



Supplement of

MAGARA: a Multi-Angle Geostationary Aerosol Retrieval Algorithm

James A. Limbacher et al.

Correspondence to: James A. Limbacher (james.limbacher@noaa.gov)

The copyright of individual parts of the supplement might differ from the article licence.

MAGARA expanded aerosol and surface retrieval description

The following bullets succinctly describe the MAGARA retrieval approach:

- 1) The algorithm assumes that the daily averaged fine- and coarse-mode components are constant and retrievable for a given day. This means that we assume fine- and coarse-mode properties, apart from FMF and AOD, are fixed for a given day. This is always the case for the first two retrieval steps, corresponding to Sections 2.2.1 and 2.2.2.
 - a) **For Section 2.2.3** (i.e., the 10- to 15-min retrievals of FMF and AOD), if 10- to 15-min retrieved FMF is < 0.1 , we supplement the daily averaged fine-mode component along with its fraction for that time and day with a fixed fine-mode component as described in Section 2.2.3. This is done because MAGARA has minimal ability to discern fine-mode particle properties for such a low FMF. The same is done for the daily averaged coarse-mode component if the coarse-mode fraction (CMF) is < 0.1 , as MAGARA would have minimal ability to discern coarse-mode particle properties.
- 2) The algorithm assumes that BRF^{Surf} at any specific time of day does not vary much from day to day over a given retrieval period (typically on order 7-14 days). The algorithm accounts for slight variations in MAIAC BRF^{Surf} over land due solely to small changes in solar geometry from day to day. Over water, retrieved BRF^{Surf} is constant from day to day for a specific time of day.
- 3) There is enough information content in the 5-10 (1-2 ABIs) TOA BRFs to perform FMF and AOD retrievals at the cadence of the measurements with no temporal averaging.

Because FMF is allowed to vary at the cadence of the measurements, 1) is not a particularly stringent constraint on the result, as the algorithm still retrieves varying amounts of the fine- (e.g., smoke) and coarse-mode component fraction (e.g., dust) in the same region during the same day (Section 2.2.3). 2) is generally a reasonable assumption, except for snowfall and melt, rapid greening, rapid senescence, or flooding. ***If we assume 2) is valid, given an assumed aerosol optical model and AOD, we analytically calculate BRF^{Surf} at any time of day without using a BRF^{Surf} model.*** However, we must know aerosol properties and loading to get an accurate constraint on BRF^{Surf} . Once we have daily averaged fine- and coarse-mode component fractions and a well-constrained surface, we retrieve the generally more temporally variable fine-mode fraction and AOD at the cadence of the ABI measurements.

The first of these three retrieval steps involves iterating between analytical calculations of surface BRF, retrievals of time-averaged AOD, and retrievals of daily averaged fine- and coarse-mode aerosol component fractions. The second retrieval step uses the daily averaged fine- and coarse-mode aerosol component fractions to refine the analytical calculation of the surface BRF and time-averaged AOD. This is done via iteration between retrievals of time-averaged AOD and retrievals of surface BRF. The third and final step uses the refined surface BRF and the retrieved daily averaged fine- and coarse-mode component fractions from the first step to obtain the AOD and FMF AOD at the cadence of the ABI measurements.

This series of retrieval steps runs significantly faster than attempting to solve the system of linear equations $\sqrt{\mathbf{w}} \cdot \mathbf{A} \cdot \vec{x} = \sqrt{\mathbf{w}} \cdot \vec{b}$ for all aerosol and surface parameters. For a 10-day run of MAGARA with 2 ABI instruments (GOES-16 + GOES-17) and at least 50 retrievals over the course of a day, these 5000 observations (5 bands, 2 ABIs, 10 days, 50 times) yield 1000 pieces of information about AOD and FMF, 170 pieces of information about aerosol particle properties, and 500 pieces of information about retrieved BRF^{Surf} . In the case of a simultaneous retrieval, \mathbf{A} would represent the derivative of TOA modeled BRFs with respect to all retrieved parameters (at an initial starting guess for our solution vector), resulting in a matrix size of 5000 by 1670 elements. Of the 8.35 million elements in the matrix (for a single pixel), 8.25 million elements would be populated with 0. Even if sparse solvers were 100 percent computationally efficient, just populating the matrices would take an inordinate amount of time. As such, the retrievals presented here are broken down into independent retrievals of different parameters using a much smaller subset of the total TOA BRFs.

MAGARA Expanded Daily Fine- and Coarse-Mode Component Fraction Retrieval

The following paragraphs represent a detailed portrayal of the flow chart presented in Figure 2. The rest of this section describes how the algorithm iterates to a solution for daily averaged fine- and coarse-mode component fractions. In the process, it must also iterate towards a more refined retrieval of time-averaged AOD and BRF^{Surf} . Output for all these parameters is generated in this code segment, but only the daily averaged fine- and coarse-mode component fractions are final.

After regridding and initialization, the rest of this portion of MAGARA is summarized as a series of independent retrievals of BRF^{Surf} , aerosol loading, and daily averaged fine- and coarse-mode component fractions *for an individual pixel*. The following bullet points briefly summarize the retrieval aspects of this section, with significantly more detail presented in Figure 2.

- a) **For retrievals of BRF^{Surf} and AOD only**, this section of the algorithm first builds daily aggregate (based on component fraction) LUTs of all RT parameters *at every AOD* found in our LUT (for each time, band, and ABI). This is done solely for speed reasons, as it allows us to use a modified (1/17th the size) LUT with one aggregate component for BRF^{Surf} and AOD.
 - i) The initial guess for this daily aggregate model is AOD = 0.10, FMF = 0.9, with the fine mode comprised of component 7 and the coarse mode comprised of component 16.
 - ii) Regardless of the previously prescribed/retrieved aggregate daily model, the algorithm includes the following aggregate model for all AOD < 0.5 in the LUT: FMF = 0.75 with component 14 as the fine mode and component 16 as the coarse mode. This model represents 100% of the weight for LUT parameters at an AOD = 0 and 0% of the weight at an AOD = 0.5, with linear weighting in between. The initial value choices made in i) and ii) are explained below.

- 65 iii) These *daily aggregate models are used only for the surface and AOD iterations in this subsection*. The output daily averaged fine- and coarse-mode component fractions can deviate significantly from both i) and ii) above, even when daily averaged AOD is low.
- 70 iv) These daily aggregate models will have minimal impact on 10–15-minute retrievals of fine-mode fraction, especially if $AOD > \sim 0.3$, but should weigh the distribution of daily retrieved component fractions towards the values in ii) if daily aerosol loading is minimal.
- b) ***Once these daily aggregate LUTs have been created/updated***, the algorithm uses these aggregate LUTs combined with the initial cloud-screening weights to ***analytically*** calculate the BRF^{Surf} for a given time of day, spectral band, and ABI.
- 75 c) The algorithm then retrieves time-averaged AOD with different time averages for different iterations using the BRF^{Surf} retrieved in b).
- d) ***For retrievals of daily averaged fine- and coarse-mode component fractions***, the algorithm interpolates RT LUTs for each aerosol component in Table 1 to the time-averaged AOD retrieved in c) above, using the BRF^{Surf} retrieved in b) above to estimate the surface contribution to the TOA signal for each component. The algorithm then retrieves daily component fractions for a given day using all spectral and ABI data for every time of day (150+ total observations).
- e) The algorithm updates the initial aggregate model based on d). Initial cloud weighting is updated with the cloud weights negatively correlated to retrieved AOD and Cost.
- 80 This iterative process, looping from a) - e), occurs multiple times and results in ***output daily averaged fine- and coarse-mode component fractions, a crude estimate of BRF^{Surf} , and a crude estimate of daily-averaged AOD***. The particle-property choices for both the initial guess (FMF=0.9, fine-mode component 7) and the fine mode at low AODs (FMF=0.75, fine-mode component 14) are educated guesses based on our experience with the MISR RA. Both of these fine-mode components correspond to 0.12- μm effective radius aerosol models, either non-absorbing or with a SSA of 0.9, that the MISR RA often
- 85 selects by as natural background, pollution, or smoke aerosol. ***The algorithm is weighted towards these models because the algorithm does not typically have the sensitivity to retrieve these properties when aerosol loading is low***. If we did not use aerosol particle property weights, we would see significant, unphysical scatter (i.e., errors) in pixel-to-pixel retrievals of aerosol particle properties, which would propagate into retrievals of FMF and AOD. We provide some retrieval statistics for the Camp Fire example in the next section to offer some evidence for the selection of these parameters. The MAGARA code
- 90 corresponding to the daily averaged fine- and coarse-mode component fraction retrieval first takes the trimmed RT LUT described earlier and converts from solar/viewing geometry to the geometry of our pixel corners (so we can bilinearly interpolate), as done in the MISR RA, for the dates/times corresponding to the GOES ABI acquisition times. The rest of this portion of MAGARA is nested within a 50-pixel loop in the x direction and a 50-pixel loop in the y direction. Spatial interpolation of our RT LUT parameters is only performed on every 5th pixel in the spatial x direction and on every 5th pixel

95 in the spatial y direction. This results in an additional 25x speedup for interpolation because solar/viewing geometry varies minimally over such small spatial scales.

The algorithm then begins a loop over an index we call `iterInd`, iterating from 1 to 3, with the main output being the daily averaged fine- and coarse-mode component fractions. This `iterInd` index controls the length of two nested loops within the code, with the lengths of both loops varying from 2, 3, and 5, for `iterInd` = 1, 2, and 3. As the retrieved surface albedo
100 becomes more refined, we are more confident the algorithm will iterate towards a correct solution, so we allow the algorithm to iterate more times. After spatially interpolating to a given pixel's x and y location, the algorithm applies the correction $1.0 / (1.0 - s_\lambda * A_\lambda)$ to account for multiple light reflections off of the Earth's surface (e.g., Equation 6). The surface albedo A_λ initially comes from MAIAC or is assumed to be 0 initially if over water, and it is further refined upon every increment of `iterInd`. **For all but `iterInd` == 3, MAGARA makes an initial guess of daily average fine- and coarse-mode component fractions (see
105 bullets a_i and a_{ii} in section 2.2.1), with an initial guess of AOD = 0.10. These daily component fractions are used only for the surface and AOD iterations, not for retrieval of output daily aerosol particle properties. For `iterInd` == 3, MAGARA uses the output from the previous iteration for both AOD and component fraction. These prescribed/retrieved aerosol component fractions are weighted at low AOD as described in bullets a_{i-ii} .** The algorithm then performs an initial BRF^{Surf} retrieval and begins the aerosol/surface retrieval as presented in the flow chart.

110 The algorithm begins two nested loops, with indices labeled `outerInd` (outer loop) and `innerInd` (inner loop), each starting at 1 and looping through to 2, 3, and 5, for `iterInd` = 1, 2, and 3. The logic of this double nested loop is to better control retrieval of AOD on the coarse LUT grid (see the AOD bins in Table 1) or use Newton's method to **quickly** iterate to the next AOD, which requires a cost function calculation at only two AOD grid points. The downside to Newton's method is that it requires the initial guess to be close to the solution. Therefore, the algorithm *initially* calculates a cost function for every point
115 in our LUT unless `outerInd` is > 1, `innerInd` > 2, or `iterInd` == 3, in which case the AOD should already be significantly more refined. The complexity of this *if* statement requires that the AOD retrieval be in this doubly nested loop. Even within the coarse grid AOD solver, MAGARA still performs an iteration of Newton's method after identifying the best-fitting AOD on our coarse-grid, meaning this algorithm is significantly slower than Newton's method alone. Aside from the algorithmic differences between the coarse LUT AOD solver and using Newton's method alone, the Newton's method solver retrieves ***a***
120 ***weighted daily average AOD***, rather than an AOD using a ± 16 -time bin, 2.5–4.0 hour average, as is performed for the coarse LUT AOD. The reasoning for the daily average retrieval using Newton's method is both speed and accuracy. The coarse grid solver, even with a ± 16 -time bin average, is sufficient to ensure that the Newton's method solver iterates towards a refined solution. After one AOD iteration, the algorithm analytically solves for the best fitting BRF^{Surf} using the just retrieved AOD (shown below) and, if `innerInd` == 1, retrieves the daily averaged fine- and coarse-mode component fractions, as presented in
125 Equation 9. We perform only a single retrieval of daily averaged fine- and coarse-mode component fractions for every loop of `outerInd`, as this prevents the algorithm from bouncing around between local minima. Because this conditional is only activated if `innerInd` == 1 (i.e., for the first iteration of the inner loop) and this daily averaged fine- and coarse-mode component fraction retrieval is located at the end of both nested loops, this retrieval must remain in both loops due to loop structure, as seen in the

130 flow chart. Both nested loops end at this point, and initial retrieval weights that are negatively correlated with retrieved AOD and calculated cost are updated. The final operation of the loop containing iterInd is updating the retrieved surface albedo as the average value of the spectral BRF^{Surf} over all days, times, views (i.e., ABIs), and output aerosol/surface parameters.

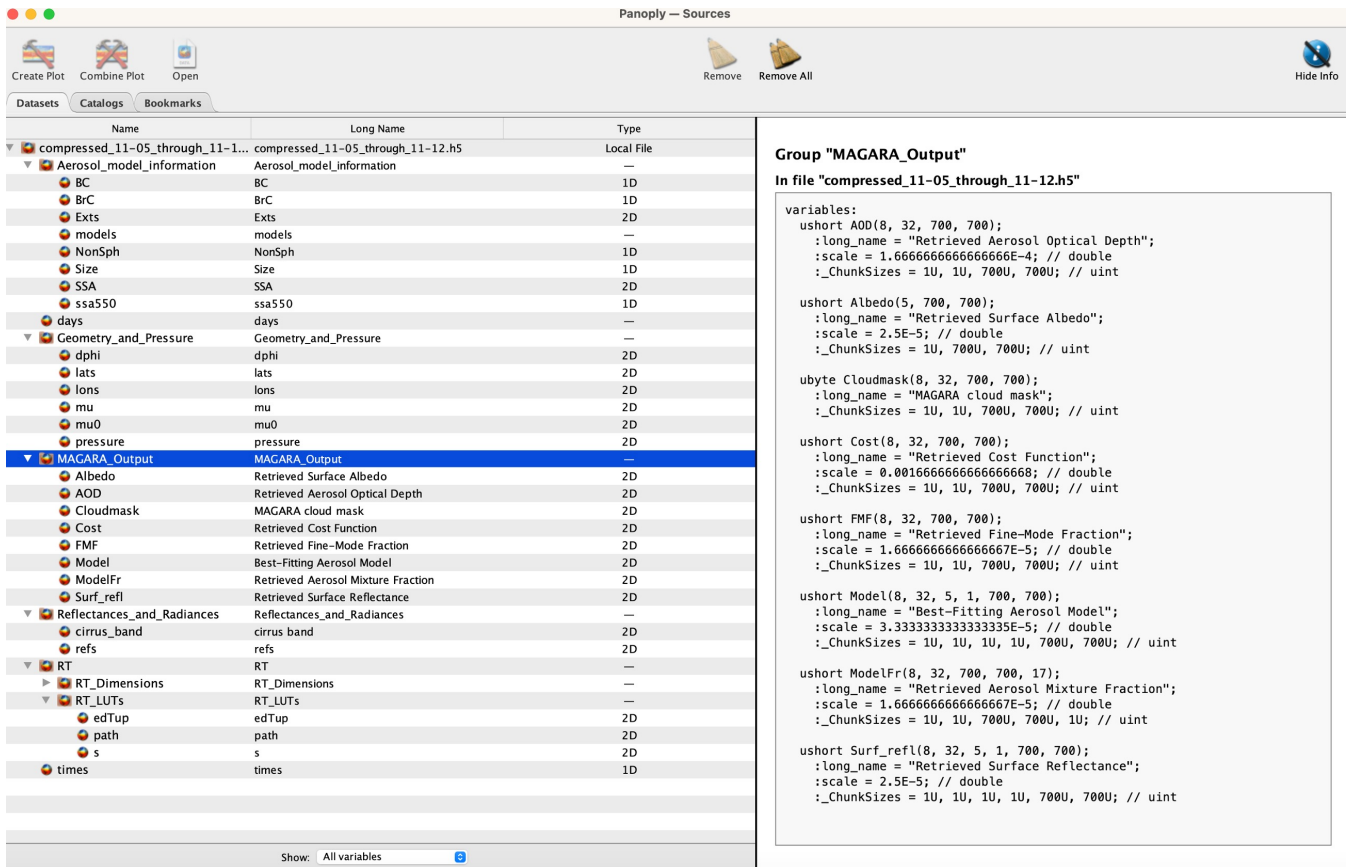
Aerosol model retrieval statistics for Camp Fire case

135 For Camp Fire AOD retrievals, where the **maximum 15-minute** AOD was < 0.1 (very low aerosol loading) for any given day (number of coincidences > 27 million), 70% of retrieved 550-nm extinction comes from fine-mode aerosol. Of the 550-nm AOD, 20% came from our 0.06- μm models (any SSA), 36% came from our 0.12- μm models, and 14% came from our 0.26- μm models. For the 58 million Camp Fire retrievals with $0.3 < \text{maximum daily AOD} < 0.5$, the average FMF is 72%. The contributions to total 550-nm extinction become more skewed towards the 0.12- μm aerosol, even though the 0.12- μm weighting is decreasing because AODs are increasing: 19% for the 0.06- μm models, 44% for the 0.12- μm models, and 9%
140 for the 0.26- μm model. For the 8.4 million Camp Fire retrievals with $0.5 < \text{maximum daily AOD} < 1.0$, average FMF increases to 84%, with the 0.12- μm models now accounting for 50% of total 550-nm extinction and mean retrieved 550-nm SSA equaling 0.93. In this AOD regime, the initial aerosol component fraction weights matter little, as these weights are only used at AOD < 0.5 , so the fact that the fine-mode statistics are converging towards our initial and weighted fine-mode model selection implies that our initial model selection and low-AOD model weighting are likely reasonable.

145

MAGARA Output Description

After all MAGARA output has been generated **for all 50- x 50-pixel regions**, GOESMAGARAMultiCore.py is run again to aggregate output from all 50- x 50-pixel regions into one file. An output HDF-5 file is created, and all output parameters are
150 converted to scaled unsigned 16-bit integers, except latitude and longitude, before compression. In addition to latitude and longitude, we save information on solar/viewing geometry and surface pressure. The radiative transfer LUTs presented in the subsections above are also saved, as well as the TOA BRFs used by MAGARA (called *refs* in the output file), and the BRFs of the cirrus band (band 5; called the *cirrus_band* in the output file). As for direct MAGARA output, we save the following information: 550-nm AOD, 550-nm FMF, cost function (Cost), spectral surface albedo (Albedo), BRF^{Surf} (called *Surf_refl* in
155 the output file), TOA model BRF (Model), and component fraction (ModelFr). An example output file for the Camp Fire, where MAGARA was run for 8 days (with 32 times in each day), as viewed in Panoply (<https://www.giss.nasa.gov/tools/panoply/>), is presented as Figure S1.



160 Figure S1) Example MAGARA output file as viewed in Panoply. This output file represents the output for the Camp Fire.

Optimal spatial averaging window for comparison of MAGARA and NOAA Bias Corrected Product with AERONET


165 For comparison of MAGARA and the NOAA product with AERONET, we must identify the optimal spatial averaging window to use (nearest neighbor, 3x3, 5x5, etc.), as the decreasing noise reduction associated with larger averaging windows will be countered by increasing aerosol spatial variability as window size increases. Additionally, AERONET itself should act to significantly reduce stray clouds if the spatial filter is small enough, as AERONET L1.5 cloud screening is excellent. For the analysis below, where we identify the optimal averaging windows for both MAGARA and the NOAA product, we require a

170 good quality MAGARA and/or NOAA aerosol retrieval centered on the AERONET site, as this will result in the same coincidences with AERONET even as the averaging window increases. For an apples-to-apples comparison of MAGARA


with the GOES-16/17 bias corrected product (looking at the same AERONET coincidences), we require both MAGARA and NOAA to have a quality-assessed retrieval centered on the AERONET site.

175 Comparisons of MAGARA and the NOAA G16/17 bias corrected error statistics for averaging windows of 1 to 51
pixels, in both the x-and-y directions, are presented as four panels in Figure S2. A couple of important points are observed in
Figure S2. Both retrievals show improvement due to spatial averaging, but the improvement for MAGARA is maximized for
3x3 pixel averaging, and even this improvement is very small. However, the NOAA bias corrected product shows remarkable
decreases in both root-mean-squared error (RMSE, ~25% improvement) and median-absolute error (MAE, ~30%
improvement) for the same small change in averaging window. Figure S3, which shows the 2-D histograms of MAGARA vs.
180 AERONET and the NOAA bias corrected product vs AERONET, for the same 1x1 and 3x3 averaging windows, clearly shows
the same phenomenon. Here, it seems very likely that spatial averaging improves the error statistics for the bias corrected
product via noise reduction. The fact that MAGARA's error statistics are quite a bit better than the NOAA bias corrected
product for the case of no spatial averaging, combined with the fact that MAGARA does not see much improvement with
spatial averaging here, provides some evidence that MAGARA's errors may be driven by systematic biases (such as systematic
185 errors in the surface reflectance) rather than random noise, at least for these 3 case studies. It is also likely that MAGARA is
greatly benefitting from both AERONET's and NOAA's cloud/quality masking.

Because both the NOAA product and MAGARA need to be assessed independently to identify the optimal spatial
averaging window, we repeat the analysis of Figure S2 while only requiring that MAGARA **or** the NOAA product have a
quality assessed pixel centered on the AERONET station. This means that Figure S4 is no longer an apples-to-apples
190 comparison of MAGARA and the NOAA bias corrected product, as only GOES-16 was available for the Camp Fire analysis,
and QA NOAA bias corrected G16 aerosol retrievals did not extend up into the Central Valley of California (where smoke
was thickest). Comparing Figure S4 to Figure S2, it appears that both MAGARA and the NOAA bias corrected product benefit
substantially from the cloud/QA filtering provided by each others' algorithms. The upper left panel of Figure S4 also shows a
large drop in RMSE, suggesting that MAGARA may benefit from spatial averaging when aerosol loading is high. The three
195 statistics indicate that the optimal averaging window for MAGARA should be between 5^2 and 15^2 pixels in size. In contrast,
it should be between 19^2 and 21^2 pixels in size for the NOAA bias corrected algorithm in this case.



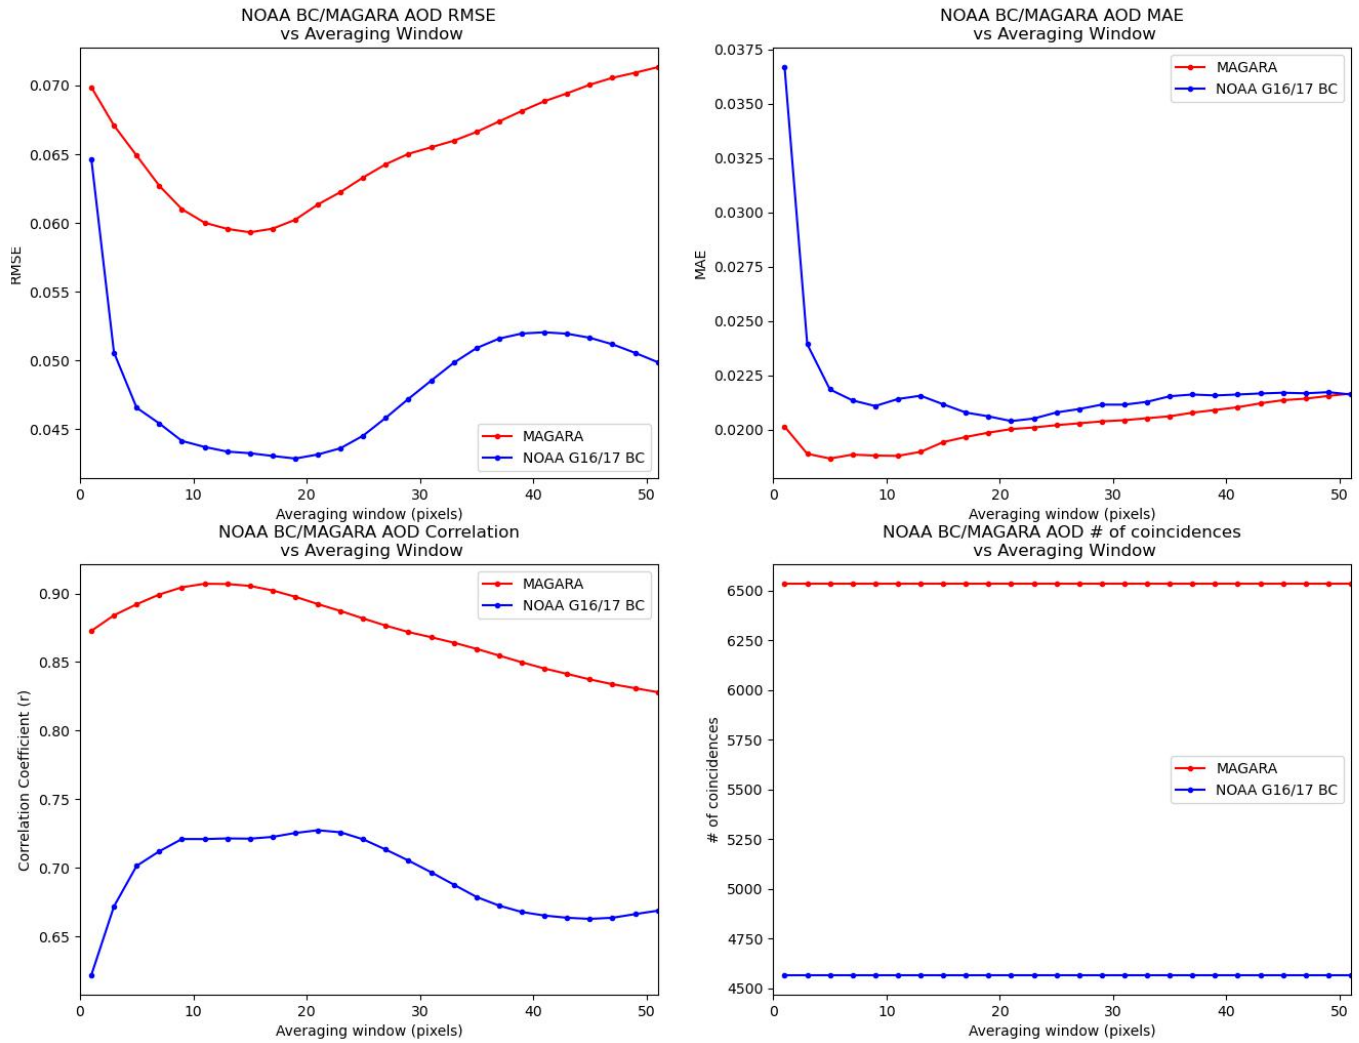
200 **Figure S2) MAGARA and NOAA bias corrected error statistics as a function of averaging distance (box average in the x-and-y directions). Root-mean squared error (RMSE) is presented in the upper left, median-absolute error (MAE) is presented in the upper right, Pearson linear correlation coefficient (r) is presented in the lower left, and the number of coincidences is presented in the lower right panel.**



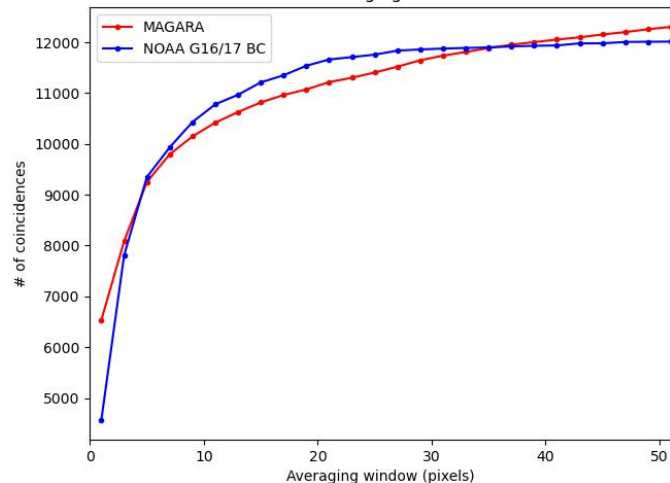
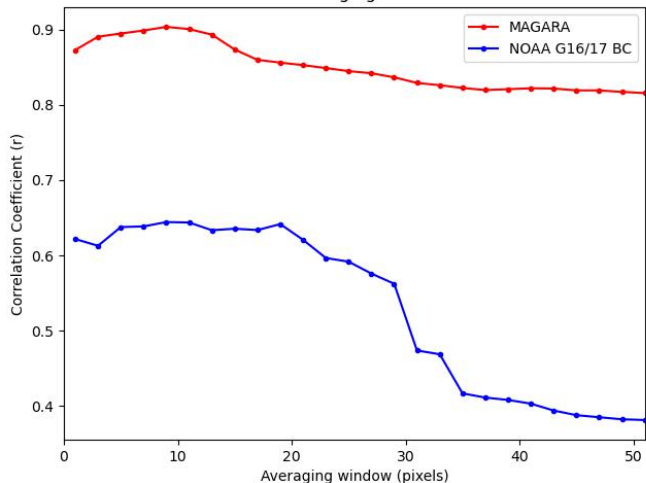
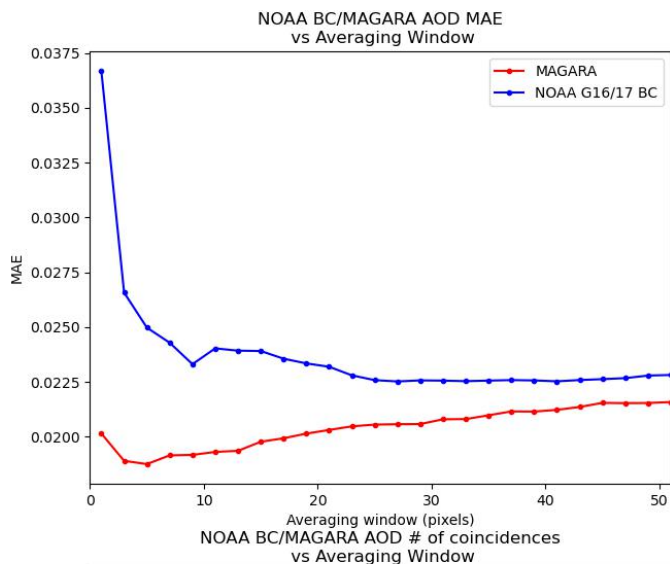
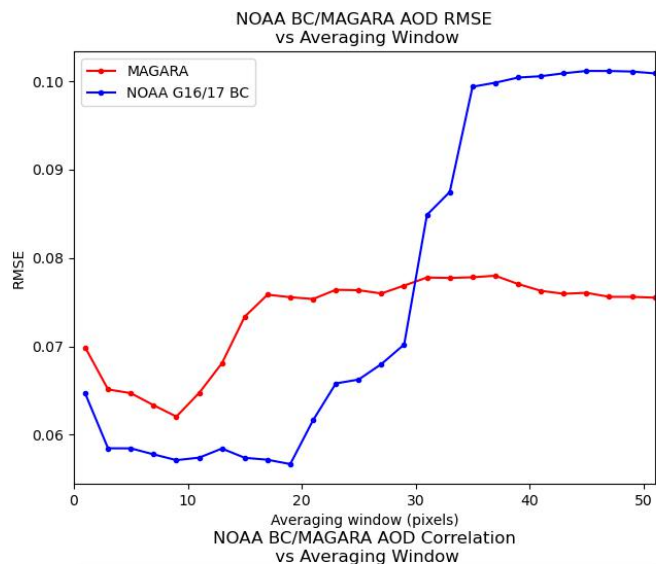
205 Figure S3) 2-D histograms of MAGARA and NOAA bias corrected 550 nm AOD retrievals vs AERONET 550 nm AOD. Upper left
 panel represents the MAGARA vs AERONET 550 nm AOD 2-D histogram using only the MAGARA retrieval centered on the
 AERONET site. Upper right panel represents the NOAA G16/17 bias corrected vs AERONET 550 nm AOD 2-D histogram using
 only the NOAA retrieval centered on the AERONET site. Bottom rows represent the same as the top row, but with 9 x 9
 210 pixel spatial averaging for the satellite retrievals. A standard $\pm (0.03 + 0.15 * AOD)$ uncertainty envelope is provided for reference, with the
 percent meeting that threshold indicated. Statistics for each panel are presented in the title.

We also provide the same statistical analysis as Figure S5, this time not requiring a quality assessed pixel centered on the AERONET station (just one good retrieval for the region). This type of filter is more commonly applied to satellite-based aerosol remote-sensing retrievals, as it can substantially increase the amount of data available for comparison with AERONET.

215 Here, the statistics suggest an averaging window of 3^2 - 9^2 pixels for MAGARA and 9^2 - 19^2 for the NOAA bias corrected product. Considering all of the data presented in Figures S2-S5, a 9×9 spatial averaging window seems ideal for both MAGARA and the NOAA bias corrected product for comparison with AERONET.



220 **Figure S4)** Same as Figure S2, but only requiring a good QA retrieval centered over the AERONET station for either MAGARA or NOAA.



225 **Figure S5)** Same as Figure S2, but without requiring a good QA retrieval directly over the AERONET station for either MAGARA or NOAA.