the helpful feedback from the two anonymous reviewers. This paper is long and detailed and most certainly took an inordinate amount of time to review. We appreciate the time investment, and we feel as though the comments were very helpful for improving the quality of the manuscript. All reviewer comments start with 'comment:', while the responses begin with 'response:', and are in regular font.

Anonymous Referee #1

comment: The manuscript describes a new version of AIRS/AMSU cloud products obtained with the help of improved retrieval algorithm. Compared to a previous version, cloud temperatures are determined on a 3 times finer spatial grid (13.5 km vs 45 km), C6659

accounting for variability within complex cloud scenes. Version 6 data set (v6) also contains new products: cloud thermodynamic phase, ice cloud optical thickness, and ice cloud effective diameter, and ice cloud top temperature. This, combined with a broad spatial coverage of the AIRS sounder makes v6 a high demand product for the community.

Overall, the paper is well organized, well written, and is detailed enough to understand the approach and the results. I had difficulties only with two parts, which I address in "General comments" section. I believe that these parts are essential for the quality of the paper and they both require additional calculations, so I have chosen "major revision" from the list of options. I am looking forward to see an updated version of the manuscript.

Response: The authors appreciate the kind words on our research efforts and on the paper itself. We will address the two suggestions by reviewer #1 below.

comment: 1) The manuscript describes improvements to existing methodology of the cloud parameter retrieval from the AIRS observations, in comparison to an earlier NASA retrieval. However, the most interesting and most difficult methodological part is underrepresented (the case studies described in page 7 and 8 are already dealing with measured radiance). Section 2.1 "What is new in Version 6" qualitatively describes the new algorithm, but it does not provide a self-consistent study, which would be more convincing than 75 lines of text. Without this study, it is difficult to estimate the quality of the retrieval algorithm itself, especially since the authors note in lines 26-27 of page 6 that "there is the potential to create a non-existent cloud layer that in practice only fits noise". If the retrieval algorithm can take noise signal for a real cloud instead of filtering it out based on a large r.m.s. of the deviation of radiances, this raises certain questions regarding the approach. I would suggest supplementing Section 2.1 with the following exercise: two cloud layers are considered, full pixel coverage is assumed, four test situations are modeled – two layers, lower cloud, upper cloud, no clouds. Forward radiance is calculated in the corresponding channels; realistic noise is added to each

channel; standard retrieval procedure is performed. Radiance deviations (one can sort them in an ascending mode for the presentation purposes) should be shown for each case, and the solution search described in lines 29-31 of page 6 should be presented graphically. This exercise can be repeated for two or three values of cloud optical depth and shown in one plot.

Response: When we stated that 'there is the potential to create a non-existent cloud layer that in practice only fits noise', we were referring to very small values of effective cloud fraction (ECF). These values are typically much less than 0.01 (and many of those clouds may in fact be real).

The major challenge of this study is characterizing geophysical 'noise' whose source and magnitude are not fully understood. Our conjecture is that this 'noise' in the ECF retrievals arises from the assumption of horizontally uniform surface temperature and emissivity fields, and constant temperature and water vapor profiles in the AIRS/AMSU FOR (the 3x3 AIRS FOVs) that are used for cloud clearing. Any non-cloud horizontal variability may be misidentified as ECF.

Non-uniformity is a bigger worry than instrument noise. In the revised paper, we are including each channel's noise level in K. Our use of multiple channels beats the effective noise level well below the 0.1 K level. Therefore, a 'phantom' cloud on the level of this noise is vanishingly small. The bottom line is that the cloud retrieval accuracy is not 'noise limited'.

In solving for two cloud pressures uniformly over the AIRS/AMSU FOR (older version 5), many more of these very small ECF values are produced. When using the version 6 algorithm described in this paper, these small ECF values are retrieved more rarely. This does not fully characterize the contributions of sub-FOV variations in surface temperature, emissivity, or temperature and water vapor profiles, but it does demonstrate that when the small-scale cloud configurations are taken into account, the solutions of ECF are improved and fewer erroneous cloud retrievals are obtained.

C6661

While helpful in showing sensitivity of our method to cloud-cleared radiances, the experiment proposed by the reviewer would not address a more fundamental concern: characterizing the 'geophysical noise' produced by the cloud clearing process. If we understand the reviewer's comment correctly, the suggestion is to try and fit a single layer or double layer cloud for different values of optical thickness with the channels we are using for the cloud parameter retrieval with the addition of some amount of "realistic" noise. This does not address the fundamental error source in the cloud clearing process.

Also, the AIRS-LMD retrieval does not attempt a two-layered cloud solution, and it uses the AIRS L2 temperature and water vapor profiles as-is in conjunction with the TIGR database to account for atmospheric thermodynamic structure. So, the AIRS-LMD approach suffers from some of the same limitations.

Characterizing the sub-FOV variability and its impacts on geophysical retrievals such as temperature and water vapor is an ongoing fundamental problem with AIRS retrievals. That is why a multi-sensor retrieval is appealing if one were interested in effects of very small cloud/temperature/water vapor variations on, for instance, minor/trace gas retrievals (CO2, CO, O3, N2O, CH4, SO2, HDO, etc.). We argue that these effects are less of an issue with much larger geophysical signals.

There are plenty of other 'indirect' lines of evidence regarding 'self-consistency'. The radiative consistency (Nasiri et al., 2011) between AIRS and MODIS cloud products is identical between v5 and v6, as described in the paper. Therefore, the improved cloud top height agreement between AIRS, CloudSat and CALIOP implies improved ECF values. On a more basic level, the 'closure' between observed and calculated Tb in the AIRS cloud-clearing algorithm is reported in the variable 'CBTmOBT1231'. This is shown in Fig. 1 for January. These are the residual BTs from computed BT minus L1B observed BT using the full AIRS retrieved state including clouds. They show that we generally get a good closure but also highlight where the problems are (over some land areas, and in the cloudier regions of the tropics and mid/high latitudes). However,

the mean residual is only a few tenths of a K, and most regions are within 1 K or less.

Also, some of the process-based examinations in the context of the spatial and temporal variability of cloud thermodynamic phase, the diurnal variability of convective ice clouds, and mid-latitude cyclone composites provide further evidence of this 'selfconsistency'.

comment: 2) I believe that the manuscript would benefit from extending the demonstrated dataset for at least one month in boreal summer (or better yet for a whole year). Otherwise, it creates a feeling that the retrieval is so slow and complicated that only one month of data could be produced during the time of the manuscript preparation. Having four seasons in two hemispheres will give a representative picture of the updated dataset and will create a better image of the work as a whole.

Response: This is a good suggestion and we have processed additional data to make this change possible. We have processed one years' worth of AIRS data (2007) as opposed to only January 2007. This has led to changes in several figures and some of the text. The figures that now show all of 2007 include Figs. 2, 5, 6, 8, 9, 10, 12, and 14 (formerly Fig. 13). Figs. 2, 5, 6, and 8 show global maps of various cloud properties. With the inclusion of a full year's data set, the cloud features are now much smoother with less hemispheric asymmetry. We have adjusted and supplemented the body of the text accordingly. Figures 15-17 (formerly Figs. 14-16) are now used to demonstrate the interesting hemispheric differences in the context of mid-latitude cyclone composites. We decided against including the full 2007 data set in these figures since authors Kahn and Naud are pursuing similar work independently, and the extent and detail of this work is well beyond the scope of the present article. The new Fig. 13 is very similar to Fig. 12 and shows the seasonal zonal variations in ice cloud top temperature, effective radius and optical thickness.

comment: The authors refer to other cloud datasets, as well as to an internationally coordinated assessment activity. The latter has led to a database which is available

C6663

since the beginning of this year. To show how the results of the presented dataset compare with the other existing datasets, especially with the one derived from the same AIRS observations (on a statistical basis), it would be important to include a short comparison, for example in the form of a table.

Response: This is another good suggestion and we took the reviewer's advice to explore the GEWEX data set in the limited time we have for a paper response and revision. We compared monthly global maps of the ice cloud optical thickness, effective diameter/radius, and cloud top temperature between AIRS-NASA (the current retrieval) AIRS-LMD, MODIS-ST, MODIS-CE, ATSR, PATMOSx, and POLDER. See Figs 2, 3, and 4 at the end of the response for example maps during January 2007. (The other 11 months of 2007 are not shown.) Frankly, we were quite shocked at how poor the agreement was between the various data sets. In the case of AIRS-NASA and AIRS-LMD, which should be the most consistent, we found that AIRS-LMD had a tendency for a significantly lower optical thickness than AIRS-NASA. With regard to the cloud top temperature, AIRS-LMD was quite a bit colder than AIRS-NASA, especially in the Tropics. For reasons described in the paper, the rather warm cloud top heights are seen in some subtropical areas in the AIRS-NASA data set. However, the most troubling agreement was seen for the effective diameter comparisons. The patterns look somewhat coherent for AIRS-NASA and are described at length in the current paper draft. The AIRS-LMD maps look rather noisy with very little resemblance with AIRS-NASA, except in the coarse mean sense (if all of the pixels globally were averaged to one value and the two numbers were compared for the two data sets). This is evidenced in Fig. 5, which shows zonal mean averages for 2007 between the two AIRS data sets. The lack of agreement is quite surprising.

We think that this sort of inter-comparison effort deserves a lot more attention, but it is impossible to add anything meaningful to the paper in the form of a short table that would be useful to either data set. We have decided against adding any figures, tables, or this particular analysis to the final draft of the ACP paper. It deserves its own separate analysis.

A more revealing approach that the lead author would like to take is to make pixelscale comparisons between the two AIRS retrievals matched in time and space. This has not been done to date, and unfortunately this level of detailed work is well outside of the scope of this paper. However, the lead author has downloaded daily pixel-scale files from the LMD FTP server for one 16-day orbit cycle and will pursue this work in the near future in conjunction with similar efforts that are being made to compare MODIS-ST and CALIOP/IIR products with AIRS-NASA at the pixel scale.

comment: The assessment activity has shown that distributions of cloud properties like tau and De depend strongly on retrieval filtering: the authors introduce a quality flag for the retrieval of these quantities and show the statistics of retrieved data for different quality flags (Table 4). One has to note here that using only good or best quality retrievals introduces a bias in the statistics towards, for example, optically thicker or optically thinner clouds (please, see the specific comments to page 46).

Response: The authors completely agree with the reviewer. In the paper, we detailed very carefully the impacts on the distributions of ice cloud top height, effective diameter and optical thickness in Fig. 10 in the paper. For optical thickness, one can see that the "best" retrievals have distributions skewed the largest, and "good" is a little less, with "bad" the smallest. A similar story is found for the other two variables. The authors would also like to point out that choosing which particular AIRS FOV to retrieve ice cloud properties on and which to ignore can also lead to large differences in the sample set. We suspect that this is a big contributor to the differences found between AIRS-LMD and AIRS-NASA. If we understand correctly, the approach used by LMD uses a CTT colder than some nominal value to choose which FOVs to retrieve cloud properties, and which to ignore. In the present paper, the authors argue that using a physically based cloud thermodynamic phase algorithm is a large step forward for a more consistent retrieval with AIRS-NASA. The ice cloud detection approach was well validated in the study of Jin et al. (2013). This topic will be pursued in future work by

C6665

the lead author with regard to the FOV-scale comparisons between AIRS-NASA and AIRS-LMD.

comment: Page 1, line 1 - I don't know an internal policy of naming the datasets in this case. Perhaps, the authors have to specify that this is a "Science Team Version 6" to distinguish it from other retrievals?

Response: Where necessary, we will make clear that these products are generated by NASA/JPL.

comment: Page 2, line 18 – IPCC AR4 is not defined. I would suggest putting all the abbreviations to the Annex for the sake of readability. In this case, one can skip some definitions in the text and just add a note in the beginning.

Response: done

comment: Page 3, line 18 – perhaps, one has to complete the sentence: "climate sensitivity to radiative forcing" since this is the first time when climate sensitivity is mentioned.

Response: done

comment: Page 4, line 27 – The AIRS cloud retrieval algorithm described in [Kahn, et al., ACP, 2008] required far less channels. Please, explain.

Response: In Kahn et al. (2008), the retrieval approach was based on a look-up table that relied on searching for the minimum chi-squared value in the look-up table space. In the present work, we discuss the channel selection process in Section 3.2.2 and that reduced retrieval convergence rates were obtained with successively smaller channel data sets. (The same channel set in Kahn et al. (2008) was tested, but performed poorly. That channel data set did not include CO2-slicing channels, which we suspect is a source of error.) In short, in the current retrieval implementation, we were able to obtain retrieval convergence almost 100% of the time when we extended the channel set to several dozen (See Section 3.2.2 and Fig. 7). Similar convergence

rates and nearly identical retrieval values were obtained for some tests with about \sim 200 channels, but these retrievals were much more computationally expensive. We do not know the full set of reasons why retrieval convergence rates are reduced for smaller channel data sets. We speculate this may be tied to the information content of the channels chosen as well as prior guesses and constraints. The lead author is hoping to quantitatively investigate these issues further if future funding is available.

comment: Page 5, line 2 – how does fixing the surface temperature affect the retrieval of low clouds? What is meand under "fixing" atmospheric parameter here? Please, explain.

Response: This is in reference to the ice cloud retrieval algorithm tacked on to the AIRS Standard algorithm which performs the cloud clearing and provides the one- or two-layered cloud top temperature/pressure and effective cloud fraction values discussed in Section 2. We are not retrieving low liquid cloud optical thickness or effective radius, only ice cloud properties identified as likely containing ice by the cloud thermodynamic phase algorithm. 'Fixing' is just another way of saying that the parameter is "held constant". (Will change to 'held constant')

comment: Page 6, lines 24-31 – please, see the general comments.

Response: Please refer to our response regarding the first general comment.

comment: Page 7, lines 13-15 – "more accurate determination" should be supplemented with the actual numbers.

Response: The improvements in spectral emissivity can be quantitatively assessed against the Masuda model over the ocean, and v6 was shown to be superior to v5. These results are described in full in Susskind and Blaisdell (2008) as currently cited and written. Furthermore, the land surface emissivity is more poorly known in terms of its spectral behavior, so an absolute standard is harder to come by. However, validation efforts over land surfaces with strong spectral features that have been measured with

C6667

in situ field campaign data (e.g., desert surfaces in the southwest U.S. and Namibia) have shown that more realistic surface spectral emissivity has been obtained in v6. Hulley et al. (2009) describe the methodology and this work was repeated for v6.

comment: Page 8, line 3 – the "summary" is "summarized". Please, re-phrase.

Response: fixed

comment: Page 8, line 4 – are the values in Table 1 area weighted? If not, this can bias the statistics. Please, specify.

Response: That is correct. We weight by latitude. We clarified in the Table caption.

comment: Page 8, lines 6 and 19 - what is meant under the term "cloud signal" here?

Response: For sake of clarity, this was changed to 'cloud amount' in both instances.

comment: Page 9, lines 2-16 – this is only an indirect proof of the improved methodology. No change is required here, I just draw the authors' attention to a lack of direct demonstrations of the quality of the retrieval approach using single retrievals.

Response: We understand the reviewer's point. We did not include some figures we produced that are similar to Kahn et al. (2008) that show 2-d pdfs of height comparisons broken up by cloud type. Some of the cloud types saw a reduction in a bias when compared to CloudSat and CALIOP in version 6. We are investigating this topic further as a larger data set is produced.

comment: Page 10, lines 16, 21 and below (i.e. page 14, line 6) – the names of the fields alone can be used in the description of a data product. However, if they enter the formulas (like formulas 3-5), it creates a strange mixture of scientific and programming styles. Please, consider some simplification/unification of the notation.

Response: On pp. 10 and 14, it is appropriate to state the actual field name since there are multiple fields/variables that can be confused, and it is certainly helpful to be as concise as one can be (from extensive previous experience). Furthermore, some

of these field names, like the initial prior guesses or scalar averaging kernels, may not have accepted symbolism (e.g., 'tau' for optical thickness). Table 3 was intended to be the descriptor of all variable names, as they appear commonly and as they appear in the AIRS L2 product files, and each is associated with a concise description.

comment: Page 11, line 29 – "60% of all liquid clouds are identified by AIRS as unknown" – this is a large value. Is it related to a previous version? How is this issue addressed in a new version? Please, specify the name of the algorithm in the text.

Response: We have added a reference (Jin and Nasiri, 2013, sub judice) that is from the doctoral work/dissertation of H. Jin (2012). The cloud thermodynamic phase algorithm they developed was intended to be a sensitive ice detection algorithm. This was developed for Version 6 and did not exist in any version previous to it. It has greatly succeeded in that case (see Jin and Nasiri, 2013). As may be obvious to some (and not to others), low warm clouds have a much weaker infrared signature than cold high ice clouds because of a greatly reduced thermal contrast with the ground. Despite this fact, most uniform stratocumulus are detectable as 'liquid' (exploiting the index of refraction differences between liquid and ice in the 8-12 micron window region), whereas almost all trade cumulus (the dominant cloud type in the low latitude ocean) or open cell cumulus appear as 'unknown' because they exert and even smaller radiance signature than stratocumulus because of their small spatial coverage. This is discussed at length in Jin (2012) and Jin and Nasiri (2013), and also in Section 4.2. Keep in mind that these cloud phase choices are obtained from radiative signatures, not assumptions about cloud top temperature, height, etc. If the latter were included to supplement the former, almost all of these 'unknown' clouds would be called 'liquid'.

comment: Page 13, line 9 – "are retrieved in log space to prevent negative values". I did not understand this explanation since the log space and a "normal" one are in one-to-one correspondence. Is it related to some kind of extrapolation? Please, explain.

Response: In some cases where an atmospheric parameter can vary widely, it is pos-

C6669

sible that the solver in linear space can test a parameter in an iteration with a negative number, which would cause errors or crash the forward model. (Imagine if you put in a negative optical depth into the forward model, for instance.) However, if the solver is working with a parameter in logarithmic space, it's assured that the parameter will be > 0 when one takes the exponential to use in the forward model. By retrieving logarithmically, things can be set quite small (or effectively zero), but never truly zero or negative.

comment: Page 16, lines 17 and 21 – these sentences contradict with each other. If the computational expense is high then it would be logical to optimize the selection of the channels.

Response: On line 17, the lack of 'optimization' refers to the channel selection for information content, not to the computational expense. Please also see our response to the comment about p. 4, line 27.

comment: Page 17, line 8 – please, see the general comments. I would suggest showing more examples to balance the methodological part of the manuscript.

Response: We now have a full year of 2007 results in the paper, and have included the seasonal variations of the three ice cloud property parameters sorted by land, ocean, day, and night.

comment: Page 46, Table 4 – why does De retrieval never produce "Best" quality retrievals? Please, supplement this table with averaged tau values for each column.

Response: As discussed in Section 3.2.2, we chose not to assign a 'best' value to De because it is a very difficult parameter to retrieve, and correlative data sets that are considered 'truth' are difficult to come by. With regard to average values for each parameters and each QC value, we think that the full histograms are much more interesting and informative. These are found in Fig. 10.

comment: Page 47, line 2 – color bar titles are barely readable (the same is true for

Fig. 5,6, and 8)

Response: fixed

comment: Page 52, Fig. 6 and lines 4-7 – the "phases" here are not "natural" phases like "liquid/ice", so it can mislead the reader. The numbers should be explained in the legend or, perhaps, one can also introduce some names in the text. Response: In Section 3.1, the four ice tests and two liquid tests were described in great detail, and the caption of Fig. 6 clearly says that these values indicate the sum of the different tests passed (liquid is negative, ice is positive). We added 'Please refer to Sect. 3.1 for an explanation of each phase test' in the figure caption for clarity. Page 55, Fig. 9 – is the binning always linear or the "Tau" histogram has a log scale binning?

Response: the binning is always linear (bin widths are the same from tau=0.1 to 10.0), while the plots are simply plotted in log scale. Presenting the results on a log scale made the most sense, while using linear binning made sense because we did not want to weight different parts of the histograms more than others.

comment: Page 56, Fig. 10 - same as above.

Response: see above

comment: Page 59, Fig. 13 – Please, add the latitude and longitude values to the axis since the contours are not always visible.

Response: done

Reviewer #2 comment: The authors of this paper faced the ambitious task of introducing a new release of the cloud property retrieval algorithm for the AIRS instrument. The AIRS sounder was the first of its kind, a hyperspectral infrared sounder, in space and has been operational in polar-orbit since 2002. This is the sixth version of the retrieval algorithm.

Response: While it is true that this is the sixth version of the operational AIRS algo-

C6671

rithm, it is the first version containing ice cloud properties and cloud thermodynamic phase. Furthermore, it is also the first version with cloud top pressure retrievals performed at the AIRS FOV. Technically, the first release of AIRS to the public was Version 3, so there have only been 4 versions in 10+ years.

comment: Since 2002, three hyperspectral infrared sounders have been launched in 2006, CrIS on Suomi-NPP in 2011, and IASI on Metop-B in 2013. With four infrared sounders in operational polar-orbit, hyperspectral sounding retrievals offer an unprecedented characterization of the vertical atmospheric column in space and time. Of the 2374 available channels in AIRS, only half are still useful today and this number is continually decreasing due to what is commonly referred to as channel popping.

Response: In order to address this comment, the following reply is via a series of personal communications from Steve Broberg, Margie Weiler, and Denis Elliott, members of the AIRS instrument cal/val and engineering teams. The AIRS infrared instrument has 2378 channels. There are currently 2213 channels with a NEdT < 1.0 K at a scene BT of 250 K, but almost all of these channels have a NEdT < 0.5 K (median=0.197 K). Thus, 165 channels are not usable or have excessive noise (some are caused by 'popping'). This is a smaller number of bad channels than at launch (around 200) because of multiple channel calibration table uploads to the instrument in the last few years.

'Popping' is a specific noise phenomenon that is a characteristic of mostly higher noise channels where the baseline level changes suddenly. Popping (or telegraph noise) usually means abrupt shifts in the output level of any electronic device with fairly random up and down intervals but often fairly regular amplitudes. A very small percentage of AIRS channels (probably less than 10) are regular 'poppers' by this definition (a change in level during one scan) and only 70-100 channels per day get a pop flag (flags are generated for each of the 240 granules a day). Since the latest calibration table upload, this is now closer to 70 than 100. These numbers are higher than we saw earlier in the mission, probably because of a gradual accumulation of total radiation dose in the readout electronics (affecting the bias generators for the detectors).

However, the number of useful channels has actually increased since launch.

The AIRS instrument has two detectors per channel (A and B). If both detectors are working, then both A and B are averaged for square root of 2 increased noise performance. There are some channels where one detector may not be working. If someone is looking to use a set of channels that is good for the entire duration of the mission, then that number would be decreasing with time, as there is a period of time where an A detector or a B detector has gone bad, and the channel has not recovered (many do recover) or been restored to its best state by switching to the remaining good detector. This period of time where a channel is not good can be quite long; however, the AIRS Team has recently started restoring channels more frequently. Given this, only a small proportion of the 2213 usable channels have one detector (A or B) functioning. Thus, for most channels, the noise performance has remained rather steady over the length of the mission.

One last comment about data releases. When it comes to comparisons with IASI, the AIRS team actually releases the Level 1B data. It is out understanding that the IASI Team does not – they only release a 'corrected' Level 1C. So, it is difficult to independently ascertain the true instrument performance, like it is easily done with AIRS.

comment: This said however, the strong advantage AIRS has is its decadal data record. This could perhaps be the dominant motivation for a new algorithm release and reprocessing of the full data record.

Response: As mentioned above, these cloud products are new for Version 6, which is the dominant motivation for publishing this particular paper. It is not intended to be an instrument overview or a discussion on the cloud-clearing retrieval system.

comment: The value in generating an accurate and long-term cloud retrieval set is recognized by the authors in the first few paragraphs of the introduction and again in the conclusion. They state that AIRS offers "unique decadal scale and global snap-shots

C6673

of clouds within the diurnal and annual cycles at all latitudes" (line 14, p14510). This means that the (Version 6, reprocessed) AIRS cloud retrieval record has the potential to help characterize "long-term, global-scale cloud trends that has been uncertain and difficult to determine" thus far (line 18, p14480).

An algorithm release is a difficult type of paper to write; even more so when the algorithm is the sixth version in a series of algorithms for an instrument with a life cycle spanning a decade. Over and above communicating the scientific integrity of algorithm improvements, it is the type of paper that needs to reach a wide audience and make a case for one data set (V5) over another (V6). In essence, a paper like this needs to fulfill three requirements; provide (1) an overview and summary of the instrument, its retrieval algorithm, products, and usefulness thus far, (2) technical description of algorithm changes and additions, as well as (3) a thorough examination and justification of improvements by way of comparison with other relevant datasets.

Response: We agree with the reviewer that an overview instrument or cloud-clearing algorithm paper must satisfy a minimal set of requirements, and the suggestions of (1)-(3) listed above are well taken. However, this paper is not intended to be such a paper. It is only focused on the cloud products, most of which are new for version 6 (ice cloud properties and thermodynamic phase), and for the remaining cloud properties that have been in previous versions of AIRS products (cloud top temperature and effective cloud fraction), there have been important changes and improvements over previous versions.

With regard to comparisons to other relevant data sets, this work is currently in progress. It is straightforward to take the GEWEX cloud inter-comparison results and compare to the current AIRS algorithm results (see response to reviewer #1). However, these sorts of mean plots aren't entirely insightful. The lead author is currently making comparisons at the pixel-scale to MODIS and CALIOP/IIR retrievals of ice cloud optical thickness and effective diameter. This work is rather involved and well beyond the scope of the current work, which is simply focused on the new cloud products that

AIRS has to offer.

comment: In my opinion the paper in its current form almost succeeds in addressing this ambitious set of requirements. They give a thorough technical description of new algorithms and updates on old ones. They do a thorough evaluation of the quality flags and discuss the impact quality control can have on the final results. They spare no detail in describing how the quality flags, that are "neither absolute nor quantitative" (line 1, p14496), should be applied to derive at meaningful results. This is often overlooked in algorithm papers, but key in making the data products useful to the community. Then they present and analyze a rich set of results at different spatial and temporal scales; e.g., localized Tc patterns (V5 versus V6) at a single-granule scale over the Pacific Ocean (Fig. 1), global statistics of Tc and ECF (Table 1), global distributions of AIRS V6 Tc and ECF (Fig. 2), a comparison of global ocean and land cloud heights between V5, V6, Caliop and Calipso (Fig. 3 & 4), global patterns of V6 cloud phase (Fig. 5), a global evaluation of the sensitivity of the cloud phase results to algorithm thresholds (Fig. 6), AIRS spectrum (in brightness temperature units) sensitivity to changes in the three ice cloud products, optical depth, effective diameter, and effective temperature. This gives insight into the challenges faced during algorithm design (Fig. 7), global distribution, statistical analysis, effect of quality flags, zonal averages and diurnal patterns of ice cloud properties (Fig. 8-13), and lastly a mid-latitude cyclone evaluation with the different V6 cloud products (Fig. 14-16). I want to commend the authors on their thorough evaluation and presentation. Response: The authors appreciate the kind words offered by the reviewer. However, I do have a couple of major concerns and questions. My overall impression is that the paper lacks a clear goal and seems haphazard in its collection of results presented. It does not form a coherent whole and instead comes across as a collection of parts. More specifically, given what I consider to be the requirements for a paper like this, it lacks a meaningful summary of the AIRS algorithm evolution and accompanied pivotal algorithm papers.

Response: We want to reiterate that this is not an AIRS algorithm paper that describes

C6675

end-to-end the entire retrieval system. It is a description of the cloud products produced by this retrieval system. The retrieval method of the new ice cloud properties (which uses optimal estimation) is completely new and different from the standard AIRS retrieval (which is based on cloud clearing).

comment: In addition, there's a conspicuous dearth of citations on how AIRS data have been used till now. The reader is left with the conclusion that even though there's been a V5 cloud retrieval data set, it has not been used much in analysis.

Response: As stated before, only cloud top temperature/pressure and effective cloud fraction were produced in v5, and the cloud top temperature/pressure was produced over the AIRS/AMSU 3x3 field of regard (FOR). For v6, these products are now produced at the AIRS FOV scale. This has led to some significant improvements in cloud detection and height characterization, as evidenced by the comparisons with CloudSat and CALIPSO. We counted at least 17 references that describe observations, simulations, retrievals, or some application of AIRS-related cloud products. There are many other papers in the literature, but this is not intended to be a review paper. We cited the most relevant work regarding AIRS cloud products. The reviewer suggested some additional papers that are very much relevant at the end of the review, and we have added these to the manuscript (please see reply after last review comment).

comment: Why then do the authors envisage V6 retrieval products to be applied with more enthusiasm?

Response: Because the new cloud thermodynamic phase and ice cloud property products are new and unique. And based on current collaborations between the lead author and members of the scientific community, these new products will be essential for evaluating the new generation of climate models that contain explicit representations of ice microphysics.

comment: Yes, compared to Cloudsat and CALIOP, there's now an improvement in cloud height and frequency estimates. This is encouraging and clearly demonstrated.

But, given the fact that the authors make a case for using the decadal AIRS cloud record in long-term trend studies it remains unclear how the V6 AIRS record adds "unique" information to the suite of existing decadal datasets.

Response: See previous comments about the new ice cloud and thermodynamic phase properties produced by a completely new retrieval approach. We have also used an optimal estimation retrieval approach and have begun to attack the error estimation problem with rigor. Error estimation must be addressed if one is to use these data for trends.

comment: How does a global day of cloud retrievals compare between V5 and V6?

Response: For CTP/CTT and ECF, they are substantially improved. Low clouds have a much more realistic cloud top pressure (the high bias is reduced). Heterogeneous clouds have more structure to them due to the improved spatial resolution of CTP/CTT. The CloudSat/CALIPSO comparisons show that the vertical histograms of cloud height are more realistic in v6 compared to v5. Table 1 also summarizes the global statistics between v5 and v6 for land, ocean, and the entire globe.

comment: Or between V6 and other existing cloud products (notably from imagers such a MODIS)?

Response: We addressed this in a response to reviewer #1s suggestion to compare against the GEWEX cloud inter-comparison data set. The comparisons between cloud top height, ice cloud effective diameter and optical thickness comparisons are quite poor, even between the current AIRS retrieval and the AIRS-LMD retrieval. However, we argue that showing differences in the mean fields is not that insightful. The lead author and colleagues are pursuing cloud regime dependent comparisons that are built from the pixel-scale. These offer more promise for some insight on the differences between different instrument and algorithm approaches. We have performed some limited comparisons for MODIS Collection 6 retrievals (now shown). When the AIRS FOV is covered by 100% ice cloud (according to the MODIS IR cloud phase/cloud

C6677

mask), the comparisons have very little bias for optical thickness, which is remarkable considering the observing wavelengths and algorithmic approaches are vastly different. The reason the maps disagree so much is that outside of these scenes with 100% ice cloud fraction, there is a whole suite of other cloud configurations where the agreement falls apart because the different instruments are sensing different types of clouds, or different aspects about the same clouds.

comment: Apart from improved cloud height, is there a significant improvement in depicting spatial patterns and their diurnal/seasonal variation? This is especially relevant since one of the primary improvements in V6 is that cloud retrievals are now performed at single-FOV scale. Given the objective of depicting cloud trends on a global-scale, do single-FOV retrievals really matter?

Response: We show that the diurnal variability is quite reasonable for ice cloud properties in the tropical Western Pacific Warm Pool region (Section 4.1) and compares well with the TRMM literature of the convective diurnal cycle. We also showed that composites of NH vs. SH mid-latitude cyclones show very revealing and realistic cloud structures similar to previous work done with CloudSat, CALIPSO, and MODIS data (see references by C. M. Naud, et al.).

For depicting trends, which is not the sole purpose of a long data record (accumulating statistics over a wide variety of regimes is also another one – that way much better composites such as that shown in Sect. 4.2 can be made), we don't know if single FOV retrievals are important or not. However, we do note that different data sets such as ISCCP, HIRS, and AVHRR have much different cloud trends, and sensor resolution is only one factor (of many more) that varies between these data sets.

comment: The V6 products suite includes, for the first time, the capability of characterizing cloud thermodynamic phase and ice cloud properties with an infrared sounder. This is a major contribution and promises to serve the user and scientific community well. However, it remains unclear why a user should switch from using imager-retrieved to sounder- retrieved microphysical properties.

Response: We are not arguing that data users should do such a thing. We are simply describing a new data set that may complement existing data sets and perhaps offer some new insights on cloud properties. The fact that the GEWEX data set comparisons compare so poorly in the mean sense is plenty of motivation to continue doing research on this very important remote sensing problem.

comment: The authors make a brief statement on how V6 AIRS phase retrievals compare with those from ISCCP and MODIS, but I find this statement to be a shallow and largely meaningless.

Response: If the reviewer was referring to p. 11, lines 6-9, this statement is simply referring to several cited papers showing cirrus cloud frequencies, and the intent was to make a very blunt statement about relative amounts of ice cloud among several satellite instruments. All we were trying to say was that AIRS and HIRS are quite similar in ice cloud frequency, ISCCP is known to have a lower ice cloud because many of the high clouds are reported as mid-level clouds, CALIOP reports slightly lower ice cloud frequencies than AIRS, mostly because the lidar spot is much smaller than all other observations (and probably represents a more realistic frequency than all others), and also MODIS is lower than AIRS because the operational product does not handle thin cirrus below an optical thickness of 0.3-0.4 (Ackerman et al., 2008, JAOT).

comment: Were the data scaled to a similar spatial resolution? Were any attempts made at neutralizing instrument (AIRS and MODIS) and product (ISCCP is a multi-instrument product) differences before comparison?

Response: Please see previous response. These comments were based on published climatology maps that were cited. We are not trying to force different instruments to be the same as each other. We are merely citing the refereed literature and making a broad-brush statement about it. There is nothing new or controversial about this at all.

C6679

comment: In a time when multi-instrument analysis of the environment and climate becomes not only possible but also essential, it is more important than ever to be clear and specific so that results can be reproduced, tested and scrutinized. How do sounder microphysical retrievals improve on broadband retrievals?

Response: This isn't the type of paper we intended to write. This is a paper about the new and existing AIRS products. As stated in previous responses, as a separate effort, we are currently making comparisons at the pixel-scale to assess the relative performance of sensors like MODIS and AIRS as a function of different scene and cloud types. We hope to say something meaningful at some point whether spectral or spatial resolution is more important for cloud climatology, but we are not trying to answer that question in this paper.

comment: Or stated differently, what specific types of information do the V6 AIRS cloud retrievals add that can compliment existing efforts at existing methods of cloud characterization?

Response: This is currently being assessed. This is a tremendously complex question to answer in a quantitative manner. It requires years of cross-comparisons, in situ validation, radiative transfer closure studies, etc. However, we would argue that the new cloud thermodynamic phase and ice cloud property products appear quite interesting and should contribute something new to the large number of available satellite-based cloud records.

comment: Given the fact that the authors made the introductory statement that CrIS, and IASI (Metop-A and Metop-B) "will continue the AIRS legacy into the future" (lines1– 5, p14483), I find it odd that no further mention is made of CrIS and IASI in the remainder of the paper. Do the authors have a comment on how data continuity among the four hyperspectral sounders in space can be established or promoted?

Response: We have not applied this approach to CrIS and IASI at this time. The authors do think that they could be used to retrieve similar cloud properties, although

these efforts could require significant resources and effort to advance. One idea that is possible is to use the JPL NPP Sounder PEATE as a test bed to try some CrIS cloud property retrievals, perhaps tacked on as a post-processor using the Sounder PEATE retrievals of temperature, water vapor and surface properties as fixed priors in the optimal estimation algorithm. Probably the same can be applied to IASI. Joel Susskind is working on v6-like clouds from CrIS. NOAA NUCAPS produces v5-like clouds from IASI and CrIS. All of these are possibilities.

comment: I am at issue with the list of factors (lines 14–19, p14482) the authors cite as reasons why sounder retrievals have not been used widely in atmospheric science. Response: The authors did not write that 'sounder retrievals have not been used widely in atmospheric science'. Rather, on p. 14482 and lines 14-19, we wrote that 'hyper-spectral IR cloud retrievals have proved elusive until now because of...'. (i) Excessive computational expense. Perhaps this was an issue in 2002 but computational capabilities have improved significantly since then. Could you specify the system for which AIRS processing will be prohibitively expensive?

Response: An optimal estimation approach that retrieves the entire kitchen sink is prohibitively expensive. This is why the cloud clearing (or principal component) approaches are so appealing. They are able to churn through a large data set in fairly quick order.

comment: Apart from computational capabilities, there are well-known statistical methods for compressing data that preserve the information content, most notably principal component analysis. This has been used with great success in hyperspectral sounding retrievals as well as radiative transfer modeling (Huang and Antonelli 2001, Antonelli et al. 2004, Xu et al. 2005).

Response: This is very true and we have acknowledged these efforts by adding these references and appropriate text.

comment: (ii) Complexity and variability of cloud geometry. This is a cloud remote

C6681

sensing problem in general and not limited to infrared hyperspectral sounders.

Response: We agree this is a widely encountered problem in cloud remote sensing.

comment: (iii) Uncertainties in underlying surface and atmospheric state. Again, this is an atmospheric remote sensing problem in general and not limited to infrared hyper-spectral sounders.

Response: We agree this is a widely encountered problem in cloud remote sensing.

comment: (iv) The necessity of using a large channel set. Almost every paper describing a hyperspectral retrieval algorithm has a section on channel selection that focuses on reducing the channel set and improve the signal-to-noise ratio.

Response: The optimization of a channel set is still an active area of research. If future funding is available, we will address this from a formal Shannon information content point of view.

comment: There are also numerous papers describing methods for quantifying the information content of soundings as a means to identify suitable channels for parameter retrieval, or compressing the relevant information into a small number of super channels (or principal components) (references are too numerous to list here).

Response: We have recognized these previous efforts and added the suggested references to the manuscript. Per a previous response to reviewer #1, the fact that we had to use so many channels for the algorithm to converge suggests we need to further explore the prior variances and a priori in the retrieval. Again, this is an active area of investigation. And lastly, according to my best knowledge, the full channel set of AIRS (2374 in total) has never been employed in retrievals (an optimal set of âLij1500 was used for a long time). Moreover, the number of functional channels continues to decrease due to channel popping. Response: See our previous response regarding the instrument characteristics. This is not true.

comment: (v) The ongoing difficulty in developing an automated, rigorous, and fast

retrieval that is applicable to 10+ years of data at single FOV. This is true, but the difficulties have been rigorously addressed and largely overcome with the development of a regression retrieval algorithm (Huang et al. 2004, Weisz et al. 2007, Smith et al. 2012)

Response: While we agree that the regression retrieval algorithm has its advantages, we want to point out that we are not aware of a rigorous error estimate that is associated with each FOV obtained by the regression approach. An optimal estimation approach may be much more expensive, but the errors are quantitative and the attribution of the error to particular inputs or retrieval parameters can be assessed. Both algorithmic approaches have their advantages and disadvantages, as well as their applications.

A list of some minor issues: comment: (1) The term "pixel" is typically not used to refer to the field of view (FOV) of sounders. The authors use "pixel" and "FOV" interchangeably in this paper.

Response: we fixed AIRS to 'FOV'

comment: (2) What is the difference between "process based evaluation of climate models" and the study of "long-term global-scale cloud trends" (p14480)?

Response: Two examples of process-based evaluation are found in Sects. 4.1 and 4.2, e.g., the evaluation of the diurnal variability of convection, or the composite structure of mid-latitude cyclones in climate models and in remote sensing data. Its basically another way to look at how well a climate model represents physical processes, which are inevitably tied to climate sensitivity (e.g., trends).

comment: (3) What are "adverse scale-dependent behaviors in clouds" (line 28, p14481)? âÅĺ

Response: The cited paper describes some modeling experiments that show effects of cloud parameterizations on global-scale cloud organization. For clarity, we removed the reference and associated statement about it.

C6683

comment: (4) The language used is often too informal for a scientific paper, e.g., "Coarser resolution broadband type measurements", (line 12 p14482), "native temporal and pixel scales" (line 8, p14482), "hyperspectral retrievals have proved elusive until now" (line 14, p14482), etc.

Response: We changed these statements to 'broadband observations', 'instantaneous temporal and pixel resolution', and 'hyperspectral retrievals are difficult to calculate'

comment: (5) Can the authors be more specific as to why the TES algorithm was adopted for ice cloud property retrievals? Perhaps a short discussion on the strength of the TES algorithm will suffice.

Response: The TES algorithm was chosen because of the extensive experience at JPL. Also, Bill Irion had been experimenting with an OE retrieval for AIRS (based on the TES retrieval). Naturally, our work on clouds grew out of this effort. We cite Bowman et al. (2006) as the definitive source for the TES algorithm.

comment: (6) If the ice cloud physical property retrievals are to be a legacy for the CrIS and IASI instruments, then a clear description of the channel selection process is required. Currently a PhD dissertation is cited as reference but this is not a widely accessibly source.

Response: We added an additional reference (Jin and Nasiri, 2013) on the cloud thermodynamic phase algorithm and the logic behind the channel selection. With regard to the channel selection for the ice cloud retrieval, see our response to Reviewer #1.

comment: Which channels (in wavenumber units) are currently being used for these retrievals? How will the authors deal with the challenge of channel popping in retrievals that relies on such a small set of core channels?

Response: We added an Appendix that contains all of the channels used in the ice cloud and thermodynamic phase retrieval, and the current specifications for these channels.



Fig. 1. Residual observed-calculated brightness temperature residual for January 2009. The mean global value is -0.299 K, with nearly all places globally within +/- 1K.



Fig. 2. Mean ice cloud effective diameter during January 2007 for the MODIS Science Team, the MODIS CERES Team, the AIRS-LMD, the current AIRS-NASA retrieval, ATSR, and PAT-MOSx.

C6687



Fig. 3. Mean ice cloud optical thickness during January 2007 for the MODIS Science Team, the MODIS CERES Team, the AIRS-LMD, the current AIRS-NASA retrieval, ATSR, and POLDER.



Fig. 4. Mean ice cloud top temperature during January 2007 for the MODIS Science Team, the MODIS CERES Team, the AIRS-LMD, the current AIRS-NASA retrieval, ATSR, and PATMOSx.

C6689



Fig. 5. Zonal averages of AIRS-NASA and AIRS-LMD ice cloud top temperature, optical thickness, and effective diameter for the entire year 2007.