

# Data Fusion of AIRS and CrIMSS Near Surface Air Temperature

P. Kalmus<sup>1</sup>, H. Nguyen<sup>1</sup>, J. Roman<sup>1</sup>, T. Wang<sup>1</sup>, Q. Yue<sup>1</sup>, Y. Wen<sup>2</sup>, J. Hobbs<sup>1</sup>,  
and A. Braverman<sup>1</sup>

<sup>1</sup>Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 91109

<sup>2</sup>Department of Geography, University of Florida, Gainesville, FL, 32611

## Key Points:

- We have developed a method for fusing any number of two-dimensional remote sensing datasets which estimate the same observable
- We introduce a new daytime and nighttime fused near-surface air temperature product from satellite hyperspectral sounders over CONUS
- The fused product decreases bias and RMSE by 1 K and 25% respectively relative to input datasets, averaged over the domain of the study

---

Corresponding author: Peter Kalmus, [peter.m.kalmus@jpl.nasa.gov](mailto:peter.m.kalmus@jpl.nasa.gov)

**Abstract**

We present a near surface air temperature (NSAT) fused data product over the contiguous United States using data from the Atmospheric Infrared Sounder (AIRS), on the Aqua platform, and the Cross-track Infrared Microwave Sounding Suite (CrIMSS), on the Suomi National Polar-orbiting Partnership (NPP) platform. We create the fused product using a fast python implementation of Spatial Statistical Data Fusion (SSDF) along with weather station data from NOAA’s Integrated Surface Database (ISD) which is used to estimate bias and variance in the input satellite datasets. Our fused NSAT product is produced twice-daily (one daytime and one nighttime estimate per day) and on a 0.25-degree latitude-longitude grid. We provide detailed validation using withheld ISD data and ERA5-Land reanalysis. The fused gridded product has no missing data; has improved accuracy and precision relative to the input satellite datasets, and comparable accuracy and precision to ERA5-Land; and includes accurate uncertainty estimates. Over the domain of our study, the fused product decreases daytime bias magnitude by 1.7 K and 0.5 K, nighttime bias magnitude by 1.5 K and 0.2 K, and overall RMSE by 35% and 15% relative to the AIRS and CrIMSS input datasets, respectively. Our method is computationally fast and generalizable, capable of data fusion from any number of datasets estimating the same quantity. Finally, because our product removes bias, it produces long-term datasets across multi-instrument remote sensing records with improved stationarity for climate trend analysis, even as individual missions and their data records begin and end.

**1 Introduction**

From the point of view of scientific analysis, satellite remote sensing datasets present several challenges. Many satellite remote sensing datasets are released as “Level 2” (L2) products, geophysical quantities retrieved from directly observed radiances. Instantaneous snapshots are obtained at a great number of spatial and temporal fields of regard, and data coverage can be spatially incomplete due to gores (spaces between orbit tracks determined by orbital and sensor geometry), clouds, downlink limitations, or other issues. Satellite retrievals suffer from uncertainties and errors due to information and algorithm limitations, while uncertainty estimates, if reported at all, are not always reliable. Drifts of orbits and spectral channels, and even sudden changes, make the use of data records from satellites challenging in climate studies by causing bias nonstationarity that must be separated from real signals. While L2 satellite data brings invaluable information to scientific analysis, using it appropriately requires significant expertise and involves serious limitations.

Data fusion is the combining of multiple datasets into a single dataset with better properties than any of the individual input datasets (for a recent review, see Ghamisi et al. (2019)). Here, we demonstrate a data fusion method, called Spatial Statistical Data Fusion (SSDF) that addresses each of the above issues (Nguyen et al., 2012, 2014). We use SSDF to create a fused near-surface air temperature (NSAT) product. NSAT is a critical remote sensing product for climate studies of extreme heat, as well as for many science applications areas of great importance to society such as health, agriculture, urban planning, hydrology and water management, ecology and conservation, and fire management. Our SSDF NSAT product combines two remote sensing data products: L2 NSAT from the Atmospheric Infrared Sounder (AIRS) on the Aqua platform, and L2 NSAT from the Cross-track Infrared Microwave Sounding Suite (CrIMSS) on the Suomi National Polar-orbiting Partnership (NPP) platform, which are furthermore created using two independent retrieval algorithms. We also use information content from in situ weather station networks (NOAA’s Integrated Surface Database, or ISD) to determine uncertainties in the two remote sensing datasets which are needed to perform fusion. The fused NSAT product is produced on a twice-daily basis (one

65 daytime and one nighttime estimate per day), and covers the contiguous United States  
66 (CONUS) and adjacent parts of North America.

67 Our fused SSDF NSAT product has the following key advantages over either of  
68 the input remote sensing datasets:

- 69 1. SSDF fills spatial gaps (e.g., due to orbital gores or clouds);
- 70 2. SSDF produces estimates on a regular 0.25-degree spatial grid;
- 71 3. SSDF reduces bias and variance;
- 72 4. SSDF produces uncertainty estimates that characterize the actual error with  
73 more skill than the input datasets;
- 74 5. SSDF improves long-term bias stationarity relative to the input datasets, facil-  
75 itating creation of climate records over changing instrument epochs.

76 The rest of the paper is organized as follows. We first describe the input datasets  
77 and methodology. Then we present the SSDF NSAT product, and the results of  
78 validation against withheld ISD surface station data. We also compare the SSDF fused  
79 NSAT product to the individual input remote sensing datasets, and to ERA5-Land  
80 reanalysis. In the process of validating our SSDF product, we also produce the most  
81 thorough validation study to date of the AIRS V7 and CrIMSS-CLIMCAPS V2 NSAT  
82 products over CONUS. We conclude with a discussion of advantages, limitations, and  
83 potential future work.

## 84 2 Data and methods

85 Performing and evaluating SSDF involves five major steps: (1) Obtaining and  
86 filtering input remote sensing datasets that estimate the same quantity; (2) Matching  
87 the remote sensing datasets to a reference in situ dataset in space and time; (3) Using  
88 these matched data (“matchups”) to characterize the input datasets via estimation of  
89 their bias and variance relative to the reference estimate; (4) Performing the SSDF  
90 calculations; and (5) Validating the results using withheld data from the reference  
91 estimate. The method and the specific datasets used in our NSAT dataset are described  
92 in the following subsections.

### 93 2.1 Satellite NSAT data

94 The input satellite datasets come from two generations of hyperspectral infrared  
95 sounders and retrieval algorithms. The Aqua platform that carries AIRS launched  
96 in 2002 in a sun-synchronous polar orbit, with equator crossing times of 1:30 P.M.  
97 and 1:30 A.M. for ascending (south to north) and descending (north to south) nodes,  
98 respectively. AIRS is an infrared grating spectrometer with 2378 channels, spanning  
99 3.7 to 15.4  $\mu\text{m}$  (Chahine et al., 2006). Power to critical channels of the Aqua satellite’s  
100 Advanced Microwave Sounding Unit (AMSU)-A2 was lost in September 2016 (Yue  
101 et al., 2017), which was used to complement the AIRS instrument in atmospheric  
102 temperature and moisture profile retrievals.

103 We use the AIRS version 7 L2 “infrared-only” temperature retrieval algorithm  
104 (Susskind et al., 2014). This retrieval uses the Stochastic Cloud Clearing Neural Net-  
105 work (SCCNN) which is trained to ECMWF fields (Blackwell, 2005) as a first guess,  
106 then refines to a final estimate. It also uses information from the satellite’s other mi-  
107 crowave sounder, AMSU-A1 (Yue et al., 2020). The retrieval uncertainty is estimated  
108 via a regression model using eleven retrieval diagnostic quantities as predictors; the re-  
109 gression coefficients are trained on two days of retrievals (9/29/04 and 2/24/07) using  
110 ECMWF 3-hour forecasts as a reference dataset (Susskind et al., 2014; Thrastarson et  
111 al., 2020). Each individual retrieval has a nominal horizontal resolution of 45 km, and

112 each swath contains 30 retrievals across its width and 45 along track. The product is  
 113 organized nominally in 240 “orbital granules” per day (AIRS Project, 2020).

114 The Cross-track Infrared Sounder (CrIS) and the Advanced Technology Mi-  
 115 crowave Sounder (ATMS) instruments launched onboard the NPP platform in 2012.  
 116 NPP is in the same orbital plane as Aqua, but at a higher altitude (824 km as opposed  
 117 to 705 km), with equator crossing times also nominally of 1:30 P.M. and 1:30 A.M.  
 118 for ascending and descending nodes, respectively. We use the Community Long-term  
 119 Infrared Microwave Coupled Atmospheric Product System (CLIMCAPS) Version 2  
 120 L2 temperature retrieval, which uses an optimal estimation methodology with a first  
 121 guess from the Modern-Era Retrospective Analysis for Research and Applications ver-  
 122 sion 2 (MERRA2) (N. Smith & Barnet, 2020), and information from both instruments.  
 123 CLIMCAPS uncertainty is estimated and propagated sequentially via error covariance  
 124 matrices in stages (N. Smith & Barnet, 2019). CLIMCAPS produces a combined  
 125 infrared and microwave retrieval at two spectral resolutions: Nominal Spectral Res-  
 126 olution (NSR) and Full Spectral Resolution (FSR). We use the CLIMCAPS-SNPP  
 127 NSR product to create our SSDF product. In what follows, we refer to this product  
 128 as “CrIMSS-CLIMCAPS” or simply “CrIMSS.”

129 For both instruments, NSAT is obtained from the vertically-resolved temperature  
 130 profile (100 pressure levels) by interpolation or extrapolation with pressure to the  
 131 surface pressure for each field of regard (Olsen et al., 2017). The profile temperatures  
 132 immediately above and below the surface are used for the interpolation, unless the  
 133 level above is within 5 hPa of the surface pressure. In that case, the two levels above  
 134 the surface are used. We ingest only L2 NSAT retrievals from AIRS V7 IR-only and  
 135 CrIMSS-CLIMCAPS products with data quality flags ‘good’ or ‘best’ in our data  
 136 fusion procedure.

## 137 2.2 In situ NSAT data

138 The National Oceanic and Atmosphere Administration (NOAA) Integrated Sur-  
 139 face Database (ISD) is a global database of near-surface meteorological observations  
 140 compiled from over a hundred sources (A. Smith et al., 2011). The record extends back  
 141 to the 1950s, although new stations have been added on a continual basis as available,  
 142 improving coverage over time. Today ISD consists of more than 35,000 surface weather  
 143 stations globally, 14,000 of which remain active. Figure 1 shows the spatial coverage  
 144 of ISD stations in North America.

145 We use sub-hourly 2 m NSAT measurements gathered from over 7000 stations in  
 146 North America as our reference dataset, for bias and variance estimation and for val-  
 147 idation. Naturally ventilated screened surface station air temperature measurements  
 148 are accurate to  $\pm 0.1^\circ\text{C}$  in most circumstances (Harrison & Burt, 2021). ISD data come  
 149 with a set of ten data quality flags, indicating various problems and levels of quality.  
 150 We only use ISD data flagged as highest quality, i.e., data must be flagged with either  
 151 1 (‘Passed all quality control checks’) or 5 (‘Passed all quality control checks, data  
 152 originate from an NCEI data source’).

## 153 2.3 Reanalysis

154 We also compare the SSDF NSAT results to European Centre for Medium-Range  
 155 Weather Forecasts (ECMWF) Reanalysis 5 (ERA5)-Land reanalysis data. The ERA5  
 156 is the fifth-generation global atmospheric reanalysis from ECMWF, replacing the ERA-  
 157 Interim reanalysis which stopped being produced on August 31st, 2019. In addition,  
 158 newly reprocessed datasets along with recent instruments have been assimilated into  
 159 the ERA5 that could not be ingested into the ERA-Interim (Hennermann & Berris-  
 160 ford, 2019). We note that some AIRS spectral channels under clear conditions are

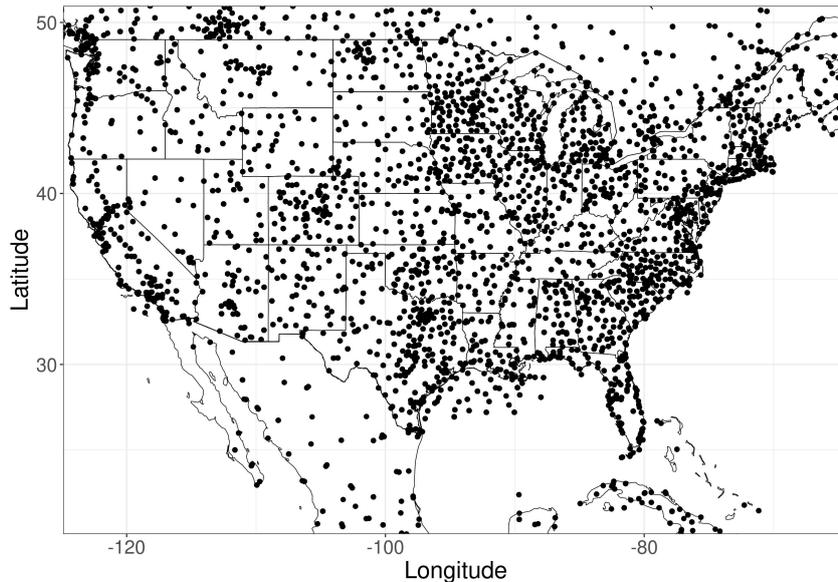


Figure 1: Spatial coverage of the ISD stations over North America.

161 incorporated into ECMWF reanalysis (McNally et al., 2006), but that ISD data are  
 162 not.

163 We use hourly ERA5-Land output which is a high-resolution version ( $\sim 9$  km)  
 164 of the land component of the ERA5 climate reanalysis. ERA5-Land was chosen over  
 165 the full ERA5 reanalysis for its finer spatial resolution of  $0.1 \times 0.1^\circ$ . Hourly 2 m air  
 166 temperature output was selected for our comparison.

#### 167 **2.4 Bias and variance estimation**

168 Biases and variances of input data sources are crucial for proper data fusion. The  
 169 SSDF methodology assumes input data are unbiased, and weights them by the inverse  
 170 of their respective variances. This minimizes output errors of the fused estimates.  
 171 Therefore, data are bias-corrected before SSDF ingestion, and the quality of the final  
 172 fused product is largely determined by the quality of uncertainty estimates for the  
 173 inputs.

174 To estimate bias and variance for satellite footprints, we create an ensemble  
 175 of “matchups”: matched pairs of satellite and ISD station estimates that are close in  
 176 space and time. For a given period, the matchups are sorted into 240 km ( $\sim$ two-degree)  
 177 diameter hexagonal spatial bins based on satellite footprint location, with three-day  
 178 time bins (day of interest, along with preceding and following days). This binning  
 179 is the basis for quantifying bias and variance for all satellite footprints in a given  
 180 space-time cell. We experimented with using longer and shorter time bins to explore  
 181 the trade off between sample size and capturing rapid changes in conditions affecting  
 182 retrieval bias, and found that the three-day bin delivered the lowest average biases and  
 183 variances over CONUS. Before starting, we randomly selected 1% of the ISD matchups  
 184 to withhold for validation. We chose a relatively small amount to withhold in order to  
 185 maximize the information content for the SSDF product. In this subsection, the term  
 186 “ISD” refers to the non-withheld ISD data.

187 To obtain the matchups we apply the following steps.

- 188 1. Given an ISD observation at location  $\mathbf{s}$  and time  $t^I(\mathbf{s})$ , select the AIRS granule  
189 (1 of 240) with the closest time to  $t^I(\mathbf{s})$ .
- 190 2. Within this granule, select all L2 retrievals within 100 km of  $\mathbf{s}$  and 1 hour of  
191  $t^I(\mathbf{s})$ .
- 192 3. If Step 2 results in more than 1 retrieval, select the one closest in spatial distance.

193 Note that the Steps in 1-3 will result in a one-to-one match between an ISD  
194 observation and a single AIRS footprint. Some ISD observations may have no corre-  
195 sponding AIRS match, in which case we return a null result. We next tessellate a fixed  
196 hexagonal spatial grid over CONUS and find the biases and variances using matchups  
197 aggregated over 3 days within each grid cell. That is,

- I. To compute a bias on day  $d$  and mode  $j$  (day or night) and in hexagonal grid  
198 cell  $i$ , we find the set of all valid (i.e., non-null) AIRS-ISD matchups from Steps  
199 1-3 above such that,  
200  
201 (a) the AIRS data come from mode  $j$ ,  
202 (b) the AIRS footprint belongs within the grid cell  $i$ ,  
203 (c) the ISD date is in  $(d - 1, d, d + 1)$ .
- 204 II. The bias and variance for day  $d$ , mode  $j$ , and grid cell  $i$  are then computed  
205 using the set of paired ISD-AIRS matchups.

206 Bias and variance estimation for CrIMSS follows the same procedure. For bias  
207 correction, given an instrument observation at location  $\mathbf{s}$  on day  $d$  and mode  $j$ , we  
208 compute the corresponding bias within the grid cell which contains  $\mathbf{s}$  for day  $d$  and  
209 mode  $j$ , and we subtract it from the instrument’s NSAT value. For more detail on the  
210 bias and variance estimation process, please refer to the Appendix.

## 211 2.5 Data fusion methodology

212 In this section we review the framework of Spatial Statistical Data Fusion (SSDF;  
213 Nguyen et al., 2012) on two satellite NSAT datasets. Remote sensing data in general  
214 are heterogeneous. By this we mean that they may have different footprints, mea-  
215 surement error characteristics, and sampling patterns. We account for this by using a  
216 spatial statistical model that captures the spatial dependence between the true quan-  
217 tity of interest at a particular location and the observations from all data sources. In  
218 particular, the issue of different footprint sizes and shapes is known as a *change-of-*  
219 *support problem* (e.g. Gotway & Young, 2002), and we will address this using SSDF  
220 as described in Nguyen et al. (2012).

221 Consider a discretized domain where  $\{Y(\mathbf{s}) : \mathbf{s} \in D\}$  is a hidden, real-valued  
222 spatial observable. The domain of interest is  $\cup\{A_i \subset \mathcal{R}^d : i = 1, \dots, N_D\}$ , which is  
223 made up of  $N_D$  fine-scale, non-overlapping, areal regions  $\{A_i\}$  with locations  $D \equiv$   
224  $\{\mathbf{p}_i \in A_i : i = 1, \dots, N_D\}$ . Nguyen et al. (2012) call these fine-scale regions Basic  
225 Areal Units (BAUs), and they represent the smallest resolution at which we will make  
226 estimates with the model.

227 For a given day and mode ( $d$  and  $j$  using the notation of the previous subsection),  
228 denote the vector of NSAT data at all locations by  $\mathbf{Z}^k$ , where  $k = 1$  for AIRS and  
229  $k = 2$  for CrIMSS:

$$\mathbf{Z}^k = (Z^k(B_{k1}), Z^k(B_{k2}), \dots, Z^k(B_{kN_k}))',$$

where  $\mathbf{Z}^k$  is  $N_k$ -dimensional,  $B_{kq}$  is the  $q$ -th footprint from the  $k$ -th dataset and is  
made up of BAUs with locations indexed by  $D \cap B_{kq}$ . We assume that data observed  
at an arbitrary areal region  $B$  follow the “data model” in which the true observable is

averaged over the areal region plus an independent error term. That is,

$$Z^k(B) = \frac{1}{|D \cap B|} \left\{ \sum_{\mathbf{s} \in D \cap B} Y(\mathbf{s}) \right\} + \epsilon^k(B); \quad B \subset \mathfrak{R}^d. \quad (1)$$

230 where  $Y(\cdot)$  is a geophysical observable (here, NSAT) that is common to both datasets,  
 231 and  $\epsilon^k(\cdot)$  is an independent but non-identically distributed Gaussian random variable.  
 232 That is, we assume that the  $q$ -th error in the  $k$ -th dataset is distributed as  $\epsilon_q^k \sim$   
 233  $N(b_q^k, v_q^k)$ . In general,  $b_q^k$  is not zero, however, in our case  $b_q^k$  is assumed to be zero  
 234 because we performed bias correction as described in the previous subsection, and  
 235  $v_q^k$  are calculated from the hexagonal-cell-specific mean and variance estimates (see  
 236 Appendix for details).

237 Our fused estimate for a region centered at location  $B_0$  is a linear combination  
 238 of  $\mathbf{Z}_1$  and  $\mathbf{Z}_2$ . That is,

$$\hat{Y}(B_0) = \mathbf{a}'_1 \mathbf{Z}_1 + \mathbf{a}'_2 \mathbf{Z}_2, \quad (2)$$

239 where  $\mathbf{a}_1$  and  $\mathbf{a}_2$  are  $N_1$  and  $N_2$  dimensional vectors, respectively. These vectors are  
 240 unknown and are estimated in a way that minimizes the expected squared error relative  
 241 to the true observable. That is, we choose  $\mathbf{a}_1$  and  $\mathbf{a}_2$  to minimize,

$$\begin{aligned} E((Y(B_0) - \hat{Y}(B_0))^2) &= \text{Var}(Y(B_0) - \mathbf{a}'_1 \mathbf{Z}_1 - \mathbf{a}'_2 \mathbf{Z}_2) \\ &= \text{Var}(Y(B_0)) - 2\mathbf{a}'_1 \text{Cov}(\mathbf{Z}_1, Y(B_0)) \\ &\quad - 2\mathbf{a}'_2 \text{Cov}(\mathbf{Z}_2, Y(B_0)) \\ &\quad - 2\mathbf{a}'_1 \text{Cov}(\mathbf{Z}_1, \mathbf{Z}_2) \mathbf{a}_2 \\ &\quad + \mathbf{a}'_1 \text{Var}(\mathbf{Z}_1) \mathbf{a}_1 + \mathbf{a}'_2 \text{Var}(\mathbf{Z}_2) \mathbf{a}_2 \end{aligned}$$

242 subject to the unbiasedness constraint that the elements of  $\mathbf{a}_1$  and  $\mathbf{a}_2$  add up to 1.  
 243 That is,

$$\mathbf{1} = \mathbf{a}'_1 \mathbf{1}_{N_1} + \mathbf{a}'_2 \mathbf{1}_{N_2}, \quad (3)$$

244 where  $\mathbf{1}_{N_k}$  is an  $N_k$ -dimensional vector of ones. The solution to the minimization  
 245 problem in (3) can be found via the method of Lagrange multipliers; but it requires  
 246 knowledge of the spatial covariance structure  $C(B_i, B_j)$ , which can be expanded in  
 247 terms of the BAU covariances:

$$C(B_i, B_j) = \frac{1}{|D \cap B_i| |D \cap B_j|} \sum_{\mathbf{u} \in D \cap B_i} \sum_{\mathbf{v} \in D \cap B_j} C(\mathbf{u}, \mathbf{v}). \quad (4)$$

Typically, the covariance structure in kriging-based approaches is estimated from  
 the data, but the formulation in Equation 4 makes estimation intractable for non-linear  
 covariance classes. We make use of the Spatial Mixed Effects model (SME; Cressie  
 & Johannesson, 2008), which assumes that the true observable, here NSAT, can be  
 written as the linear mixed model,

$$Y(\mathbf{s}) = \mathbf{t}(\mathbf{s})' \boldsymbol{\alpha} + \mathbf{S}(\mathbf{s})' \boldsymbol{\eta} + \xi(\mathbf{s}). \quad (5)$$

248 where  $\mathbf{t}(\cdot) \equiv (t_1(\cdot), \dots, t_p(\cdot))'$  is a vector of  $p$  known covariates, such as geographical  
 249 coordinates or other physical variables. The vector of linear coefficients,  $\boldsymbol{\alpha}$ , is unknown  
 250 and will be estimated from the data. The middle term captures the spatial dependence  
 251 as the product of an  $r$ -dimensional vector of known spatial basis functions,  $\mathbf{S}(\mathbf{s})$ , and an  
 252  $r$ -dimensional Gaussian random variable,  $\boldsymbol{\eta}$ . Here, we assume that with  $\boldsymbol{\eta} \sim N(\mathbf{0}, \mathbf{K})$ .  
 253 Similar to the implementation in Nguyen et al. (2012), we implement these using  
 254 multi-resolution bisquare basis functions centered at different resolutions of the Inverse  
 255 Snyder Equal-Area Projection Aperture 3 Hexagon (ISEA3H) type within the Discrete

256 Global Grid (DGGRID) software (specifically, resolutions 2, 3, and 5 of ISEA3H, for  
 257 details see Sahr, 2019). The last term,  $\xi(\cdot)$ , describes the BAU-scale variability of the  
 258 process. We assume that  $\xi(\cdot)$  is an independent Gaussian process with mean zero and  
 259 variance  $\sigma_\xi^2$ .

260 The SME model in Equation 5 has useful change-of-support properties, which  
 261 makes computation of the spatial covariance function straightforward. In particular,  
 262 Nguyen et al. (2012) shows that

$$\text{cov}(Z(B_i), Z(B_j)) = \mathbf{S}(B_i)' \mathbf{K} \mathbf{S}(B_j) + \sigma_\xi^2 \frac{|D \cap B_i \cap B_j|}{|D \cap B_i| |D \cap B_j|} + v_i^k I(i = j), \quad (6)$$

where

$$\mathbf{S}(B_i) \equiv \frac{1}{|D \cap B_i|} \sum_{\mathbf{u} \in D \cap B_i} \mathbf{S}(\mathbf{u}).$$

263 Notice that Equation 6 allows us to express the covariance between spatial averages  
 264 *explicitly* in terms of the spatial dependence parameter  $\mathbf{K}$ . This allows for straightfor-  
 265 ward estimation of it from footprint data.

266 Another advantage of the SME model is its scalability. For a general covariance  
 267 structure, solving for  $\mathbf{a}_1$  and  $\mathbf{a}_2$  requires inverting a  $(N_1 + N_2) \times (N_1 + N_2)$  covariance  
 268 matrix, which has computational complexity  $O((N_1 + N_2)^3)$ . For large datasets such as  
 269 AIRS and CrIMSS where the data size is on the order of tens of thousands, this matrix  
 270 inversion is computationally infeasible. However, the model in Equation 5 implies the  
 271 following full covariance matrix:

$$\begin{aligned} \boldsymbol{\Sigma} &\equiv \text{var}((\mathbf{Z}^{1'}, \mathbf{Z}^{2'})') \\ &= \mathbf{S}' \mathbf{K} \mathbf{S} + \mathbf{U}, \end{aligned}$$

272 where  $\mathbf{S}$  is a matrix constructed by appending the spatial function  $\mathbf{S}(\cdot)$  over all the  
 273 footprints in both datasets,  $\mathbf{U}$  is the *sparse* covariance matrix for the fine-scale pro-  
 274 cesses  $\xi(\cdot)$ , and the measurement-error processes  $\epsilon^k(\cdot)$  at the given data locations (for  
 275 more details, see Equation 4 of Nguyen et al., 2012). Using the Sherman-Morrison-  
 276 Woodbury formula (e.g., Henderson & Searle, 1981), the matrix inverse is given by,

$$\boldsymbol{\Sigma}^{-1} = \mathbf{U}^{-1} - \mathbf{U}^{-1} \mathbf{S}' (\mathbf{K}^{-1} + \mathbf{S} \mathbf{U}^{-1} \mathbf{S}')^{-1} \mathbf{S} \mathbf{U}^{-1},$$

277 Note that the inversion above, and hence the calculation of the coefficients  $\mathbf{a}_1$  and  
 278  $\mathbf{a}_2$  for the fused estimate, is very fast because it only requires inversion of the *sparse*  
 279  $(N_1 + N_2) \times (N_1 + N_2)$  matrix  $\mathbf{U}$ , which is typically very sparse, and inversion of  $\mathbf{K}$   
 280 and  $(\mathbf{K}^{-1} + \mathbf{S}' \mathbf{U}^{-1} \mathbf{S})$ , both of which are  $r \times r$  matrices ( $r \ll N_1 + N_2$ ).

281 The methodology described in this section is a scalable variant of Gaussian pro-  
 282 cess prediction (Cressie, 2015). It has been applied to fusion of total column CO<sub>2</sub>  
 283 concentration (XCO<sub>2</sub>) from AIRS and OCO-2 and aerosol optical depth from MISR  
 284 and MODIS (Nguyen et al., 2012, 2014). Hammerling et al. (2012) used another  
 285 variant called local kriging to produce Level 3 estimates of XCO<sub>2</sub> from the GOSAT  
 286 instrument.

287 There are two important advantages of Gaussian process prediction over other  
 288 approaches currently in use such as binning or nearest neighbor interpolation. First,  
 289 our fused estimates are *best linear unbiased estimates*. That is, the standard errors are  
 290 guaranteed to be the smallest possible because the estimates are derived through an  
 291 algorithm that minimizes errors relative to the unknown true process. Such estimates  
 292 are called best linear unbiased estimates, and are optimal in that sense. It is easily  
 293 shown that within the class of linear estimators, this method produces the smallest  
 294 prediction errors. The second advantage is that SSDF provides a statistically princi-  
 295 pled method for estimating uncertainties (that is,  $\text{Var}(\hat{Y}(B_0) - Y(B_0))$ ). Quantifying  
 296 and minimizing uncertainties in this manner is crucial for creating data products for  
 297 scientific analyses that involve making inferences about geophysical observables.

298

## 2.6 Dataset preparation for validation

299

300

301

302

303

304

305

306

307

308

309

310

We validate our SSDF product using a randomly chosen reserved 1% of the ISD dataset. We match up SSDF, AIRS, CrIMSS, and ERA5 estimates to withheld ISD data using a 100 km and 1 hour matchup criterion (see Section 2.4 for more detail). This matchup procedure generates multiple paired datasets: ISD-AIRS, ISD-CrIMSS, ISD-SSDF, and ISD-ERA5. These matchup datasets might differ in their coverage; for instance, an SSDF estimate might be matched to an ISD observation at a location where there are no nearby AIRS or CrIMSS estimates. Therefore, to mitigate the effect of biases due to differing spatial and temporal coverage in these matchup pairs, we also require that SSDF estimates are also close to (within the same matchup distance and time) at least one datum from the comparison dataset. This allows us to compare, for example, AIRS and SSDF(AIRS) datasets which have the same number of samples, all of which are collocated in space and time within the matchup criterion.

311

312

313

The choices of a 1% test ISD dataset and this matchup scheme results in over 4000 AIRS-SSDF sample pairs and over 13,000 CrIMSS-SSDF sample pairs for 2013, a typical year.

314

## 3 Results

315

### 3.1 SSDF product overview

316

317

318

319

320

321

322

323

324

We produced fused NSAT using two satellite input datasets over North America between 25 N and 50 N, from November 28 2012, when CrIMSS-CLIMCAPS first becomes available, through the end of 2020. During this time period, there were 34 days and 36 nights with no AIRS data (approximately half of which occurred in 2020), and 24 days and 28 nights with no CrIMSS-CLIMCAPS data. In the cases with only one input satellite dataset, the SSDF product is created from only the single dataset, thus creating a continuous record. There was one day/night period (November 7, 2020) without either AIRS or CrIMSS-CLIMCAPS data; we did not create SSDF product for this day.

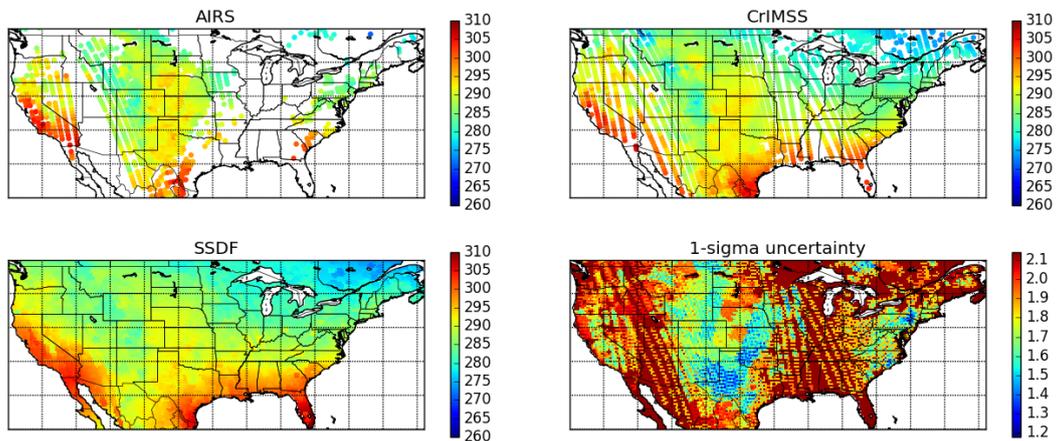


Figure 2: Sample data fusion satellite NSAT inputs, SSDF NSAT results, and uncertainty estimates for 2015 October 31, day. The top two plots show maps of the input satellite NSAT data ingested into the SSDF product, with AIRS on the left and CrIMSS on the right. The bottom-left plot shows the SSDF fusion results. The bottom-right plot shows the uncertainty estimates on the SSDF fusion results at the 1-sigma level. All units are Kelvin.

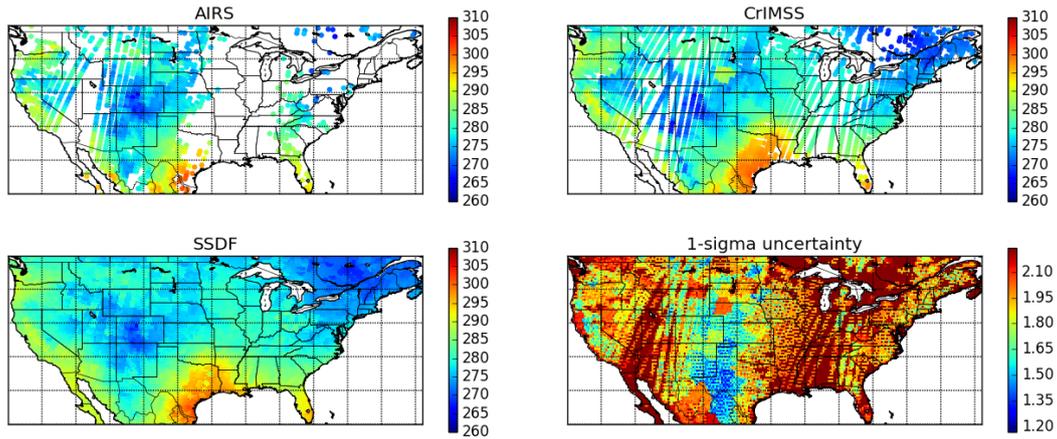


Figure 3: Same as Figure 2 but for night. All units are degrees K.

325 Figures 2 and 3 provide maps representing one arbitrarily chosen day and night  
 326 of the SSDF product. For both the day and night cases, the top two plots show maps  
 327 of the input satellite data ingested into the SSDF product, with AIRS on the left and  
 328 CrIMSS on the right; the bottom left plot shows the SSDF fusion results; and the  
 329 bottom right plot shows the uncertainty estimates on the SSDF fusion results at the  
 330 1-sigma level. These sample maps demonstrate how our SSDF method fills in missing  
 331 data in the input datasets by exploiting spatial correlations to provide a complete  
 332 gap-filled, gridded product. They also provide a first look at the SSDF uncertainty  
 333 estimates. Note that the estimated uncertainties are higher in regions that contain no  
 334 observations, contain observations from only a single input dataset, or in which the  
 335 two input datasets have relatively poor agreement.

### 326 3.2 Bias, standard deviation, and RMSE comparison

337 We now turn to validation against withheld ISD reference data to quantify im-  
 338 provement in the SSDF products. We examine bias, standard deviation, and RMSE,  
 339 calculated from the withheld matchups, of AIRS, CrIMSS, ERA5-Land, and the cor-  
 340 responding matched SSDF data. In what follows, analyze daytime and nighttime  
 341 separately, as daytime and nighttime biases differ significantly.

342 We first show maps of bias, RMSE, and standard deviation relative to the 1% of  
 343 withheld (testing-only) ISD reference data, based on the matchups aggregated into the  
 344 hexagonal bins. Figure 4 shows maps of bias (retrieval - ISD) for AIRS, CrIMSS, and  
 345 SSDF, for the 2013-2020 period in total, and for day-only and night-only. Individual  
 346 bias estimates for retrieval-ISM pairs are aggregated into 2-degree hexagonal cells.

347 Overall, in the mean over CONUS and over the entire time period, SSDF provides  
 348 a reduction in the magnitude of daytime bias of 1.7 K and 0.5 K relative to AIRS and  
 349 CrIMMS, respectively. At night, SSDF is essentially unbiased in the mean over the  
 350 domain and provides a reduction in the magnitude of bias of 1.5 K and 0.2 K relative  
 351 to AIRS and CrIMMS, respectively.

352 AIRS shows a strong cold bias in daytime over the mountainous West, which  
 353 is also present in CrIMSS, although less severe. AIRS shows a near-constant warm  
 354 bias over the entire Eastern CONUS at night, while CrIMSS shows a sharp warm bias  
 355 over small regions of the mountainous West at night. SSDF corrects all of these biases  
 356 (through the bias-correction procedure described above) and produces estimates with  
 357 lower biases than either of its input satellite data sets over the domain.

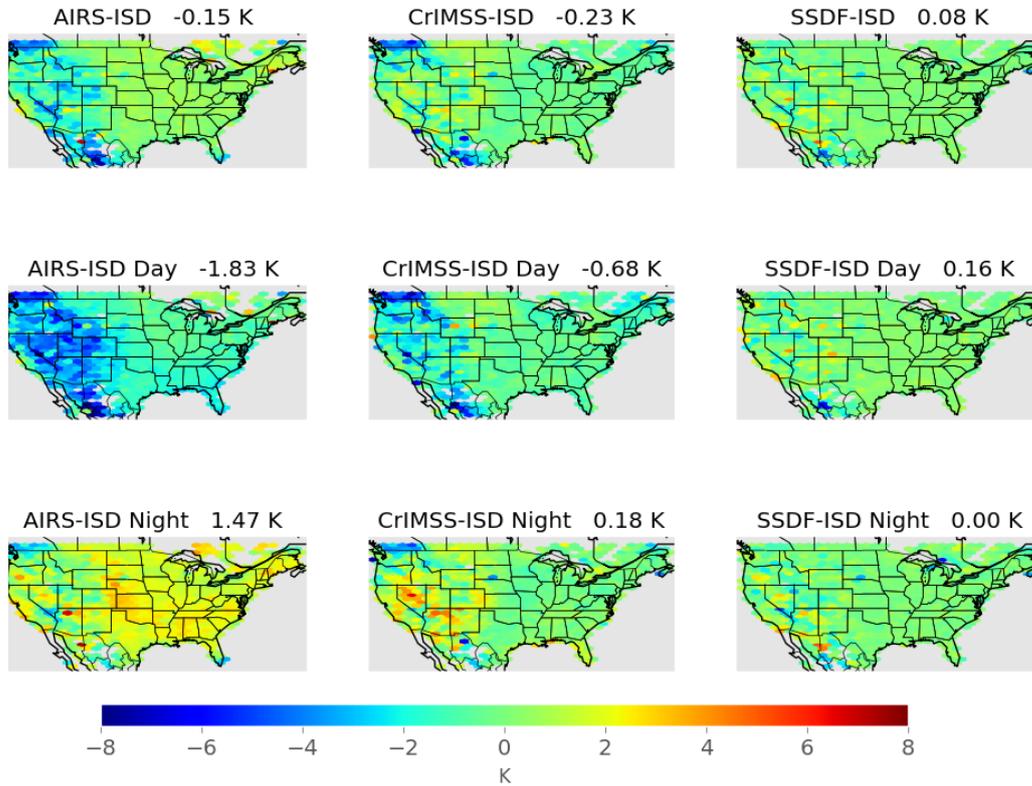


Figure 4: Maps of bias (retrieval - ISD) over the product period of 2013-2020, created against the withheld ISD test data, for AIRS (first column), CrIMSS-CLIMCAPS (second column) and SSDF (third column), for both day and night together (top row), for day only (second row) and for night only (third row). Individual bias estimates for retrieval-  
ISD matchup pairs are aggregated over 2-degree hexagonal cells. The mean bias over CONUS for the entire time period is shown in the title for each map.

358 Figures 5 and 6 show maps of standard deviation and RMSE for AIRS, CrIMSS  
 359 and SSDF, for the 2013-2020 period, and for daytime only and nighttime only. Stan-  
 360 dard deviation and RMSE tell a similar story to that of bias. Overall, in the mean  
 361 over CONUS and over the entire time period, SSDF provides a reduction in RMSE of  
 362 35% and 15% compared to AIRS and CrIMSS, respectively.

363 CrIMSS has high RMSE over the mountainous West in both day and night,  
 364 but low RMSE over the eastern two-thirds of the continent. Similarly, AIRS has  
 365 relatively high RMSE over the entire domain, but especially over the mountainous  
 366 West. Mountainous regions pose particular challenges for remote sensing of surface  
 367 quantities, and of NSAT in particular, which can vary greatly depending on e.g., north-  
 368 facing versus south-facing mountain surfaces. Furthermore, variations in topographic  
 369 features between ISD stations and their matched remote sensing retrievals can lead  
 370 to random errors, increasing RMSE and variance estimates. However, SSDF NSAT  
 371 shows a clear decrease in bias over all regions, including in the mountainous western  
 372 CONUS, although there is potential for improvement in the SSDF product over the  
 373 West.

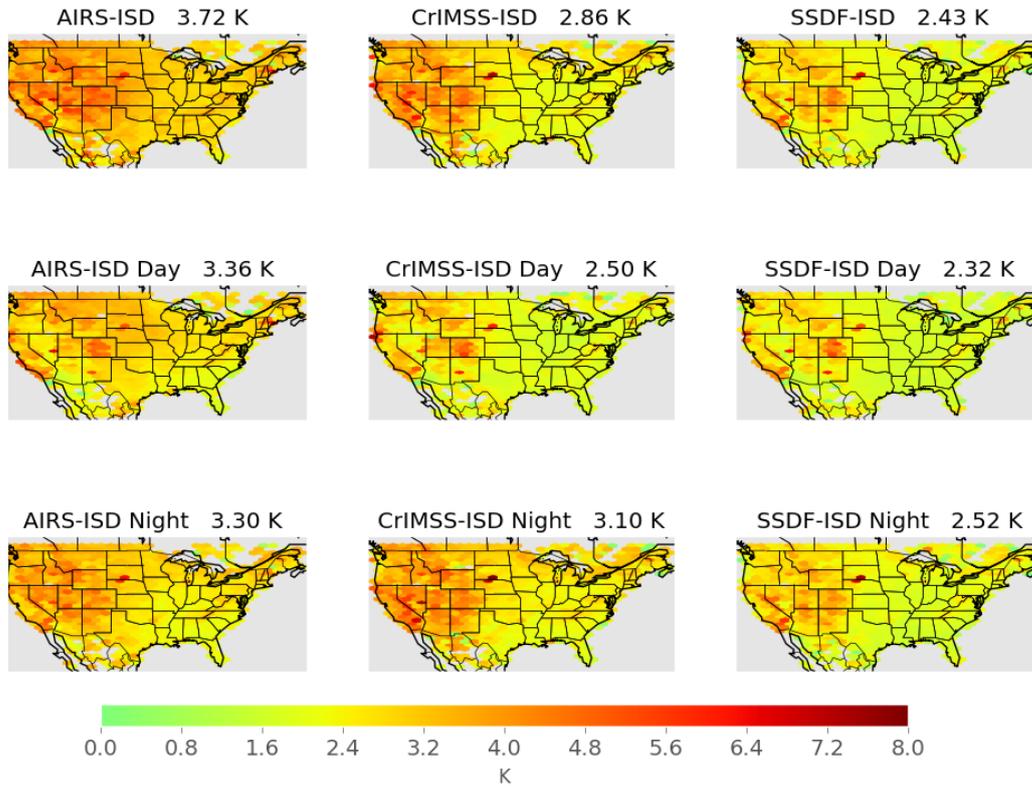


Figure 5: Standard deviation maps. The nine panels are similar to those in Figure 4 but for standard deviation.

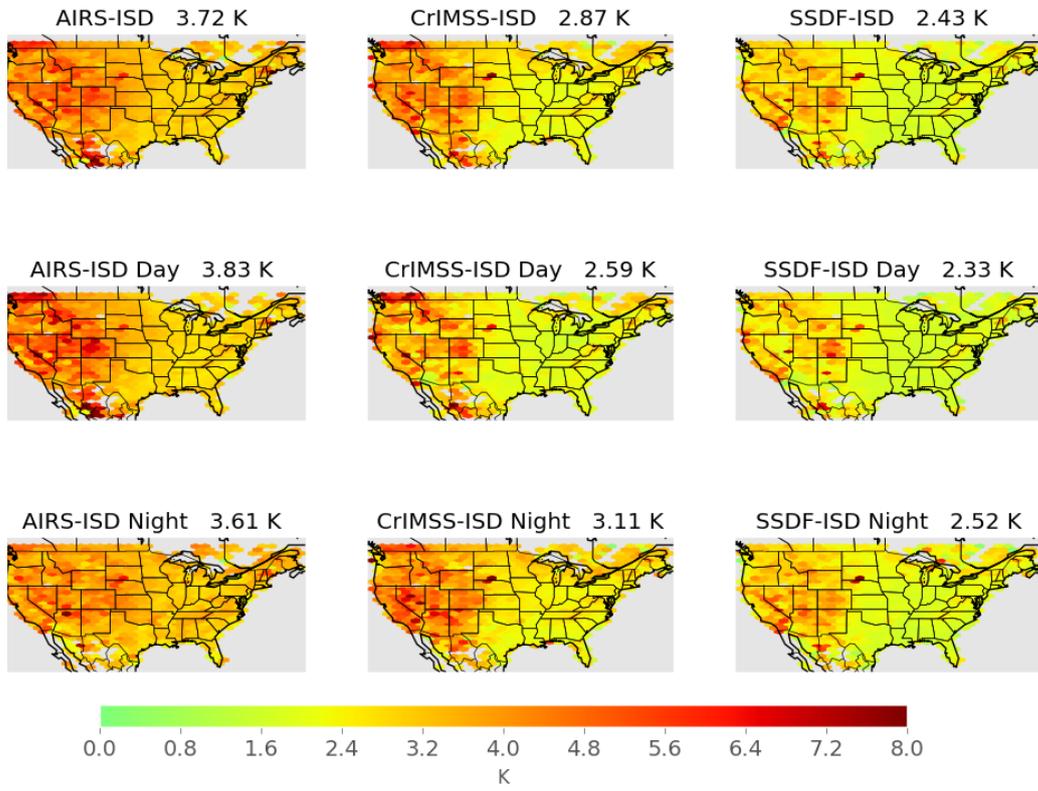


Figure 6: RMSE maps. The nine panels are similar to those in Figure 4 but for RMSE.

374 We repeated this analysis over CONUS and the 2013-2020 period for the SSDF  
 375 product created with AIRS alone, without CrIMSS. We found similar improvements in  
 376 bias, standard deviation, and RMSE. The mean bias of the AIRS-only SSDF product  
 377 over the entire domain was  $-0.08$  K for daytime only, and  $-0.03$  K for nighttime only.  
 378 The overall RMSE was  $2.52$  K, 4% higher than the overall RMSE of the SSDF product  
 379 created from both AIRS and CrIMSS.

380 Figure 7 shows histograms of the NSAT error (retrieval/reanalysis - ISD) for the  
 381 year 2013, over CONUS only. The three comparison datasets (AIRS, CrIMSS, and  
 382 ERA5-Land) were matched separately to SSDF outputs, to ensure that the SSDF  
 383 product and each corresponding comparison dataset are considering the same scenes.  
 384 The SSDF error histograms are symmetric with a single mode and peak at 0 for both  
 385 day and night, which is consistent with the errors being unbiased. The AIRS histogram  
 386 exhibits a cold bias during the day and a warm bias at night. CrIMSS has a similar  
 387 day/night bias shift, but of a smaller magnitude. A cold bias over land, particularly  
 388 at higher temperatures, has been previously noted for both input datasets (Yue et al.,  
 389 2020, 2021), although there have been few validation studies (Ferguson & Wood, 2010;  
 390 Sun et al., 2021). The SSDF product exhibits smaller mean biases and RMSEs than  
 391 either input dataset. On average, over both input datasets, daytime and nighttime,  
 392 SSDF decreases mean bias magnitude by 81% and mean RMSE by 23% relative to the  
 393 input datasets.

394 Next, we examine the seasonality of bias and RMSE. Figure 8 shows the mean  
 395 bias (retrieval/reanalysis - ISD) by month split into day/night to examine seasonality.  
 396 There is a significant cold bias during the day for AIRS and CrIMSS that switches  
 397 to a warm bias at night. During the day, AIRS has a smaller bias during winter  
 398 months (Dec/Jan/Feb) and a larger bias during summer months (Jun/Jul/Aug). This  
 399 is switched during nighttime where a larger warm bias is observed during winter and  
 400 a smaller warm bias is observed during summer. These AIRS biases are of course also  
 401 apparent in Figure 7. The SSDF product is relatively unbiased for both day and night.  
 402 The SSDF bias magnitude is slightly larger during the day than night. From May to  
 403 December, the SSDF product has a smaller bias at night than does ERA5-Land while  
 404 during the day the reanalysis and the SSDF mean biases are of similar magnitude.

405 Figure 9 shows mean RMSE (retrieval/reanalysis - ISD) by month split by day/night,  
 406 i.e., the mean RMSE values calculated in 2-degree spatial bins. RMSE is largest for  
 407 AIRS, particularly during the day. Generally, RMSE is higher in winter and lower in  
 408 summer. During the day, the ERA5-Land has the lowest RMSE. At night, the SSDF  
 409 RMSE is comparable and sometimes lower than the ERA5-Land RMSE.

410 We next examine relative performance in hot and cold extremes. Figure 10 shows  
 411 the mean bias (retrieval/reanalysis - ISD) by ISD percentile of the ISD matchups. The  
 412 error bars are the standard error of the mean at the 95 percent confidence level. The  
 413 lighter shade of every color is the matched SSDF corresponding to the comparison  
 414 dataset. All retrievals and reanalysis do best in the mean state (25th to 75th per-  
 415 centile). At the extremes, each of the datasets being compared to ISD have warm  
 416 biases for low values (1st through the 15th percentile) and cold biases for high val-  
 417 ues (85th through the 99th); in other words, these datasets dampen out capture cold  
 418 or warm extremes represented in the ISD. The SSDF product captures the extremes  
 419 better than the input datasets, AIRS and CrIMSS. However, the reanalysis generally  
 420 does best, having the smallest bias regardless of percentile, and is better at capturing  
 421 the extremes.

422 We next examine performance at extremely high elevations. Figure 11 shows  
 423 mean biases (retrieval/reanalysis - ISD) aggregated by ISD elevation. At around 2500  
 424 meters, mean biases increase with elevation in the SSDF product, AIRS, CrIMSS, and  
 425 reanalysis. Daytime mean biases at these high elevations are larger in SSDF, although

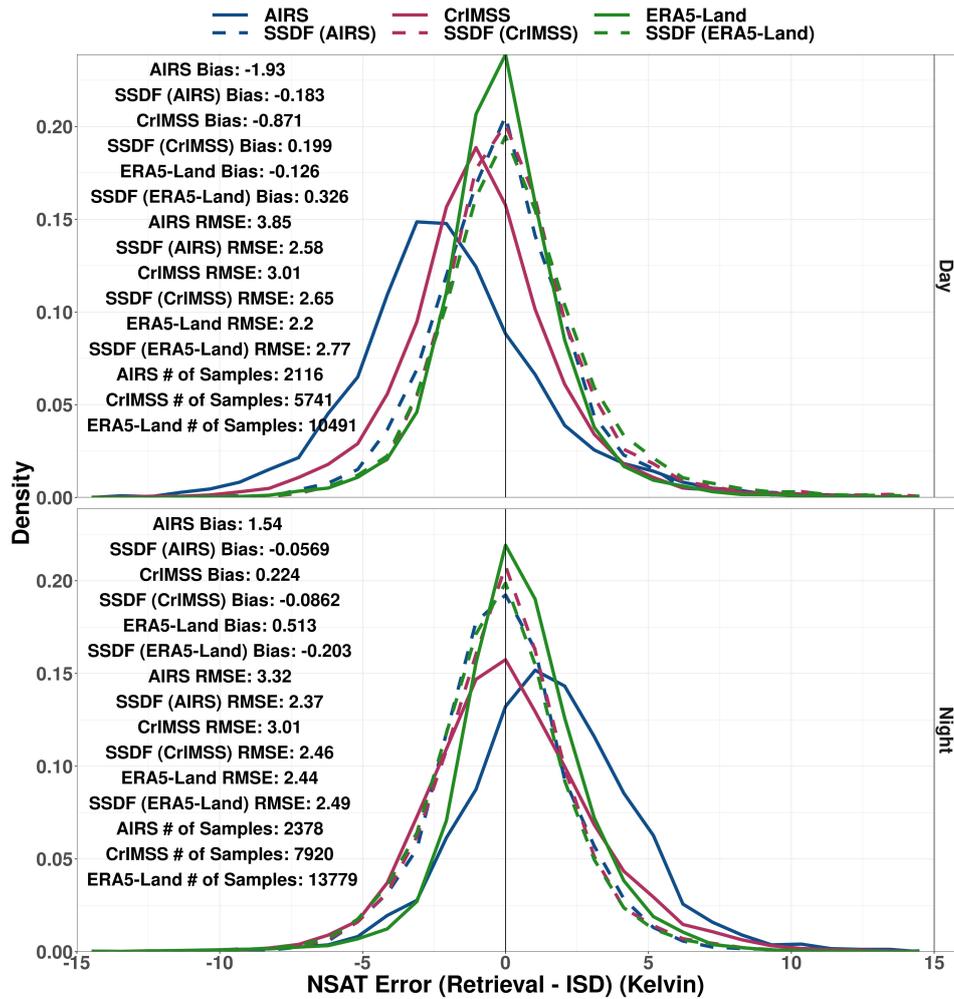


Figure 7: Histograms of errors for day (top) and night (bottom) for 2013 over CONUS, for AIRS (blue), CrIMSS (red) and ERA5-Land (green). The dashed line is the SSDF subset matched to the other datasets. Mean statistics of bias, RMSE, and the number of samples are provided.

426 we note that the sample size is small. At night, SSDF shows lower mean biases than  
 427 AIRS, CrIMSS, or ERA5-Land at high elevations.

428 In order to increase the sample size for high-elevation cases, Figure 12 shows  
 429 the mean biases aggregated by ISD elevation for elevations higher than 2000 meters  
 430 over the period 2012-2020. During the day, the SSDF bias exceeds AIRS and CrIMSS,  
 431 consistent with Figure 11. We hypothesize that this excess bias in SSDF for a very small  
 432 number of data points at very high elevations is caused by the bulk-binning method  
 433 for bias estimation. As Figure 11 shows, both remote sensing datasets exhibit a cold  
 434 bias during the daytime at lower elevations. Because the two-degree hexagonal bins for  
 435 bias estimation are dominated by lower elevations (as the problematic high elevations  
 436 are high mountain surfaces), and because both remote sensing dataset biases switch  
 437 signs from cold bias to warm bias at approximately 2500 m, the cold bias correction  
 438 calculated from the bulk bins ends up exacerbating the warm bias from the input  
 439 datasets at the highest elevations. In a future version of SSDF, we will improve the

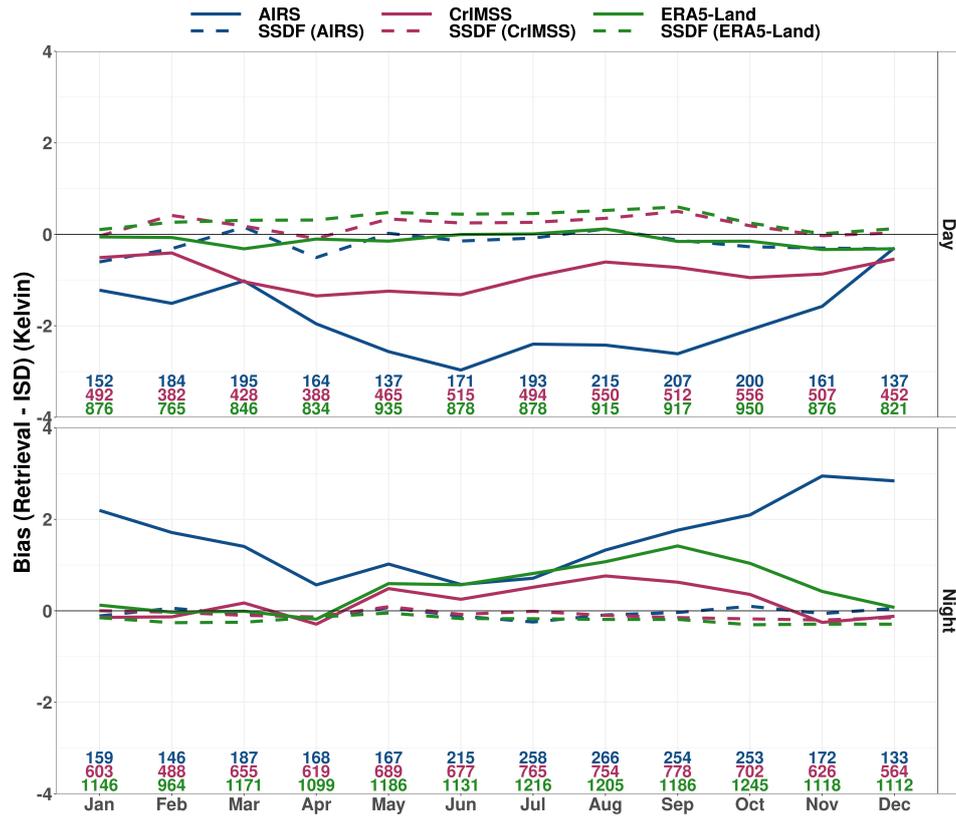


Figure 8: Mean bias as a function of month for day (top) and night (bottom) for 2013 over CONUS. Numbers at the bottom indicate the number of data points, and are color-coded according to dataset.

440 bias estimation of the input datasets, which could mitigate or eliminate this bias at  
 441 the very small number of estimates elevations above 2500 m.



Figure 9: Mean RMSE as a function of month for day (top) and night (bottom) for 2013 over CONUS. Numbers at the bottom indicate the number of data points, and are color-coded according to dataset.

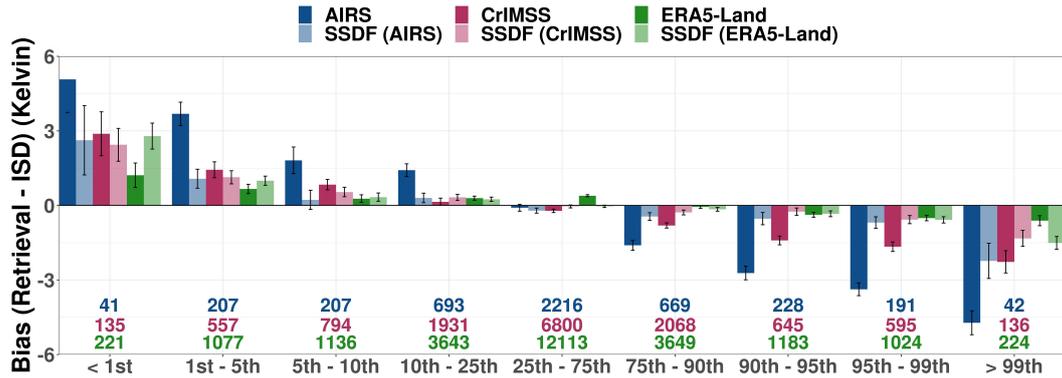


Figure 10: Mean biases as a function of ISD percentile for 2013 over CONUS. Numbers at the bottom indicate the number of data points, and are color-coded according to dataset.

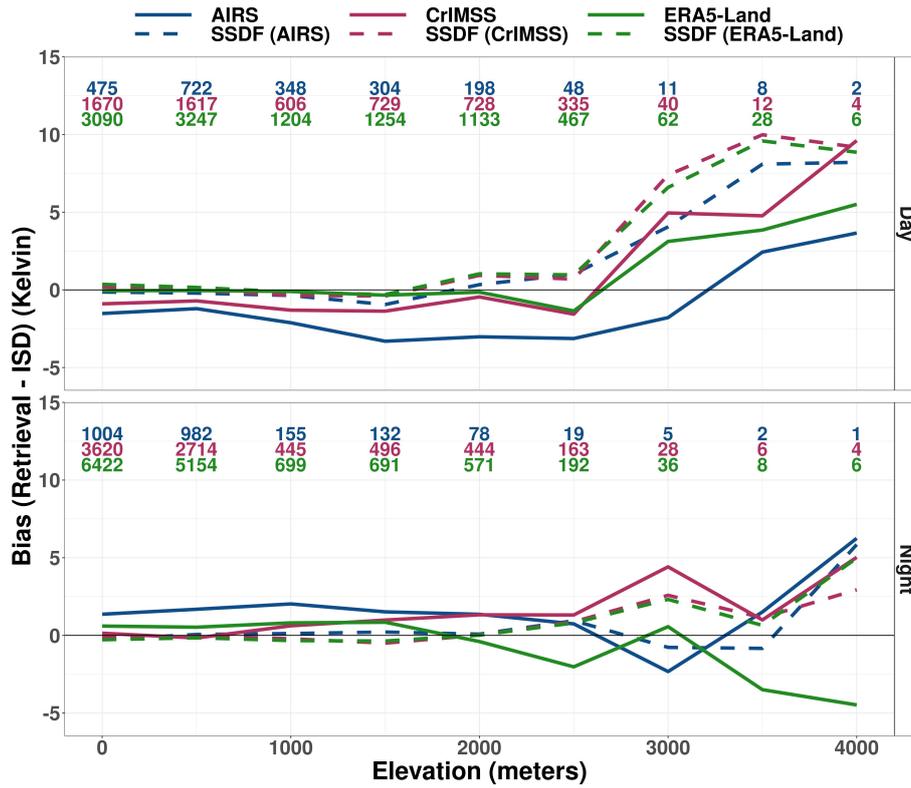


Figure 11: Mean biases as a function of ISD elevation for day (top) and night (bottom) for 2013 over CONUS. Numbers at the top indicate the number of data points, and are color-coded according to dataset.

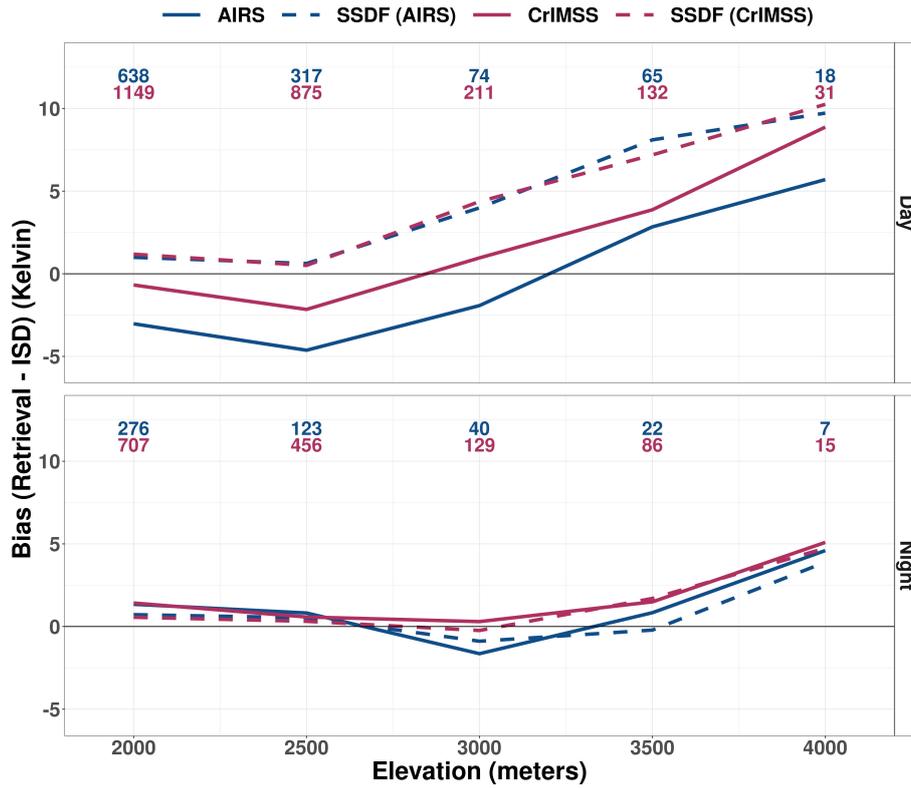


Figure 12: Mean biases as a function of ISD elevation for day (top) and night (bottom) over CONUS from 2012-2020 for AIRS, CrIMSS, and SSDF. Numbers at the top indicate the number of data points, and are color-coded according to dataset.

442

### 3.3 Validation of uncertainty estimates

443

444

445

446

447

448

449

450

The SSDF algorithm provides a mean (prediction/estimate) and standard deviation (uncertainty) of the conditional distribution of true NSAT, given the available inputs; this distribution is termed the predictive distribution. In what follows, this is a Gaussian distribution, centered at the SSDF estimate. This information can be used to construct prediction intervals for the true NSAT. Here we provide a summary and probabilistic assessment of the SSDF predictive distribution along with related information from the AIRS V7 and CrIMSS-CLIMCAPS V2 products. In the notation that follows, we use the subscript  $i$  in place of the areal unit notation  $B_i$ .

451

452

453

454

455

456

457

458

459

460

461

- In addition to each SSDF NSAT estimate,  $\hat{Y}_i$ , the algorithm also provides the conditional standard deviation of the predictive distribution, denoted  $\hat{\sigma}_{\hat{Y},i}$ .
- The AIRS V7 NSAT retrieval,  $Z_{1,i}$ , is accompanied by a corresponding uncertainty estimate, denoted  $\hat{\sigma}_{Z,1,i}$  (Susskind et al., 2014). This estimate results from a regression model for predicting the absolute retrieval error given several predictors available from the retrieval.
- The CrIMSS-CLIMCAPS V2 retrieval,  $Z_{2,i}$ , also has a corresponding uncertainty estimate, denoted  $\hat{\sigma}_{Z,2,i}$  (N. Smith & Barnet, 2020). This estimate results from a linear approximation of the posterior standard deviation of the true state given the observed radiances for a single footprint and is an output of the optimal estimation (OE) approach used in CLIMCAPS.

462

463

464

465

466

467

Figure 13 shows histograms of these uncertainty estimates:  $\hat{\sigma}_{Z,1}$ ,  $\hat{\sigma}_{Z,2}$ , and  $\hat{\sigma}_{\hat{Y}}$  across the CONUS data record. The solid line shows uncertainty estimates from AIRS (blue) and CrIMSS (red) while the dashed shows the corresponding matched SSDF uncertainty estimates. CrIMSS has a peak around 1.2 K with a narrow distribution; AIRS V7 has a peak between 1.5 and 2 K with a wide distribution. SSDF uncertainty histograms peak around 2 K.

These uncertainty estimates are properties of distributions, whereas we define error  $e_i$  as a realization of a random variable that represents the difference between an estimate and the true state. For example, the error for SSDF is  $e_{\hat{y},i} = \hat{Y}_i - Y_i$ , where  $Y_i$  is the ISD validation for collocation  $i$ . If the predictive distribution is assumed to be Gaussian, the empirical coverage of intervals of the form

$$\hat{Y}_i \pm c \hat{\sigma}_{\hat{Y},i},$$

468

469

470

471

can be assessed for the ISD matchups. In the case of an unbiased estimate, “well-calibrated” uncertainty estimates, and a Gaussian distribution; intervals with  $c = 1$  should cover the true state  $Y_i$  about 68% of the time, and about 95% of the time for  $c = 2$ .

472

473

474

475

476

477

478

479

480

481

Figure 14 shows scatterplots of the joint distribution of the uncertainty estimate (x-axis) and the observed error (retrieval-*ISD*). There are many cases for AIRS and CrIMSS where the uncertainty estimate grossly underestimates the true error; over 15% of the time for both datasets and for day and night, the true error is more than three times greater than the uncertainty estimate. However, this occurs about 3% of the time with SSDF in the day and fewer than 5% of the time at night. Overall, the CrIMSS uncertainty estimates are distributed too narrowly, and with a peak too low, to capture the true error. The AIRS uncertainty estimates also peak at a value below the peak of the error distribution, although the uncertainty estimate distribution is much wider, including a very long tail of high uncertainty estimates.

482

483

484

In general, SSDF uncertainty estimates are consistent with statistical expectations under Gaussian assumptions. For example, one would expect one-sigma uncertainty estimates to cover a standard error distribution 68% of the time, and we see

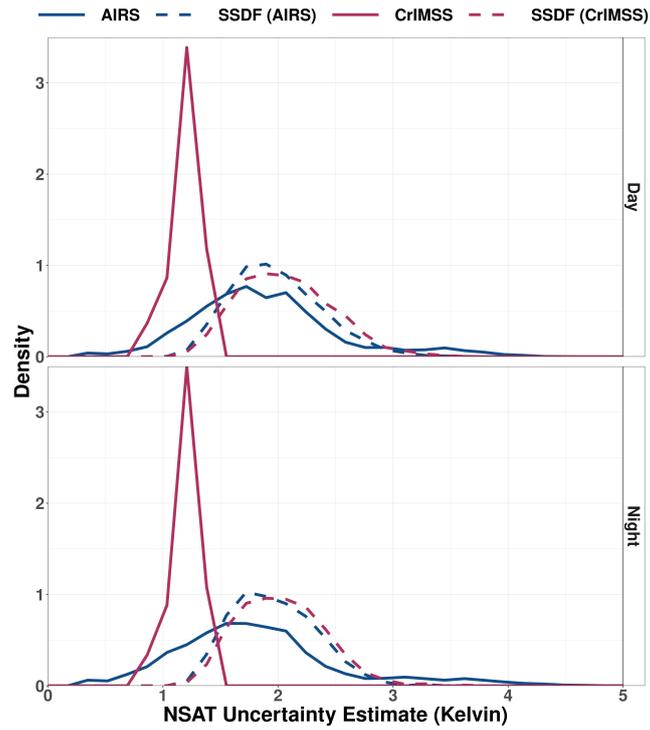


Figure 13: Histograms of uncertainty estimates for day (top) and night (bottom) for 2013 over CONUS.

485 that the SSDF uncertainty estimates do so roughly 65% of the time in daytime. Simi-  
 486 larly, one would expect the estimates to cover 95% and over 99% at the 2- and 3-sigma  
 487 levels, with SSDF covering about 90% and 97% during daytime.

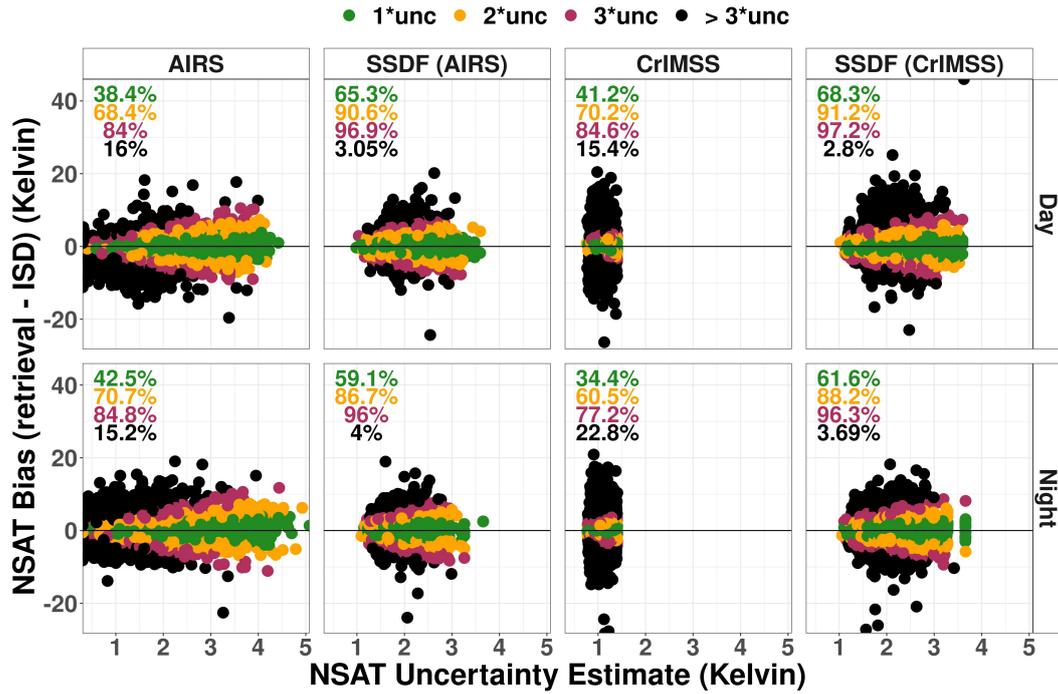


Figure 14: Observed errors (retrieval - ISD) versus uncertainty estimates for day (top) and night (bottom) for 2013 over CONUS. The colors show whether the range of each observed error was within the uncertainty bound, as described in the text:  $1\times$ uncertainty (green, should cover the true state about 68% of the time),  $2\times$ uncertainty (orange, should cover the true state about 95% of the time),  $3\times$ uncertainty (red, should cover the true state about 99% of the time) or  $> 3\times$ uncertainty (black).

488

### 3.4 Empirical distribution consistency

489

490

491

492

493

494

495

496

497

The ISD record provides a sample of the empirical distribution of NSAT over CONUS. Here, we assess the relative consistency of the SSDF empirical distribution versus the other products against the ISD reference distribution. Figure 15 shows an example of the empirical cumulative distribution (ECDF) for the ISD (pink) and AIRS (blue). While it is almost certainly the case that the products' ECDFs deviate from the ISD reference distribution in some subtle ways, we evaluate their relative consistency with ISD through a series of hypothesis tests. Figure 16 shows the difference between the ECDF of the retrieval/reanalysis to the ECDF of ISD. The AIRS ECDF has the largest difference to the ISD ECDF, particularly during the Day.

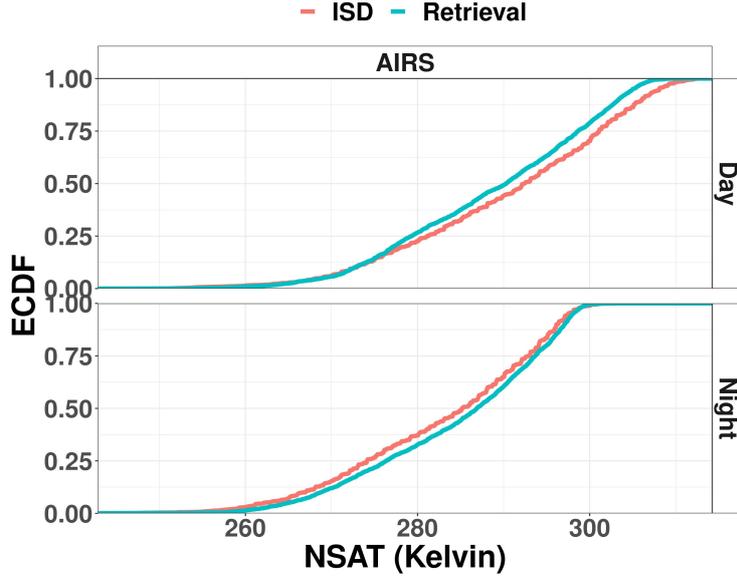


Figure 15: ECDF for AIRS (blue) and ISD (pink) for day (top) and night (bottom) for 2013 over CONUS.

498

499

500

501

502

503

504

505

The SSDF estimates are tested against each of the other products (AIRS, CrIMSS, ERA5-Land) for night and day conditions. Each assessment is carried out using a randomization or resampling test (Wilks, 2006). For this test, the null hypothesis is that the empirical distributions of SSDF and the comparison product deviate equally from the ISD reference distribution. The alternative hypothesis is that either SSDF or the comparison product have an empirical distribution that is closer to the ISD reference distribution. For this procedure, the test statistic is computed as the difference in two-sample Kolmogorov-Smirnov (KS) statistics for the products versus ISD.

For each instance of the test, we have a collection of matched triples  $\{\hat{\mathbf{Y}}, \mathbf{Z}_k, \mathbf{Y}\}$ ; where  $\hat{\mathbf{Y}} \equiv \{\hat{Y}_i\}; i = 1, \dots, n$  are the SSDF estimates,  $\mathbf{Z}_k \equiv \{Z_{k,i}\}; i = 1, \dots, n$  are the comparison products, and  $\mathbf{Y} \equiv \{Y_i\}; i = 1, \dots, n$  are the ISD NSAT. As above,  $k = 1$  for AIRS,  $k = 2$  for CrIMSS, and here  $k = 3$  for ERA5-Land. Then, test  $k$  has a test statistic

$$\gamma_k = \delta(\hat{\mathbf{Y}}, \mathbf{Y}) - \delta(\mathbf{Z}_k, \mathbf{Y}),$$

506

507

508

509

where  $\delta$  is the traditional two-sample KS statistic. The KS statistic is the maximum difference in the two ECDFs being compared. Thus, the test statistic  $\gamma_k$  for the current test is a *difference* of ECDF deviations. A negative value is an indication that the SSDF distribution is closer to ISD than the comparison product.

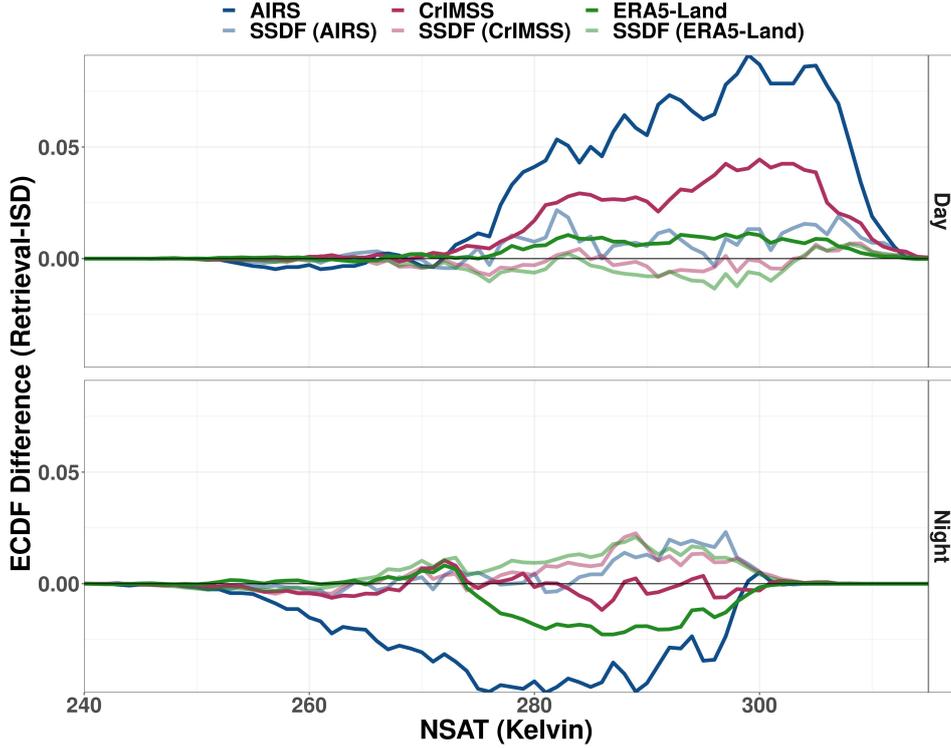


Figure 16: The ECDF difference between the retrieval/reanalysis and the ISD color coded for day (top) and night (bottom) for 2013 over CONUS.

510 The distribution of the test statistic under the null hypothesis can be estab-  
 511 lished through a resampling procedure. The procedure should preserve the inherent  
 512 dependence of the matched triples, but the assignment of the two comparison groups  
 513 can be shuffled randomly. A null distribution is generated by repeating these steps  
 514  $m = 1, \dots, M$  times:

- 515 1. Define shuffled data vectors  $\mathbf{W}_{m,1}$  and  $\mathbf{W}_{m,2}$ .
- 516 2. For each validation matchup ( $i = 1, \dots, n_k$ ), assign  $W_{i,m,1} = \hat{Y}_i$  and  $W_{i,m,2} =$   
 517  $Z_{k,i}$  with probability 0.5; otherwise assign  $W_{i,m,1} = Z_{k,i}$  and  $W_{i,m,2} = \hat{Y}_i$ . This  
 518 effectively shuffles the labels for SSDF and the comparison product for each  
 519 matchup.
3. Compute the test statistic for the randomized samples,

$$\gamma_{0,m,k} = \delta(\mathbf{W}_{m,1}, \mathbf{Y}) - \delta(\mathbf{W}_{m,2}, \mathbf{Y}),$$

The distribution of  $\gamma_{0,m,k}$  provides the null distribution of the test statistic for each test. Figure 17 displays the test statistics  $\gamma_k$  along with density plots of the null distributions of test statistics  $\gamma_{0,m,k}$  for  $M = 20,000$  resampled datasets for each test. A two-sided  $p$ -value can be computed for each test as

$$p_k = \frac{1}{M} \sum_{m=1}^M I_\gamma(|\gamma_{0,m,k}| > |\gamma_k|),$$

520 where  $I_\gamma$  is an indicator function.

521 The  $p$ -values for each of the resampling tests of SSDF versus other products are  
 522 displayed as text in Figure 17. All tests, except the night comparison of SSDF and  
 523 CrIMSS, yield  $p$ -values of 0, indicating a significant difference in consistency with the  
 524 ISD reference distribution. These results can also be seen visually as the observed test  
 525 statistics  $\gamma_k$ , shown as vertical lines, lie well outside the corresponding null distribu-  
 526 tions. The tests indicate SSDF is more consistent with ISD than AIRS for both day  
 527 and night conditions, as well as a favorable result for SSDF versus CrIMSS for day and  
 528 versus ERA5-Land at night. The positive test statistic for SSDF versus ERA5-Land  
 529 during the day indicates the reanalysis is more consistent with ISD in this case.

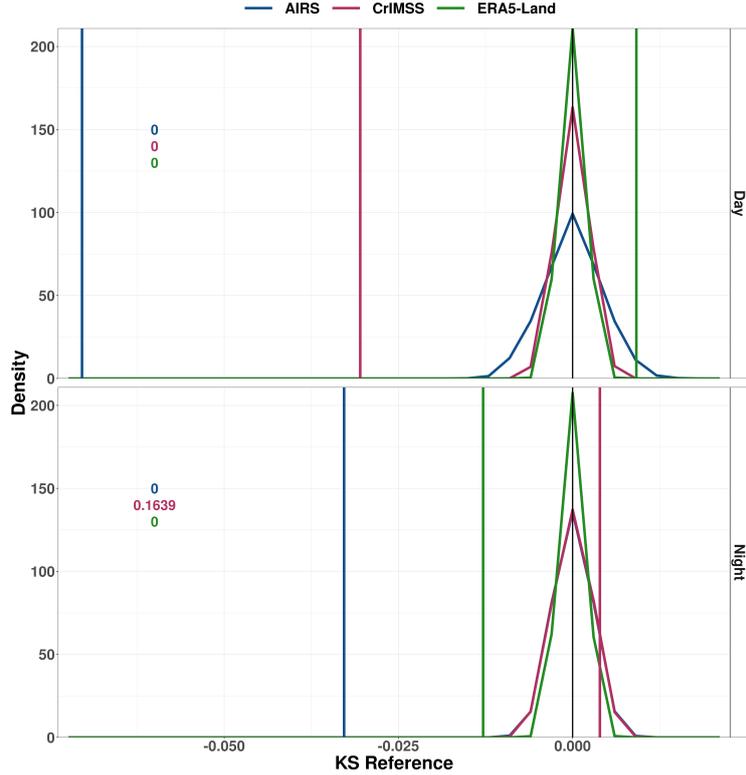


Figure 17: Histogram of the KS statistic for AIRS (blue), CrIMSS (maroon) and ERA5-Land (green), for day (top) and night (bottom) for 2013 over CONUS. The corresponding  $p$ -value is color-coded on the left side.

### 530 3.5 Long-term stationarity

531 We next assess the stationarity in the bias of the SSDF dataset. First, we exam-  
 532 ine the annual mean bias over the entire record relative to the withheld ISD reference  
 533 data. Figure 18 shows the annual mean bias (both day and night) for both the input  
 534 datasets, as well as for two periods of SSDF: the pre-CrIMSS period (2003 to 2011,  
 535 inclusive) and the post-CrIMSS period (2013 to 2020, inclusive). Shading shows two  
 536 standard deviations of these annual bias estimates, with the two SSDF periods calcu-  
 537 lated separately. We exclude 2002 as this year only includes 4 months of AIRS data,  
 538 and we exclude 2012 as this year was a mixture of AIRS-only and AIRS-plus-CrIMSS.

539 These summary data clearly show that SSDF significantly improves both the  
 540 mean annual bias, and the standard deviation in mean annual bias, relative to the

541 input datasets. The mean of these annual bias estimates are  $-0.10^{\circ}\text{C}$ ,  $-0.23^{\circ}\text{C}$ , and  
 542  $0.02^{\circ}\text{C}$  for AIRS, CrIMSS, and SSDF respectively, from 2003 to 2020 inclusive for  
 543 AIRS and SSDF and from 2013 to 2020 inclusive for CrIMSS. However, these data  
 544 also suggest a step change in SSDF mean annual bias in the pre-CrIMSS and post-  
 545 CrIMSS period. The mean of the SSDF mean annual bias estimates in the pre-CrIMSS  
 546 and post-CrIMSS periods are  $-0.020^{\circ}\text{C}$  and  $0.076^{\circ}\text{C}$ , respectively, a shift of about  $0.1^{\circ}\text{C}$ .  
 547 This shift is small compared to the biases in the input remote sensing datasets, and the  
 548 apparent downward trend in the AIRS dataset. Over-correction with the addition of  
 549 the CrIMSS dataset might be an artifact of the bias estimation bulk-binning procedure.  
 550 This small step change in bias does not occur in the AIRS-only SSDF product over the  
 551 full AIRS record. Future versions of SSDF will use improved uncertainty quantification  
 552 methods to estimate input dataset biases, which could mitigate or eliminate this small  
 553 shift in annual mean bias in transitioning from the AIRS-only SSDF product to the  
 554 two-instrument product. In the meantime, the first version of our product creates a  
 555 more coherent and stable climate record than the two input datasets taken separately.

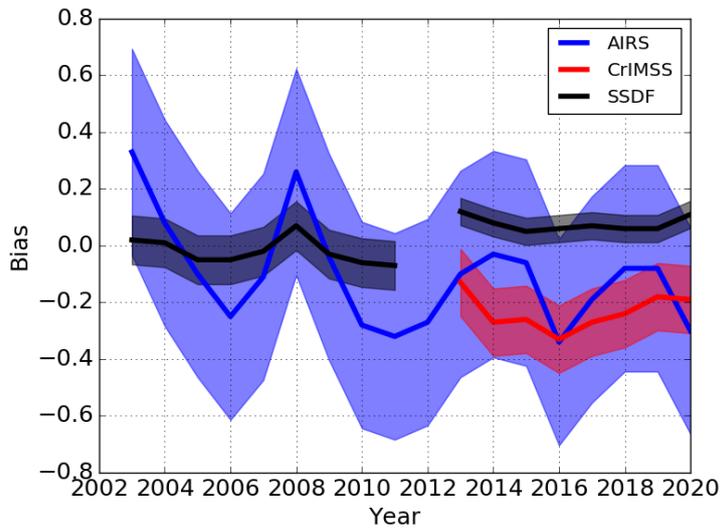


Figure 18: Annual mean bias for each year of the data record, for the SSDF product and each of the two remote sensing input products, relative to the withheld ISD data. Shading shows two standard deviations of these annual bias estimates, with the two SSDF periods calculated separately.

556 Figure 19 shows the histogram of the SSDF uncertainty estimates for 2011 (black)  
 557 and 2013 (red). The mean uncertainty is provided as text. The histograms are compar-  
 558 able. The 2011 (single instrument only) histogram is shifted slightly to the right  
 559 suggesting higher uncertainty estimates with one instrument compared to two. In-  
 560 deed, the mean SSDF uncertainty estimate is 2.15/2.21 (Day/Night) during 2011 and  
 561 decreases to 2.12/2.09 in 2013. However, this is to be expected as the additional  
 562 information from CrIMSS provides greater certainty for SSDF.

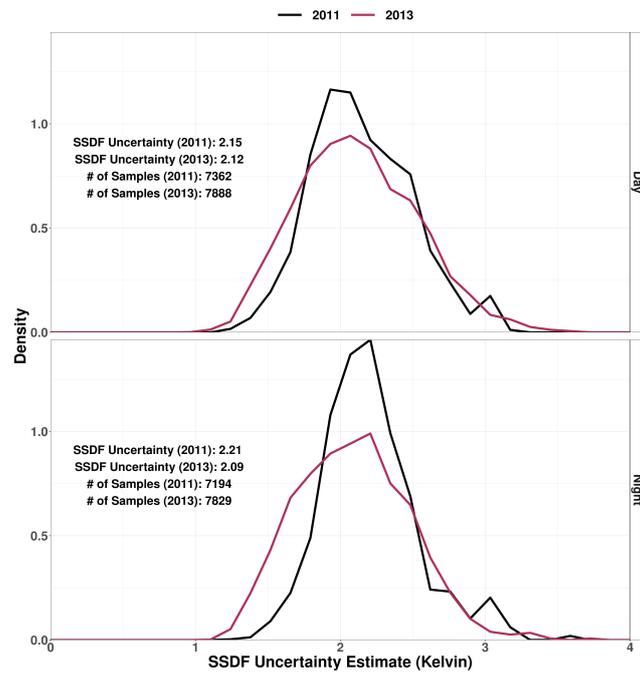


Figure 19: SSDF uncertainty histogram for 2011 (black) and 2013 (red) aggregated by day (top) and night (bottom). Summary statistics of mean SSDF uncertainty are provided as text on the upper left.

## 4 Discussion and conclusion

We have produced a new fused NSAT product over CONUS, from November 2012 through December 2020, using Spatial Statistical Data Fusion of AIRS and CrIMSS remotely sensed NSAT. We also provided detailed validation using withheld ISD data, and comparison to ECMWF ERA5-Land reanalysis. Remote sensing data provides information to span the spatial domain, in situ data provides the information to correct the remote sensing data, and SSDF provides the means to join them into something greater than the sum of their parts. The fused gridded product has no missing data (apart from one day and night without either AIRS or CrIMSS-CLIMCAPS input data); has improved accuracy and precision relative to the input satellite datasets, has comparable accuracy and precision relative to ERA5-Land and indeed significantly lower nighttime bias than ERA5-Land; and includes estimates that are more consistent with the observed errors relative to in situ ISD observations. To summarize, our NSAT SSDF pilot product is comparable in precision and accuracy to the cutting-edge ERA5-Land reanalysis, but it is a direct observational product that does not involve physical modeling. Furthermore, unlike reanalysis it could support near-real-time product creation for operational applications.

SSDF is general and could be applied to any number of datasets estimating the same observable. It could be applied across a wide range of satellite observables, such as atmospheric composition, water vapor profiles, or vapor pressure deficit (the difference between the water vapour pressure and the saturation water vapour pressure), so long as uncertainty estimates of the input datasets can be obtained. We emphasize that the quality of the SSDF product depends on the quality of the bias and variance estimates of the input datasets.

Our plans for future work include improving the bias and variance estimation using simulation-based uncertainty quantification (Hobbs et al., 2017; Braverman et al., 2021). Simulation-based uncertainty quantification has the potential to further improve the overall quality of the SSDF product. It could also mitigate or eliminate the two issues our validation has uncovered: increased bias at a small number of data points at elevations in excess of 2500 m, and a 0.1 K shift in annual mean bias when transitioning from the AIRS-only version (2002-2012) to the two input (AIRS+CrIMSS) SSDF version (2012-2020).

We also plan to create an NSAT SSDF product over global land areas, expanding beyond CONUS, and apply the SSDF method to other hyperspectral surface products (e.g., vapor pressure deficit). Finally, we plan to develop SSDF products for satellite instruments that sample observables at different points in the diurnal cycle, to enable fusion of datasets from polar-orbiting and inclined platforms to make optimal use of all available remote sensing.

## Open Research

The SSDF NSAT dataset described in this paper is available at <http://dx.doi.org/10.5067/CPXNAPA2WSQ8>.

Publicly available data were obtained from the NASA Atmospheric Infrared Sounder and the Suomi-NPP projects, the NOAA Integrated Surface Database, and the European Centre for Medium-Range Weather Forecasts reanalysis.

Aqua AIRS V7 is available from the NASA GES DISC repository (AIRS Project, 2019). The retrieved surface air temperature (TSurfAir), the corresponding error estimate for TsurfAir (TSurfAirErr), and the corresponding quality flag (QC) (TSurfAir\_QC) were obtained for the standard IR-only product.

611 CrIMSS-CLIMCAPS V2 is available from the NASA GES DISC repository (Barnet,  
612 2019). Near surface temperature (surf\_air\_temp), the corresponding QC flag (surf\_air\_temp-qc),  
613 and the corresponding error estimate (surf\_air\_temp\_err) were obtained from the NSR  
614 product.

615 NOAA ISD NSAT data is available using the rnoaa R package.

616 ECMWF ERA5-Land gridded hourly 2 m temperature means are available from  
617 the Copernicus Climate Change Service (C3S) Climate Data Store (Copernicus 2017).

### 618 **Acknowledgments**

619 The research was carried out at the Jet Propulsion Laboratory, California Institute of  
620 Technology, under a contract with the National Aeronautics and Space Administration  
621 (80NM0018D0004). We thank Evan Fishbein, Evan Manning, Erik Fetzer, and Bjorn  
622 Lambrigtsen for helpful discussions. ©2022

## Appendix A Matchups and bias estimation

In this section, we will elaborate in detail our procedure for matching between ISD and the instruments' observations, and the consequent bias estimation process. For clarity, we establish the following notation. Let  $\mathbf{s}$ ,  $\mathbf{u}$ , and  $\mathbf{v}$  be latitude-longitude locations; e.g.,  $\mathbf{s} = (lat, lon)$ . On a given day (or night) let  $Z^k(\mathbf{u})$  be the value of the  $k$ -th instrument's near-surface temperature retrieval centered at  $\mathbf{u}$ . and focus on a single ISD station at location  $\mathbf{s}$  during a single period. Let  $t_1^I(\mathbf{s}), \dots, t_M^I(\mathbf{s})$  be the times at which observations are acquired at this station during the period. These time points may be irregularly spaced, and  $M$  can change from station to station. The ISD measurements are  $Z^I(\mathbf{s}, Z_m^I(\mathbf{s}))$ ,  $m = 1, \dots, M$ .

Let  $t^k(\mathbf{u})$  be the acquisition times associated with the  $k$ -th instrument's footprints centered at location  $\mathbf{u}$ . In principle,  $\mathbf{u}$  ranges over all footprint locations for the appropriate instrument during the entire period, but in practice these locations are grouped by granules. We denote granule number during the current period by  $g = 1, \dots, 120$ , and the set of footprints belonging to granule  $g$  by  $\mathcal{G}_g^k$ . The time associated with  $\mathcal{G}_g^k$  is  $\tau_g^k$ . To ease the computational burden,  $\mathbf{u}$  ranges only over locations in the single granule with time that is closest to  $t_m^I(\mathbf{s})$ .

A matchup associates the location and time of an ISD value,  $(\mathbf{s}, t^I(\mathbf{s}))$ , with the location and time of the  $k$ -th instrument's footprint in the period:  $(\mathbf{u}^*, t^k(\mathbf{u}^*))$ . The *matchup function* is,

$$\begin{aligned} \mathbb{M}^k(\mathbf{s}, t_m^I(\mathbf{s})) &= (\mathbf{u}^*, t^k(\mathbf{u}^*)), \\ \mathbf{u}^* &= \underset{\mathbf{u}}{\operatorname{argmin}} \left\{ \|\mathbf{u} - \mathbf{s}\|, \mathbf{u} \in (\mathcal{G}_{g^*}^k \cap \mathcal{U}^{time} \cap \mathcal{U}^{space}) \right\}, \\ g^* &= \underset{g}{\operatorname{argmin}} \left\{ |\tau_g^k - t_m^I(\mathbf{s})| \right\}, \\ \mathcal{U}^{time} &= \left\{ \mathbf{u} : |t^k(\mathbf{u}) - t_m^I(\mathbf{s})| \leq 1 \text{ hour} \right\}, \quad \mathcal{U}^{space} = \left\{ \mathbf{u} : \|\mathbf{u} - \mathbf{s}\| \leq 100 \text{ km} \right\}. \end{aligned}$$

Note that, for a given instrument and period, there will only be one granule that satisfies the criterion provided by  $g^*$ .

For a given ISD station (indexed by location  $\mathbf{s}$ ) in the current period,  $p$ , we create the sets of matchup values for the  $k$ -th instrument as follows,

$$\mathcal{A}^k(p, \mathbf{s}) = \left\{ Z^I(\mathbf{s}, t_m^I(\mathbf{s})), Z^k(\mathbb{M}^k(\mathbf{s}, t_m^I(\mathbf{s}))) \right\}_{m=1}^{M(p, \mathbf{s})}$$

for all ISD time points at  $\mathbf{s}$  indexed by  $m = 1, \dots, M(p, \mathbf{s})$ .  $p$  is identified by a date and a mode (day/night) indicator, e.g.,  $p = (d, j) = (2013-01-01, \text{day})$ .  $M(p, \mathbf{s})$  is the number of ISD station values in period  $p$  at location  $\mathbf{s}$ . There is at most one AIRS and one CrIMSS footprint associated with each station-time, but the same footprint can be associated with more than one station-time. Thus,  $\mathcal{A}^k(p, \mathbf{s})$  may contain multiple elements if there is more than one ISD measurement during period  $p$  at location  $\mathbf{s}$ . They may also be empty if there are no matching AIRS or CrIMSS footprints.

After creating  $\mathcal{A}^k(p, \mathbf{s})$  for all periods and ISD locations, we create supersets of matchup value pairs by combining across three-day moving windows, by mode:

$$\mathcal{A}^{kj}(d, \mathbf{s}) = \mathcal{A}^k(d-1, j, \mathbf{s}) \cup \mathcal{A}^k(d, j, \mathbf{s}) \cup \mathcal{A}^k(d+1, j, \mathbf{s}), \quad \mathcal{A}^{kj}(d) = \bigcup_{\mathbf{s}} \mathcal{A}^{kj}(d, \mathbf{s}).$$

$j \in \{\text{day, night}\}$ . We chose the three-day time window after experimenting with shorter and longer windows. Shorter windows did not provide adequate sample sizes while longer windows failed to capture weather-related changes. Ideally, window duration would be as short as possible since longer time windows result in larger variance estimates in the fused data, relative to withheld ISD data.

The final step before actually computing estimated bias and variance for each AIRS and CrIMSS footprint is to tessellate a 240 km (approximately two degrees), hexagonal spatial grid over CONUS. We do this by creating a discrete global grid using the DGGRID software package (Sahr et al., 2003; Sahr, 2019). One of the centers, for example, is at 87.72550324 W, 40.7908839 N, near Watseka, Illinois; this center uniquely determines the tessellated grid. All elements of  $\mathcal{A}^{kj}(d)$  are sorted in to these grid cells based on the instrument’s footprint locations. Formally, let  $i \in 1, \dots, L$  index grid cell centers, and let  $1_i(\mathbf{u}) = 1$  if  $\mathbf{u}$  lies inside cell  $i$ , and zero otherwise. For grid cell  $i$ , mode  $j$ , and date  $d$ , set

$$\mathcal{A}_i^{kj}(d) = \left\{ \left\{ Z^I(\mathbf{s}, t_m^I(\mathbf{s})), Z^k(\mathbf{u}_{ms}^*, t^k(\mathbf{u}_{ms}^*)) : 1_i(\mathbf{u}_{ms}^*) = 1 \right\}_{m=1}^{M(d,j,\mathbf{s})} \right\}_{all \mathbf{s}},$$

654 where  $M(d, j, \mathbf{s})$  is the number of time points acquired by the ISD station at  $\mathbf{s}$  on  
 655 day  $d$  in mode  $j$ ,  $L$  is the total number of hexagonal grid cells, and we write  $\mathbf{u}_{ms}^*$  to  
 656 emphasize its dependence on  $m$  and  $\mathbf{s}$  via the matchup functions.

The bias assigned to all footprints from the  $k$ -th instrument observed on day  $d$  in mode  $j$  belonging to grid cell  $i$  is,

$$b_{dji}^k = \frac{1}{|\mathcal{A}_i^{kj}(d)|} \sum_{all \mathbf{s}} \sum_{m=1}^{M(d,j,\mathbf{s})} \left[ Z^k(\mathbf{u}_{ms}^*, t^k(\mathbf{u}_{ms}^*)) - Z^I(\mathbf{s}, t_m^I(\mathbf{s})) \right] 1_i(\mathbf{u}_{ms}^*).$$

The corresponding variance assigned to all footprints observed on day  $d$  in mode  $j$  belonging to grid cell  $i$  is,

$$v_{dji}^k = \frac{1}{|\mathcal{A}_i^{kj}(d)|} \sum_{all \mathbf{s}} \sum_{m=1}^{M(d,j,\mathbf{s})} \left[ Z^k(\mathbf{u}_{ms}^*, t^k(\mathbf{u}_{ms}^*)) - Z^I(\mathbf{s}, t_m^I(\mathbf{s})) - b_{dji}^k \right]^2 1_i(\mathbf{u}_{ms}^*),$$

Subtracting the biases from the satellite footprints yields bias-corrected data. Denote an footprint acquired by the  $k$ -th instrument on day  $d$  in mode  $j$ , centered at location  $\mathbf{u}$ , by  $Z_{dj}^A(\mathbf{u})$ , where we suppress the argument  $t^A(\mathbf{u})$  since, for a given date and mode, location and time are confounded. The bias-corrected value is denoted by  $Z_{dj}^{k*}(\mathbf{u})$  as follow:

$$Z_{dj}^{k*}(\mathbf{u}) = Z_{dj}^A(\mathbf{u}) - b_{dji^*}^A, \quad i^* = \underset{i}{\operatorname{argmax}} 1_i(\mathbf{u}),$$

657 with associated variance  $v_{dji^*}^k$ .

## 658 References

- 659 AIRS Project. (2019). *Aqua/AIRS L2 Standard Physical Retrieval (AIRS-*  
 660 *only) V7.0*. Goddard Earth Sciences Data and Information Services Center  
 661 (GES DISC). Greenbelt, MD, USA. (Accessed: 2019-2021) doi:  
 662 10.5067/VP1M6OG1X7M1  
 663 AIRS Project. (2020). *About the data*. Retrieved 2021-12-26, from [https://airs](https://airs.jpl.nasa.gov/data/about-the-data/granules/)  
 664 [.jpl.nasa.gov/data/about-the-data/granules/](https://airs.jpl.nasa.gov/data/about-the-data/granules/)  
 665 Barnett, C. (2019). *Sounder SIPS: Suomi NPP CrIMSS Level 2 CLIMCAPS normal*  
 666 *spectral resolution: Cloud cleared radiances v2*. Goddard Earth Sciences Data  
 667 and Information Services Center (GES DISC). Greenbelt, MD, USA. (Ac-  
 668 cessed: 2019-2021) doi: 10.5067/CNG0ST72533Z  
 669 Blackwell, W. (2005). A neural-network technique for the retrieval of atmospheric  
 670 temperature and moisture profiles from high spectral resolution sounding data.  
 671 *Geoscience and Remote Sensing, IEEE Transactions on*, 43(11), 2535-2546.  
 672 doi: 10.1109/TGRS.2005.855071

- 673 Braverman, A., Hobbs, J., Teixeira, J., & Gunson, M. (2021). Post hoc uncertainty  
674 quantification for remote sensing observing systems. *SIAM/ASA Journal on*  
675 *Uncertainty Quantification*, 9(3), 1064–1093.
- 676 Chahine, M. T., Pagano, T. S., Aumann, H. H., Atlas, R., Barnet, C., Blaisdell, J.,  
677 ... Zhou, L. (2006). AIRS: Improving weather forecasting and providing new  
678 data on greenhouse gases. *Bulletin of the American Meteorological Society*,  
679 87(7), 911–926.
- 680 Cressie, N. (2015). *Statistics for spatial data*. John Wiley & Sons.
- 681 Cressie, N., & Johannesson, G. (2008). Fixed rank kriging for very large spatial data  
682 sets. *Journal of the Royal Statistical Society: Series B (Statistical Methodol-*  
683 *ogy)*, 70(1), 209–226.
- 684 Ferguson, C. R., & Wood, E. F. (2010). An evaluation of satellite remote sensing  
685 data products for land surface hydrology: Atmospheric infrared sounder. *Jour-*  
686 *nal of Hydrometeorology*, 11(6), 1234–1262.
- 687 Ghamisi, P., Rasti, B., Yokoya, N., Wang, Q., Hofle, B., Bruzzone, L., ... others  
688 (2019). Multisource and multitemporal data fusion in remote sensing: A  
689 comprehensive review of the state of the art. *IEEE Geoscience and Remote*  
690 *Sensing Magazine*, 7(1), 6–39.
- 691 Gotway, C. A., & Young, L. J. (2002, June). Combining incompatible spatial data.  
692 *Journal of the American Statistical Association*, 97, 632–648.
- 693 Hammerling, D. M., Michalak, A. M., & Kawa, S. R. (2012). Mapping of CO<sub>2</sub> at  
694 high spatiotemporal resolution using satellite observations: Global distribu-  
695 tions from OCO-2. *Journal of Geophysical Research: Atmospheres*, 117(D6).
- 696 Harrison, G., & Burt, S. D. (2021). Quantifying uncertainties in climate data: mea-  
697 surement limitations of naturally ventilated thermometer screens. *Environmental*  
698 *Research Communications*.
- 699 Henderson, H., & Searle, S. (1981). On deriving the inverse of a sum of matrices.  
700 *SIAM Review*, 23, 53–60.
- 701 Hennermann, K., & Berrisford, P. (2019). *ERA5 data documentation*, ECMWF.
- 702 Hobbs, J., Braverman, A., Cressie, N., Granat, R., & Gunson, M. (2017).  
703 Simulation-based uncertainty quantification for estimating atmospheric CO<sub>2</sub>  
704 from satellite data. *SIAM/ASA Journal on Uncertainty Quantification*, 5(1),  
705 956–985.
- 706 McNally, A. P., Watts, P. D., A. Smith, J., Engelen, R., Kelly, G. A., Thépaut,  
707 J. N., & Matricardi, M. (2006). The assimilation of airs radiance data at  
708 ecmwf. *Quarterly Journal of the Royal Meteorological Society*, 132(616),  
709 935–957. Retrieved from <http://dx.doi.org/10.1256/qj.04.171> doi:  
710 10.1256/qj.04.171
- 711 Nguyen, H., Cressie, N., & Braverman, A. (2012). Spatial statistical data fusion  
712 for remote sensing applications. *Journal of the American Statistical Asso-*  
713 *ciation*, 107(499), 1004–1018. Retrieved from [https://doi.org/10.1080/](https://doi.org/10.1080/01621459.2012.694717)  
714 [01621459.2012.694717](https://doi.org/10.1080/01621459.2012.694717) doi: 10.1080/01621459.2012.694717
- 715 Nguyen, H., Katzfuss, M., Cressie, N., & Braverman, A. (2014). Spatio-temporal  
716 data fusion for very large remote sensing datasets. *Technometrics*, 56(2), 174–  
717 185.
- 718 Olsen, E. T., Fishbein, E., Manning, E., & Maddy, E. (2017). AIRS/AMSU/HSB  
719 Version 6 L2 product levels, layers and trapezoids. *Jet Propulsion Laboratory,*  
720 *Pasadena, CA, USA*.
- 721 Sahr, K. (2019). *Dggrid version 7.0: User documentation for discrete global grid*  
722 *software*.
- 723 Sahr, K., White, D., & Kimerling, A. J. (2003). Geodesic discrete global grid sys-  
724 tems. *Cartography and Geographic Information Science*, 30(2), 121–134.
- 725 Smith, A., Lott, N., & Vose, R. (2011). The integrated surface database: Recent de-  
726 velopments and partnerships. *Bulletin of the American Meteorological Society*,  
727 92(6), 704–708.

- 728 Smith, N., & Barnett, C. D. (2019). Uncertainty characterization and propagation  
729 in the Community Long-Term Infrared Microwave Combined Atmospheric  
730 Product System (CLIMCAPS). *Remote Sensing*, *11*(10), 1227.
- 731 Smith, N., & Barnett, C. D. (2020). CLIMCAPS observing capability for temper-  
732 ature, moisture, and trace gases from AIRS/AMSU and CrIS/ATMS. *Atmo-  
733 spheric Measurement Techniques*, *13*(8), 4437–4459.
- 734 Sun, J., McColl, K. A., Wang, Y., Rigden, A. J., Lu, H., Yang, K., ... Santanello Jr,  
735 J. A. (2021). Global evaluation of terrestrial near-surface air temperature  
736 and specific humidity retrievals from the atmospheric infrared sounder (airs).  
737 *Remote Sensing of Environment*, *252*, 112146.
- 738 Susskind, J., Blaisdell, J. M., & Iredell, L. (2014). Improved methodology for  
739 surface and atmospheric soundings, error estimates, and quality control pro-  
740 cedures: the atmospheric infrared sounder science team version-6 retrieval  
741 algorithm. *Journal of Applied Remote Sensing*, *8*(1), 084994. Retrieved from  
742 <http://dx.doi.org/10.1117/1.JRS.8.084994> doi: 10.1117/1.JRS.8.084994
- 743 Thrastarson, H. T., Manning, E., Kahn, B., Fetzer, E., Yue, Q., Wong, S., ... others  
744 (2020). AIRS/AMSU/HSB Version 7 Level 2 product user guide. *Jet Propul-  
745 sion Laboratory, California Institute of Technology: Pasadena, CA, USA*,  
746 83–92.
- 747 Wilks, D. S. (2006). *Statistical methods in the atmospheric sciences* (Second ed.).  
748 Academic Press.
- 749 Yue, Q., Lambrigtsen, B., et al. (2017). AIRS V6 test report supplement: Perform-  
750 ance of AIRS+AMSU vs. AIRS-only retrievals. *Jet Propulsion Laboratory,  
751 California Institute of Technology: Pasadena, CA, USA*. Retrieved from  
752 [https://docserver.gesdisc.eosdis.nasa.gov/repository/Mission/  
753 AIRS/3.3.ScienceDataProductDocumentation/3.3.5.ProductQuality/  
754 V6\\_Test\\_Report\\_Supplement\\_Performance\\_of\\_AIRS+AMSU\\_vs\\_AIRS-Only  
755 \\_Retrievals.pdf](https://docserver.gesdisc.eosdis.nasa.gov/repository/Mission/AIRS/3.3.ScienceDataProductDocumentation/3.3.5.ProductQuality/V6_Test_Report_Supplement_Performance_of_AIRS+AMSU_vs_AIRS-Only_Retrievals.pdf)
- 756 Yue, Q., Lambrigtsen, B., et al. (2020). AIRS V7 L2 performance test and vali-  
757 dation report. *Jet Propulsion Laboratory, California Institute of Technology:  
758 Pasadena, CA, USA*. Retrieved from [https://docserver.gesdisc.eosdis  
759 .nasa.gov/public/project/AIRS/V7\\_L2\\_Performance\\_Test\\_and\\_Validation  
760 \\_report.pdf](https://docserver.gesdisc.eosdis.nasa.gov/public/project/AIRS/V7_L2_Performance_Test_and_Validation_report.pdf)
- 761 Yue, Q., Lambrigtsen, B., et al. (2021). Version 2 CLIMCAPS-Aqua re-  
762 trieval product performance test report. *Jet Propulsion Laboratory, Cal-  
763 ifornia Institute of Technology: Pasadena, CA, USA*. Retrieved from  
764 [https://docserver.gesdisc.eosdis.nasa.gov/public/project/Sounder/  
765 CLIMCAPS.V2.Test.Report.Aqua.pdf](https://docserver.gesdisc.eosdis.nasa.gov/public/project/Sounder/CLIMCAPS.V2.Test.Report.Aqua.pdf)