

Response to Referee Comment #1 on

Characterization of dark current signal measurements of the ACCDs used on-board the Aeolus satellite

The authors thank reviewer #1 for carefully reading the paper and providing very useful comments. In the following, referee comments are repeated in green and answers by the authors are provided directly below in black.

General comments:

The manuscript is dedicated to the improvement of experimental data coming from the ALADIN lidar onboard the Aeolus satellite, which provides continuous measurements of atmospheric winds, aerosols, and clouds. The manuscript seeks to solve an important problem of identifying and fixing the experimental issue associated with the so-called “hot pixels” of the ALADIN’s ACCD detectors. The authors suggest a method for pinpointing these pixels, introduce dedicated calibration modes, correct the signals from the affected pixels, and show the results of wind retrievals from the fixed signals. The problems of this kind are not new to the experimental physics, but in this case the study was hindered by the fact that the experimental setup was not available for direct testing in the lab. Still, the authors show that it provided sufficient amount of information for performing the analysis and for fixing the problem. In general, the real state of the atmosphere and the retrieved atmospheric data are linked through a number of conversions and convolutions, each of which can affect the quality of the retrieved parameters. In the case under consideration, the biggest challenge was associated with the missing or damaged pieces of information, needed for the retrieval, namely, with pixels providing the profiles yielded from fringe-imaging or double-edged techniques in some ACCD rows.

Since the backscattered photons carrying the information about the atmospheric properties in this setup are stored in ACCD matrix, one can split the solution of the problem to several steps: (i) identifying the pixels, which are not reliable; (ii) correcting or excluding these pixels from the retrieval; (iii) depending on the previous choice, one has either use the fixed values or modify the retrieval algorithm; (iv) since the initial retrieval algorithm did not take into account the possibility of hot pixels spoiling the inputs, one has a right to impose physical constraints on the retrieved data to fix the affected points. The authors did an excellent job for (i) and then they followed the correction scheme of (ii) and ended up with (iii). From this point of view, the work is impeccable. Still, I’d suggest to consider a bigger picture and to look at the problem at a different angle. Perhaps, the authors did it in the background and found that it didn’t solve the problem, but I found no trace of it in the manuscript, so at least this is worth a discussion.

Let me explain. Looking at Fig. 1 of the manuscript and comparing it with the Fig.11, one can see that the experimental setup has a certain redundancy in a sense that the peak in the Mie signal almost never is narrower than 3 bins and, in some cases, a naked eye distinguishes 4 bins filled with non-zero signal. At the same time, Fig. 11 tells us that the situations when two adjacent horizontal pixels are “hot” are rare. Knowing that this detector is characterized by a low noise, one can make use of the remaining available information and still get a reliable result. To prove this, I performed a simple test illustrated in the attached Figure. The panel (a) shows the Mie detector mask, which is consistent with Fig. 11 of the manuscript, but converted to a binary (good/bad) form. Panel (b) shows the simulated signal, which qualitatively resembles that of Fig. 1 of the manuscript, but which passes through the hot pixels of the mask (a) for demonstration purposes. For each row, the exact position of the peak

corresponding to exact value of the wind is stored for reference. Panel (c) shows the same signal with hot pixels masked out. The Poissonian noise was added to the pixel values to imitate the detector's behavior. Then the fitting procedure based on sliding profile correlation approach similar to those used in [Goldberg et al., 2012] and [Feofilov and Stubenrauch, 2019] was applied both to a full set of input data and to a masked one. The procedure uses the knowledge of the profile of the fringe-imaged signal along the columns and this profile is supposed to be known with high accuracy. The resulting retrieval uncertainties are shown in panel (d) of the Figure. As one can see, the position of the peak retrieved from incomplete data does not change that much compared to the retrievals from the unmodified datasets, and the uncertainty in pixels converted to wind speed uncertainty is of the order of 0.03 m/s. This is just a rapid exercise, which should be done in a different way for Rayleigh signals, but it leads to an important question – even though the fixed hot pixels provide a dataset compatible with the rest of the processing chain, wouldn't it be easier and safer to exclude them from the consideration and to update the procedure? I understand that this is not what the manuscript is about, but it's a major philosophic question whether one should use fixed values from a damaged detector or use a reduced dataset profiting from the redundancy of the data. The latter approach does not diminish the significance of the work, but if it proves to give more reliable data through a simpler procedure, it should be considered.

The second question is about aforementioned step (iv) – I believe, the retrieval procedure could profit from the physical constraints of the following kinds: (a) point-to-point wind speed change and (b) point-to-point aerosol/cloud properties change. Both are easy to justify and both can serve as an additional quality control mechanism at early stage – if sudden unphysical jumps are found, the pixels are removed from the retrieval and the values are interpolated, masked, and so on.

Response to General Comments:

Thank you very much for your valuable feedback. It is highly appreciated that you carried out a simulation study to test different approaches to tackle the hot pixel issue by omitting the hot pixels in the wind retrieval.

First of all, it is important to mention that Figure 1 of the manuscript is only a simplified sketch of the ACCD to illustrate its working principle. The figure does not consider the full Mie fringe shape which is spread over all 16 ACCD pixels – even at the edges of the ACCD the values do not approach zero. For simulation studies, this could be approximated with an Airy function. On top of that, also the spectrally broad bandwidth Rayleigh signal and the solar background are part of the Mie signal. I think these aspects have to be considered in your simulation study presented in Figure 1. Apart from that, it is not clear which values for the atmospheric signal levels and hot pixel amplitudes were used in your simulation. Typical values for the Mie channel are 15 LSB and 5 LSB (both measurement level) for the atmospheric signal level and the hot pixel offset, respectively.

There are several reasons not to exclude hot pixels from the wind retrieval. An important aspect is the limited number of available pixels on the ACCD. As illustrated in Figure 1 of the manuscript, the ACCD has only 16 columns which makes each single pixel indispensable for the retrieval. As already mentioned above, each pixel contains valuable information. Although the major part of the Mie fringe is only contained in three to four pixels in the center of the ACCD, the pixels at the edge still contain valuable information which is needed to correctly derive the fringe centroid position, the fringe width and the broad bandwidth Rayleigh offset, used for the computation of the SNR and scattering ratio.

Moreover, hot pixels are not damaged and thus, still provide valuable measurement signals which can be used in the wind retrieval. Their main characteristics are increased dark current signals that are changing over time. And as introduced in Sec. 2.3 of the manuscript, DUDE measurements allow for a pixel-wise determination of the dark current signals for the correction of the increased dark signal values.

As the omission of hot pixels is not feasible, the approach of correcting hot pixels using DUDE measurements is the most suitable solution in the Aeolus. On the one hand, this method was straightforward to implement for both channels without having to redesign the well-established wind retrieval algorithms. On the other side, this method is also capable of dealing with the steadily increasing number of hot pixels without having to adjust the algorithm after each hot pixel occurrence. This is also a very important aspect for an operational satellite mission like Aeolus. In addition to that, this method is also compatible with the Aeolus L2A aerosol retrieval algorithms.

“The second question is about aforementioned step (iv) – I believe, the retrieval procedure could profit from the physical constraints of the following kinds: (a) point-to-point wind speed change and (b) point-to-point aerosol/cloud properties change. Both are easy to justify and both can serve as an additional quality control mechanism at early stage – if sudden unphysical jumps are found, the pixels are removed from the retrieval and the values are interpolated, masked, and so on.”

First of all, it's important to know that the Aeolus processing chain is strictly sequential due to near-real-time (NRT) requirement for Aeolus data products and does not contain feedback loops which means that at the L1B processing stage the L2B products are not yet available. For the NRT processing there is also the requirement to provide the L2B data products only 30 minutes after the downlink of the raw data. Since the dark signal correction is part of the L1B processor, we cannot use additional information about point-to-point wind changes at this processing stage. For the future, it is planned to implement an additional check into the L2B processor to detect hot pixel steps in the L2B winds-based comparisons with the ECMWF model background (see Sec. 3.3 of the manuscript). As demonstrated in Sec. 3.3 of the manuscript, some hot pixels induce large O-B deviations in the wind products. So, this check aims at detecting hot pixel induced bias steps which occur in-between two DUDE measurements by analyzing the wind speed difference w.r.t to the ECMWF model. Winds exceeding a certain threshold will be flagged as invalid at range bin level.

The prerequisite of masking hot pixels in the processing is the detection of hot pixel induced signal steps while the instrument is in measurement mode which is very challenging. In simple terms, the atmospheric return signal measured in the memory zone of the ACCD during wind mode is a composition of atmospheric and dark current signal. To properly detect hot pixel induced steps one must be able to distinguish between atmospheric and dark current signal induced steps. For this, the ratio between the atmospheric and the dark signal intensities is important. Thus, for the Rayleigh channel, the location of the pixel on the ACCD is the key factor. For instance, for pixels which are fully covered by the Rayleigh spots, i.e. Rayleigh column positions 1-5 and 9-13, the atmospheric signal is usually too dominant and variable compared to the dark current signal to distinguish between atmospheric and dark signal induced steps in the atmospheric return signal. Figure 1 below shall help to better understand this problem. It shows signal intensities (red line) for Rayleigh hot pixel [20, 2] which is covered by one of the Rayleigh spots at observation level. In order to minimize effects induced by changes in the range bin settings or the altitude of the range bin, the signal is median-filtered using a window size of 400 observations. The vertical dashed green lines in the plot indicate the locations of

the DUDE measurements and the blue line with the second y-axis shows the dark current signal correction value applied to the signal in LSB. The shaded areas which span certain time periods indicate the validity times of the different orbits during the day. The DUDE measurements at 05:45 UTC and 16:15 UTC indicate two jumps of about 2 LSB in the dark current signal on that day. Analyzing the atmospheric return signal (red line) shows that these jumps are not visible here.

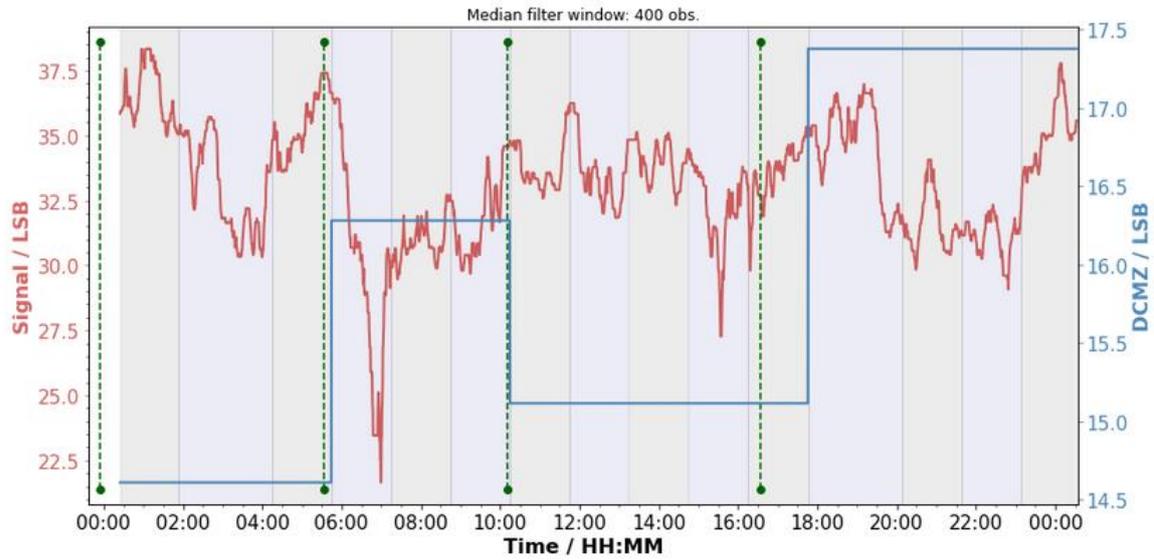


Figure 1: Median filtered (window size: 400 obs.) ALD_U_N_1A pixel intensities for Rayleigh hot pixel [20, 2] (red) on 2020-03-03 together with the DCMZ correction value in blue. The vertical green bars indicate the times of the DUDE measurements.

A general discussion of potential hot pixel correction methods was added to Sec 3.3 of the manuscript. It is made clear that the omission of hot pixels is not feasible and advantages of the implemented correction scheme are clarified:

Same as for the Mie channel the magnitude of the dark current signal-induced Rayleigh bias depends on the signal level of the backscatter signal. Generally speaking, the dark current induced bias is more constant in the Rayleigh than in the Mie channel

430 due to the more constant Rayleigh signal compared to the strongly varying Mie signal from clouds and aerosols. For the correction of hot pixels in the Aeolus NRT processing several options were considered. One way would be to omit hot pixels from the wind retrieval. However, this approach is not feasible due to several reasons. First of all, the very low number of 16 pixels per column makes each pixel indispensable in the wind retrieval. It is important to note that also the pixels at the edge of the ACCD contain valuable information necessary to retrieve the wind information. In addition, hot pixels are not damaged

435 and still contain information that can be used in the wind retrieval. Another correction method is the interpolation of hot pixels using the information from neighboring pixels. Considering the rather coarse vertical resolution of Aeolus measurement of 250 m up to 2000 m and the non-linear vertical wind shears in the troposphere, especially vertical interpolation could be highly erroneous depending on the vertical wind shear and the range bin settings.

As a result, it was decided to implement an algorithm which corrects the increased dark current signal offsets of hot pixels based on frequent dark current signal characterization measurement. This correction scheme was successfully implemented into the wind retrieval of the Aeolus operational processing chain on the 14th of June 2019.

440 Already shortly after the hot pixel issue was recognized, an on the fly correction scheme was successfully implemented into the wind retrieval of the Aeolus operational processing chain on the 14th June 2019.

For the correctionIn detail, the dark signal characterization obtained from frequently performed DUDE measurements is used

445 for a pixel-wise dark signal correction, i.e. a DSNU correction, of adjacent wind measurement signals. As this kind of correction was not foreseen before launch to be performed on a regular basis, dedicated instrument modes and algorithms had to be developed after launch. The implemented correction approach has the advantage of being applicable to both channels without having to redesign the well-established wind retrieval algorithms. Moreover, this method is capable of dealing with the steadily increasing number of hot pixels without having to adjust the algorithm after each hot pixel occurrence and is also

450 compatible with Aeolus L2A retrieval algorithms.

Moreover, the following paragraph of Sec 3.3 was modified to make clear that with the current processor configuration point-to-point wind speed changes cannot be used in the dark signal correction of the L1b processor:

This example demonstrates that even with DUDE measurements performed at high frequency, a perfect dark signal correction is not possible. It is also clear that the performance of the correction scheme depends on the behavior of the hot pixel. In case of RTS-like characteristics as shown for Mie pixel [13, 9] (see Fig. 5) the correction performs poor compared to hot pixels with sporadic shifts. Nevertheless, this approach works fine to remove the constant proportion of the dark current offset and in any case reduces periods of enhanced dark current induced bias sufficiently. In order to further mitigate hot pixel induced effects, a further check will be implemented for the Aeolus level 2B product in the future. This check is based on comparing the Aeolus with ECMWF model winds for each range gate. In case, the difference between both exceeds a certain threshold, the Aeolus winds of the affected range gate will be flagged invalid. For example, the period between 14:15 UTC and 20:30 UTC (see Fig. 10) would be flagged as invalid by the check. It is important to mention that the Aeolus processing chain is strictly sequential and does not contain feedback loops which means that at the L1B processing stage the model comparisons from the L2B products are not yet available. For the NRT processing there is also the requirement to provide the L2B data products only 30 minutes after the downlink of the raw data. Since the dark signal correction is part of the L1B processor, it is not possible to use this information to mask or flag hot pixel offsets in the L1B processing stage.

Specific comments:

Lines 200-209: if CIC noise is important, how does this fact match the "low-noise detector" statements above? Some numbers are needed here, so that the reader could make his/her own conclusions.

The paragraph you are referring to is related to the role of CIC in the generation of hot pixels. Thus, it is about CIC increasing the mean dark current signal rather than dark current signal noise. However, information about the dark current signal noise and read-out noise is added to Table 1 of the manuscript. Considering the pulse repetition frequency of 50.5 Hz and 18 pulses per measurements, the in-orbit dark current signal rates are 0.49 LSB/s and 0.24 LSB/s for the Mie and Rayleigh channel. Considering the Poisson distribution of the dark charges these values correspond to 0.78 e⁻ – 0.89 e⁻ rms dark current noise for a residence time of the signals in the ACCD of one measurement which is 0.4 s. Note that these numbers already include CIC. The numbers for the read-out noise of the detector are between 4 e⁻ and 6 e⁻ rms.

These values were added to Table 1 of the manuscript and the text of Sec. 2.1 was modified accordingly:

Table 1: Specifications and in-orbit performance of the Aeolus ACCDs

Parameters	Value
Type	Thinned backside illuminated accumulation Si-CCD
Area	Imaging zone: 0.43 mm x 0.43 mm – 16 x 16 pixels Memory zone: 0.43 mm x and 0.75 mm - 32 x 25 pixels
Pixel size	Imaging zone: 27 μm x 27 μm Memory zone: 13.5 μm x 30 μm
Operating temperature	-30 C°
Temporal resolution	2.1 μs – 16.8 μs / 250 m – 2000 m for atmospheric layers (#1 - #24) 3750 μs for solar background (layer #25)
Quantum efficiency	0.85
Charge transfer efficiency	0.9999
Radiometric gain	Mie: 0.68 LSB/e ⁻ , Rayleigh: 0.44 LSB/e ⁻
Dark current signal rate	Mie: 0.49 LSB/s, Rayleigh: 0.24 LSB/s (in-orbit values)
Dark current signal noise	0.78 e ⁻ – 0.89 e ⁻ rms (root-mean-square) (in-orbit values)
Read out noise	4 e ⁻ – 6 e ⁻ rms

Line 278: can one prove this statement about the DUDE correction with some formula or reference? At the moment, there are only qualitative statements here.

The following example shall depict the dependency of the DUDE correction on the characteristics of the hot pixels. As shown in the manuscript, we carry out DUDE calibrations four times per day. In case, there are hot pixel induced shifts in-between the DUDE measurements, the dark current of the affected pixels is not properly corrected until a new DUDE measurement is performed. As illustrated in Sec. 4.2 of the manuscript, the characteristics, i.e. the fluctuation rate and level amplitudes, are very different for the hot pixels. So, for stable hot pixels which do not change their level often, the frequency of the DUDE measurements is not very critical. But for hot pixels with high fluctuation rates, it is much more likely to have dark signal induced steps in-between the DUDE measurements which leads to biased wind results until the next DUDE calibration is carried out. This effect is also illustrated in Figure 10 of the manuscript which shows the signal intensities of the very jumpy Mie pixel [13, 9] on 2019-11-14. It shows a dark current induced signal decrease of 8.0 LSB at 14:15 UTC. As a result, the dark current signal is overestimated as the dark current calibration based on the DUDE measurement from 13:15 UTC is still active. Consequently, the winds of range bin 13 are biased for the period between 14:15 UTC and the next DUDE update at 20:45 UTC (also explained in Sec. 3.3 of the manuscript).

Lines 300-320: perhaps, it's a matter of preferences, but how does this approach compare to a simple 3-sigma test? Another approach, which could be also useful for detecting hot pixels as well as identifying the nature of the noise is building and analysing Fourier spectra of the temporal sequences for each pixel. Most probably, the spectra of hot pixels will be different from those of "normal" ones and hot pixels of a different nature will reveal this in the spectra, too.

For our analysis an approach is needed which is not only capable of differentiating between normal and hot pixel behaviour but also can find the exact temporal index of the dark current signal shifts. This information is needed to derive information about the hot pixel signal amplitudes and the time spent at the different dark signal levels. That's the basis of the categorization of the hot pixels as shown in Sec. 4.1-4.2 of the manuscript and, especially, of the analysis of the RTS characteristics (Sec. 4.2.1).

A simple 3-sigma-test would have probably been suitable to detect hot pixels but it would not have been possible to derive further information about the hot pixel characteristics (amplitudes and time spent at a dark signal level) as mentioned above. Figure 2 down below shows the dark signals of Mie pixel [20, 2] (same as Figure 15 (top) of the manuscript) at observation level together with thresholds obtained from the mean value ± 3 * standard deviation of the dark signal. It shows that the 3-sigma threshold is exceeded multiple times which would result in classifying this pixel as "hot". However, comparing Figure 2 with Figure 15 (top) of the manuscript clearly indicates the advantages of the sophisticated approach presented in the manuscript, i.e. the detection of the dark signal segments and the derivation of the dark signal levels.

I agree that an approach based on the analysis of Fourier spectra of the dark signals would also have helped to learn something about the different natures of the hot pixels. However, as mentioned above, this approach would not have provided the needed information about the indices of the switches between the dark current signal levels which is the prerequisite to derive temporal RTS characteristics (see lines 349-353 of the manuscript).

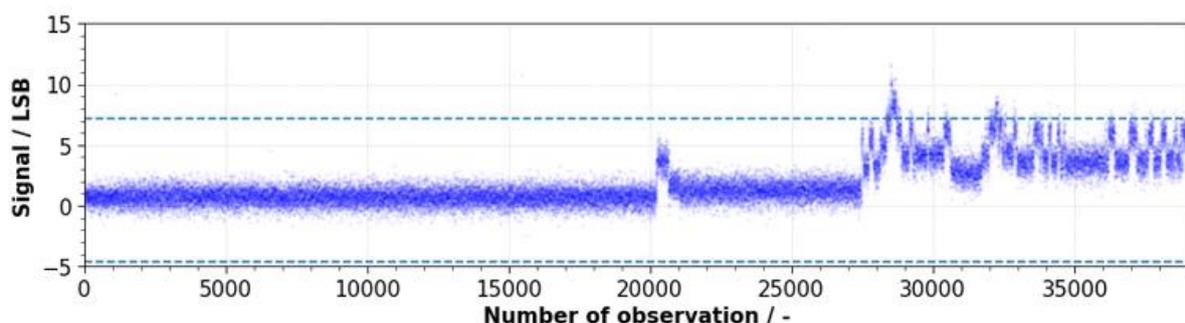


Figure 2: Dark signals of Mie pixel [20, 2] at observation level.

To clarify the need of the proposed time series segmentation algorithm, the introduction of Sec. 3.1 was changed as follows:

The motivation for a detailed characterization of permanent dark current anomalies is twofold: a) it supports investigations for the underlying root causes of the hot pixel issue and b) the number and magnitude of dark signal shifts define the impact on the wind observations and the DUDE correction. In order to fulfill b) a more sophisticated algorithm such as a simple 3-sigma test, which can differentiate between normal and hot pixels, is necessary. The presented algorithm is also capable of detecting temporal indices of the sudden shifts in the dark signal time series which is the basis of the categorization of the permanent dark current anomalies.

Line 338: again, spurious changes could have been filtered out by Fourier smoothing procedure.

In our case Gaussian Kernel Density Estimation (KDE) is used to estimate the probability distribution of the dark signal values. As depicted in the Sec. 3.1, the key parameter of the KDE is the bandwidth parameter which controls the size of the Gaussian kernel at each data point. A bandwidth value which is too high leads to an over-smoothed density estimation which probably hides important structure. On the contrary, a too narrow bandwidth will put too much emphasis on single points and thus, result in a density estimation curve with too many modes. As a result, it's quite important to find suitable values for the bandwidth. There exist several algorithms for this task. The standard approaches are Scott's rule of thumb (Scott, 1992) or Silverman's rule of thumb (Silverman, 1986). However, these are simplified approaches which assume normally distributed data. In the presence of dark signal anomalies, the underlying distribution function is not known. This is why, non-parametric approaches are needed to estimate the optimum kernel bandwidth. One established non-parametric method is Maximum-Likelihood Cross Validation (Duin, 1976) which is a purely data-driven method to derive the bandwidth parameter. Here, a metric for different values of the bandwidth is computed by estimating the kernel function on a subset of data and computing and evaluating this function on the rest of the data. The advantage of this method is that it's purely data-driven, meaning that no assumptions on the underlying data is needed. In our case this method provided reliable results. The drawback of this method is its high computational costs. However, for the application to one-dimensional data that is not a big issue. If this would have been a problem, it could have been an option to implement a fast Fourier-based Kernel density bandwidth estimation as depicted in (Gramacki and Gramacki, 2017).

To further motivate the selection of the KDE bandwidth selection method, the following was added to Sec. 3.1 of the manuscript:

On the contrary, important features may be smoothed away when applying excessive smoothing. There exist several algorithms for this task. For the purpose of analyzing dark signal anomalies, a non-parametric method which does not require any assumptions of the underlying data distribution is needed to find the optimal bandwidth parameter. Thus, maximum-likelihood cross-validation is used to determine the bandwidth parameter which is an established method for the objective, data-based derivation of the bandwidth parameter (Jones et al., 1996). This method computes a metric for different values of the bandwidth by estimating the kernel function on a subset of data and computing and evaluating this function on the rest of the data. The advantage of this method is that it's purely data-driven, meaning that no assumptions on the underlying data is needed.

Lines 400-410: see the general comments – perhaps, the discussion should be updated.

See our response to your general comments.

Lines 430-435: how does this correction compare to vertical interpolation?

So far, vertical interpolation has not yet been considered as valuable approach to correct for hot pixels. Considering the rather coarse vertical resolution of Aeolus measurement of 250 m up to 2000 m and the non-linear vertical wind shears in the troposphere, vertical interpolation could be highly erroneous depending on the vertical wind shear and the range bin settings. In case, Aeolus would be able to provide measurements with a better vertical resolution, vertical interpolation might indeed be an option.

As part of the modifications of Sec. 3.3, the mitigation approach of vertical interpolation is discussed (see screenshot under “General Comments”).

Line 439: a median correction is applied, which does not eliminate sporadic events. Even though it smooths them out, their erroneous nature is included in the results. On the other hand, gradient-based or Fourier filtering would have removed a non-physical part of the signal.

You are correct signal spikes such as dark signal introduced transient events (introduced in Sec. 3.2 of the manuscript) can influence the result of the median filtering. In this case, the correct way to perform this kind of analysis would have been to identify transient events at measurement level and remove them before averaging the measurements to observations. However, the analysis has shown that on average only 0.24 % of the measurements are affected by transient events (see Sec. 4.3 of the manuscript). This is why, this effect is considered to be negligible.

Line 500: it would be interesting to recalculate these 6% into a weighted percentage of pixels used in retrievals. For example, pixel [9,13] is used often whereas [1,1] is not.

As explained in the response to your general comments, all pixels of both channels are used in the wind retrieval. As a result, it is not possible to derive a weighted average of the pixels used in the retrieval.

Line 550: Linear trend is interesting here. If the damage is due to high energy particles hitting the ACCD then the slope should change with time, but 6% is too small a number for this to be noticed.

Yes, this is correct. In line 515 of the manuscript it is mentioned that the solar activity is currently at a minimum and it will be interesting to see if this has an influence on the rate of hot pixel generation. Moreover, it is planned to redo this kind of analysis at the end of the mission when a larger dataset of hot pixels is available (see responses to reviewer #2).

Line 689: cosmic particles partially penetrate the atmosphere, so this is not a 100% proof.

This is correct. This argument will be reformulated as follows:

715 The fact that two hot pixels Mie [16, 15] and Mie [24, 3] – both of them in different hot pixel categories and with similar characteristics of hot pixels emerged in-orbit – were already present before launch supports the hypothesis of an origin which is not solely related to the fact that Aeolus is operated in space environment with very harsh radiation conditions-might-be independent of space radiation. However, other radiation sources within the instrument or even within the ACCD package

Lines 701-702: first, we did not see this in the manuscript and second, it should be considered in the light of the exercise demonstrated in General comments.

The intention was to make this point clear in Sec. 3.3 of the manuscript. In this section, a simple calculation is provided to demonstrate the effects on the Rayleigh wind results (see lines 410-418 of the manuscript). It was demonstrated that a hot pixel induced signal elevation of only 10 LSB already results in an HLOS error of already about 2.6 m/s. In order to further clarify the correlation between the wind error and hot pixel offsets, the atmospheric return signal of Rayleigh hot pixel [11, 2] will be added to Figure 8 of the manuscript (see Figure Figure 3). The modified figure clearly demonstrates the correlation of the increased O-B bias values around 400 hPa and the dark signal induced steps in the atmospheric return signal.

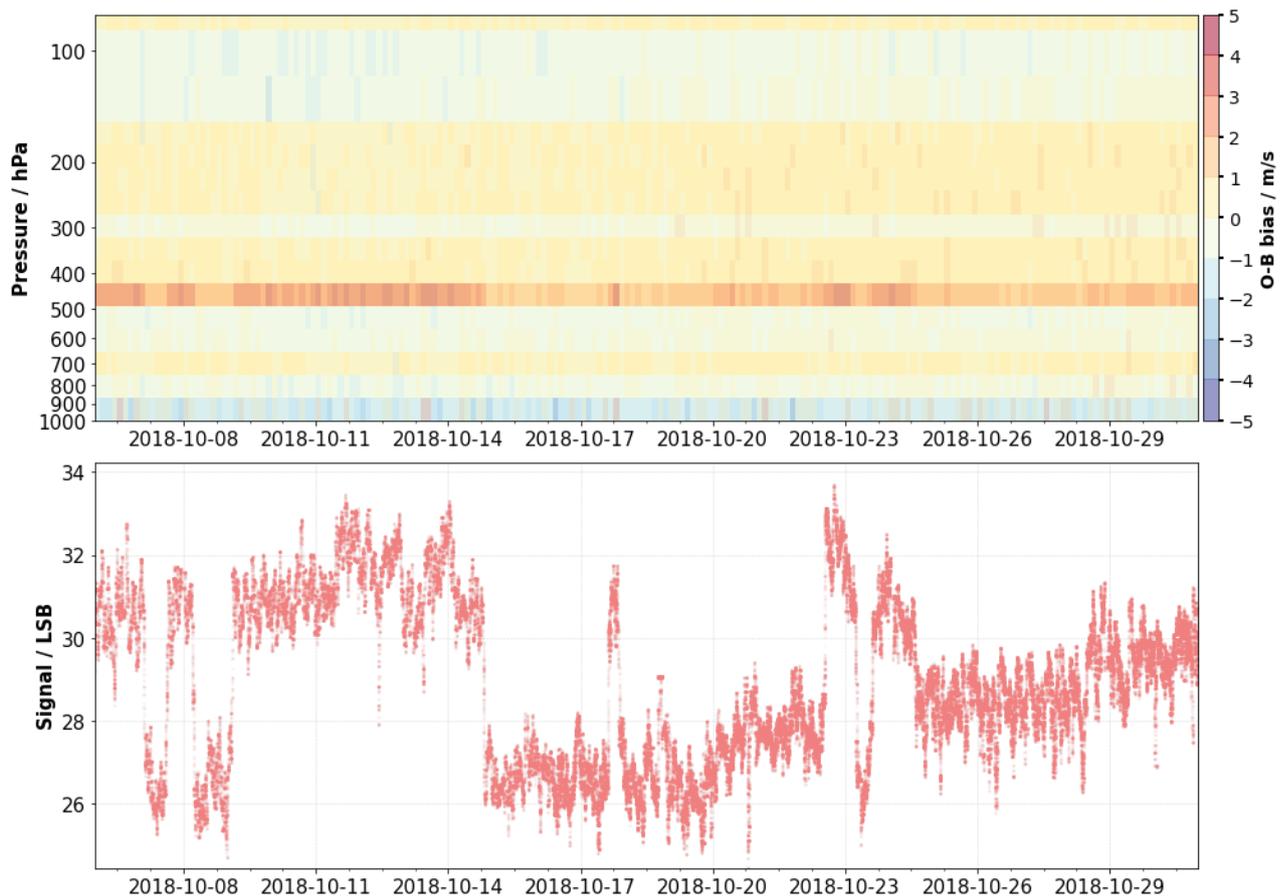


Figure 3: (Top): Comparison between Aeolus L2B Rayleigh-clear HLOS winds and the ECMWF model equivalents between 2018-10-06 and 2019-10-31. The plot shows the mean difference between the observation (O) and the background (B) (short-range forecast) model field as a function of pressure and time. (Bottom): Median filtered (window size of 400 observations) signal intensities of Rayleigh hot pixel [11, 2] during wind measurement mode.

Figure 8 of the manuscript was changed as shown in Figure 3 above. Moreover, the explanation of this figure was updated in Sec. 3.3 of the manuscript.

Line 752: numbers are missing here: uncertainty/bias after the correction vs uncertainty/bias before the correction.

As shown in the manuscript, the hot pixel induced bias mainly depends on the hot pixel characteristics. In Sec. 3.3 of the manuscript, the effect is demonstrated on the basis of a simplified example for the Rayleigh channel. Here, a realistic hot pixel dark current signal offset of 10 LSB present for Rayleigh spot A is assumed. In case, no dark signal correction is applied this would result in a wind bias of 2.6 m/s HLOS. Also, Figure 3 (Figure 8 of the manuscript) depicts hot pixel induced wind bias values of several m/s. The random error of the wind measurements is not affected by the hot pixel correction.

As mentioned in Sec. 2.2 of the manuscript, the noise characteristics and thus, the random error, are mainly driven by the shot noise of the signal.

785 A combination of dedicated instrument calibration modes and ground processors were developed to allow for a pixel-wise dark signal correction of the wind signals already shortly after launch. It was demonstrated that this correction is capable of correcting for the dark signal non-uniformity arising from hot pixels on the ACCD and thus, successfully removes hot pixel induced wind bias of up to several m/s. It is expected that this correction will work throughout the whole mission lifetime no

Lines 301, 342, and elsewhere – in some PDF viewers, the font used for Python module names looks strange.

The Python module names in the manuscript will be changed to normal font.