

Data Fusion of AIRS and CrIMSS Near Surface Air Temperature

P. Kalmus¹, H. Nguyen¹, J. Roman¹, T. Wang¹, Q. Yue¹, Y. Wen², J. Hobbs¹,
and A. Braverman¹

¹Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 91109

²Department of Geography, University of Florida, Gainesville, FL, 32611

Key Points:

- We demonstrate spatial statistical fusion for Level 2 remote sensing datasets which estimate the same observable
- We introduce a new daily and nightly 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

Abstract

We present a near surface air temperature (NSAT) fused data product over the contiguous United States using Level 2 data from the Atmospheric Infrared Sounder (AIRS), on the Aqua satellite, and the Cross-track Infrared Microwave Sounding Suite (CrIMSS), on the Suomi National Polar-orbiting Partnership (SNPP) satellite. We create the fused product using Spatial Statistical Data Fusion (SSDF), a procedure for fusing multiple datasets by modeling spatial dependence in the data, along with ground 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 and on a 0.25-degree latitude-longitude grid. We provide detailed validation using withheld ISD data and comparison with 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 improved 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 multiple datasets estimating the same quantity. Finally, because our product reduces bias, it produces long-term datasets across multi-instrument remote sensing records with improved bias stationarity, even as individual missions and their data records begin and end.

Plain Language Summary

We have used a data fusion technique called spatial statistical data fusion (SSDF) to create an improved near surface air temperature (NSAT) dataset by fusing two separate satellite datasets. NSAT is important for a variety of applications, such as drought, wildfire, and extreme heat research and prediction. The two input NSAT datasets come from the AIRS instrument on the Aqua satellite, and the CrIMSS suite on the SNPP satellite. Our fused NSAT product is produced twice daily and on a 0.25-degree latitude-longitude grid. We also performed a detailed validation using withheld reference data (which was not included in the bias-correction data) and comparison with ERA5-Land reanalysis. The new fused 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 improved uncertainty estimates. SSDF is computationally fast and generalizable, capable of data fusion from multiple datasets so long as they estimate the same quantity. Finally, because our product reduces bias, it provides a means of creating high-quality continuous long-term datasets across the years, as individual satellite missions and their data records begin and end.

1 Introduction

Data fusion is the combining of multiple datasets into a single dataset with improved properties relative to the input datasets (for a recent review, see Ghamisi et al. (2019)). Near-surface air temperature (NSAT, the air temperature at a height of 2 m above the surface) is a fundamental variable that critically affects life on the Earth’s surface, and an Essential Climate Variable. Here, we describe the use of spatial statistical data fusion (SSDF) to fuse two Level 2 (L2) satellite NSAT datasets into a single product at 0.25-degree spatial resolution on a twice-daily basis (one daytime and one nighttime estimate per day) over the contiguous United States (CONUS) and adjacent parts of North America. SSDF utilizes spatial dependence within and between the datasets to improve estimates at any given point, including at locations not covered by the input data.

63 As the Earth continues to rapidly heat due to human emissions of greenhouse
64 gases, NSAT remote sensing records are becoming increasingly important for a number
65 of critical science and applied science areas such as health, urban planning, hydrology
66 and water, ecology and conservation, and wildfire prediction. NSAT data records
67 have been produced by a variety of methods which are suited for different purposes.
68 One method is to collect NSAT measurements from ground stations; one example
69 of this type of dataset is the Integrated Surface Database, or ISD (A. Smith et al.,
70 2011). Ground station measurements are relatively accurate, but they are sparse
71 point-source measurements with some regions of the planet having less coverage than
72 others. These strengths and weaknesses make them suitable for use as reference data
73 for validation purposes. Another type of NSAT dataset can be created by filtering and
74 processing these raw NSAT ground measurements into space-filled, gridded climate
75 records useful for climate analysis and climate model validation. These climate records
76 are typically monthly mean products at low resolution, such as the 1-degree resolution
77 Berkeley Earth Monthly Land+Ocean dataset (Rohde & Hausfather, 2020). Berkeley
78 Earth is also experimenting with daily and 0.25-degree-resolution datasets. A third
79 strategy for estimating NSAT is reanalysis, which uses multiple data sources (including
80 satellite data) and dynamical weather models to create dynamically consistent gridded
81 fields. As computational power and algorithm efficiencies have increased, so have the
82 spatial resolutions of reanalysis datasets. An example is the European Centre for
83 Medium-Range Weather Forecasts (ECMWF) Reanalysis 5 (ERA5)-Land reanalysis
84 NSAT dataset (Hennermann & Berrisford, 2019), which has hourly temporal resolution
85 and a spatial resolution of 0.1 degrees, the highest available at the time of writing.
86 Finally, NSAT can be estimated from satellite remote sensing. NSAT can be retrieved
87 from imaging instruments which can estimate land surface temperature (LST) at high
88 resolutions, although obtaining NSAT from LST requires regression modeling which
89 introduces its own errors. An example of NSAT modeled from LST is the EUSTACE
90 project (Good, 2015; Rayner et al., 2020), which produced global daily NSAT at 0.25-
91 degree resolution. Another example used LST data from the Moderate Resolution
92 Imaging Spectroradiometer (MODIS) and a random forest model trained using in situ
93 data from the Global Land Data Assimilation System (GLDAS) to model daily all-
94 sky NSAT at 1 km resolution of mainland China (Chen et al., 2021). NSAT can
95 also be estimated from atmospheric temperature profiles from infrared sounders using
96 interpolation to the surface pressure level, such as the AIRS and CrIMSS products
97 used in this study and described below in Section 2.1.

98 Our data-fusion methodology, SSDF, exists within a geostatistical framework
99 which is a part of the broader area of spatial statistics. Specifically, SSDF is de-
100 signed to provide the principled error characterization and error propagation within
101 data fusion for massive remote sensing data (Nguyen et al., 2012). SSDF has been
102 demonstrated previously in the context of data fusion of L2 satellite remote sens-
103 ing datasets. L2 datasets are geophysical quantities inferred or “retrieved” from the
104 primary observations of radiances by the orbiting instruments (known as “Level 1”
105 data). The SSDF methodology we utilize here was first used to fuse L2 aerosol optical
106 depth from the Multi-angle Imaging Spectroradiometer (MISR) and MODIS aboard
107 the Terra platform. It was subsequently demonstrated in the fusion of L2 total column
108 CO₂ concentration (XCO₂) from the Atmospheric Infrared Sounder (AIRS) aboard the
109 Aqua platform and XCO₂ from the Orbiting Carbon Observatory-2 (OCO-2) (Nguyen
110 et al., 2014). In addition, an SSDF variant called local kriging was used to produce
111 fused estimates of XCO₂ from GOSAT (Hammerling et al., 2012). In the current work,
112 we describe the creation of the first long data record produced by SSDF, and the first
113 data fusion of NSAT by any method.

114 L2 datasets can present certain challenges and limitations to end users which
115 can be mitigated through data fusion. Instantaneous snapshots are obtained at a
116 large number of spatial and temporal fields of regard determined by orbital and sensor

117 geometry, and therefore do not fall on a regular grid. Data coverage is spatially
 118 and temporally incomplete due to clouds, gores (spaces between orbit tracks), and
 119 faults due to “single-event upsets” often attributed to cosmic rays. L2 data can have
 120 large errors relative for example to reanalysis datasets, and uncertainty estimates,
 121 if reported, may significantly underestimate or overestimate the error relative to a
 122 reference dataset.

123 Our fused NSAT product combines two input remote sensing datasets: L2 NSAT
 124 from AIRS, and L2 NSAT from the Cross-track Infrared Microwave Sounding Suite
 125 (CrIMSS) on the Suomi National Polar-orbiting Partnership (SNPP) platform. These
 126 L2 datasets are created using two independent retrieval algorithms with different first-
 127 guess strategies. We also use information content from in situ ground station networks
 128 from NOAA’s Integrated Surface Database (ISD) to determine uncertainties in the two
 129 remote sensing datasets which are needed to perform fusion, and to validate the SSDF
 130 product and its associated uncertainty estimates. We randomly divide the ISD data
 131 into training and testing sets to perform these two separate functions.

132 Our fused NSAT product has the following key advantages over either of the
 133 input remote sensing datasets:

- 134 1. filled spatial gaps;
- 135 2. regular 0.25-degree spatial gridding;
- 136 3. reduced bias and variance relative to a reference in situ dataset;
- 137 4. improved uncertainty estimates;
- 138 5. improved long-term stationarity.

139 The rest of the paper is organized as follows. We first describe the input datasets
 140 and methodology. Then we present the fused NSAT product, and the results of vali-
 141 dation against withheld ISD surface station data. We also compare the fused NSAT
 142 product to the individual input remote sensing datasets, and to ERA5-Land reanaly-
 143 sis. In the process of validating our fused product, we also produce the most thorough
 144 validation study to date of the AIRS V7 and SNPP-CrIMSS-CLIMCAPS V2 NSAT
 145 products over CONUS. We conclude with a discussion of advantages, limitations, and
 146 potential future work.

147 **2 Data and methods**

148 Our fusion procedure involves five major steps: (1) Obtaining and filtering input
 149 remote sensing datasets that estimate the same quantity; (2) Matching the remote
 150 sensing datasets to a reference in situ dataset in space and time; (3) Using these
 151 matched data pairs (“matchups”) to characterize the input datasets via estimation of
 152 their bias and variance relative to the reference estimate; (4) Performing the SSDF
 153 calculations; and (5) Validating the results using withheld data from the reference
 154 dataset. The method and the specific datasets used in our NSAT dataset are described
 155 in the following subsections.

156 **2.1 Satellite NSAT data**

157 The input satellite datasets come from two hyperspectral infrared sounders and
 158 retrieval algorithms. The Aqua platform that carries AIRS launched in 2002 in a
 159 sun-synchronous polar orbit, with equator crossing times of approximately 1:30 P.M.
 160 and 1:30 A.M. for ascending (south to north) and descending (north to south) nodes,
 161 respectively. AIRS is an infrared grating spectrometer with 2378 channels, spanning
 162 3.7 to 15.4 μm (Chahine et al., 2006). Power to critical channels of the Aqua satellite’s
 163 Advanced Microwave Sounding Unit (AMSU)-A2 was lost in September 2016 (Yue et

164 al., 2017). AMSU-A2 complemented the AIRS instrument in atmospheric temperature
 165 and moisture profile retrievals, and was especially informative for moisture profiles.
 166 The Cross-track Infrared Sounder (CrIS) and the Advanced Technology Microwave
 167 Sounder (ATMS) instruments launched onboard the SNPP platform in 2012. SNPP
 168 is in the same orbital plane as Aqua, but at a higher altitude (824 km as opposed
 169 to 705 km), with equator crossing times also approximately 1:30 P.M. and 1:30 A.M.
 170 Together, these two instruments are known as SNPP-CrIMSS (Cross-track Infrared
 171 Microwave Sounder Suite). SNPP-CrIS experienced an anomaly on May 21, 2021
 172 which resulted in the loss of the longwave infrared channels. Another instance of
 173 CrIMSS is flying on the JPSS-1 (Joint Polar Satellite System, also known as J1 or
 174 NOAA-20) which launched on November 2017. Data from J1-CrIMSS is not used in
 175 this study, but could be used in future SSDF products.

176 For obtaining Aqua-AIRS temperature soundings, we use the AIRS-team Version
 177 7 L2 “infrared-only” temperature retrieval algorithm (Susskind et al., 2014), a least
 178 squares estimate using singular value decomposition regularization and cloud-cleared
 179 radiances. Stochastic Cloud Clearing Neural Network (SCCNN) which is trained to
 180 ECMWF fields (Blackwell, 2005) as a first guess, then refines to a final estimate. We
 181 choose the “infrared-only” retrieval for our study due to the 2016 loss of AMSU-A2,
 182 but we note that this retrieval uses information from the satellite’s other microwave
 183 sounder, AMSU-A1 (Yue et al., 2020). The retrieval uncertainty is estimated via a
 184 regression model using eleven retrieval diagnostic quantities as predictors; the regres-
 185 sion coefficients are trained on two days of retrievals (9/29/04 and 2/24/07) using
 186 ECMWF 3-hour forecasts as a reference dataset (Susskind et al., 2014; Thrastarson
 187 et al., 2020). Each individual retrieval has a nominal horizontal resolution of 45 km
 188 comprised of nine 15 km fields of view in a 3x3 matrix, and each swath contains 30
 189 retrievals across its width and 45 along track. The product is organized nominally in
 190 240 “orbital granules” per day (AIRS Project, 2020).

191 For obtaining SNPP-CrIMSS temperature soundings, we use the Community
 192 Long-term Infrared Microwave Coupled Atmospheric Product System (CLIMCAPS)
 193 Version 2 L2 temperature retrieval, which uses a hybrid optimal estimation methodol-
 194 ogy with a first guess from the Modern-Era Retrospective Analysis for Research and
 195 Applications version 2 (MERRA2) (N. Smith & Barnet, 2020), and information from
 196 both the CrIS and ATMS instruments. Like the AIRS-team retrieval, CLIMCAPS
 197 uses nine approximately 15 km fields of view in a 3x3 field of regard of 45 km, and
 198 performs cloud clearing using L1 radiances. CLIMCAPS uncertainty is estimated and
 199 propagated sequentially via error covariance matrices in stages (N. Smith & Barnet,
 200 2019). CLIMCAPS produces a combined infrared and microwave retrieval at two
 201 spectral resolutions: Nominal Spectral Resolution (NSR) and Full Spectral Resolution
 202 (FSR). We use the CLIMCAPS-SNPP NSR product to create our SSDF product, since
 203 it begins in 2012 whereas the FSR record only begins on November 2, 2015. In what
 204 follows, we refer to this product as “CrIMSS-CLIMCAPS” or sometimes as “CrIMSS.”
 205 An overview of the AIRS-team and CLIMCAPS retrievals is available online (AIRS
 206 team, n.d.), and a detailed comparison of the two retrievals applied to AIRS L1 data
 207 is available, including relative strengths and weaknesses can be found in (Yue et al.,
 208 2021).

209 The CLIMCAPS retrieval is also applied to Aqua-AIRS radiances. For this pilot
 210 fused NSAT product, we chose to use the AIRS-team retrievals for the Aqua-AIRS L2
 211 input data to demonstrate the use of different L2 retrievals as input datasets.

212 NSAT is obtained from the vertically-resolved temperature profiles (with 100
 213 pressure levels) via interpolation to the surface pressure for each field of regard (Olsen
 214 et al., 2017). The profile temperatures immediately above and below the surface
 215 are used for the interpolation, unless the level above is within 5 hPa of the surface
 216 pressure. In that case, the two levels above the surface are used. We include only L2

217 NSAT retrievals from AIRS V7 IR-only and CrIMSS-CLIMCAPS products with data
 218 quality flags ‘good’ or ‘best.’

219 **2.2 In situ NSAT data**

220 The National Oceanic and Atmosphere Administration (NOAA) Integrated Sur-
 221 face Database (ISD) is a global database of near-surface meteorological observations
 222 compiled from over a hundred systems of ground stations (A. Smith et al., 2011). The
 223 record extends back to the 1950s, although new stations have been added on a con-
 224 tinual basis as available, improving coverage over time. Today ISD consists of more
 225 than 35,000 surface weather stations globally, 14,000 of which remain active. Figure 1
 226 shows the spatial coverage of ISD stations in North America.

227 We use sub-hourly NSAT measurements gathered from over 7000 stations in
 228 North America as our reference dataset, for bias and variance estimation and for valida-
 229 tion. No data are perfect, but the ISD errors are small relative to the errors in the input
 230 remote sensing datasets (see Figure 7). Naturally ventilated screened surface station
 231 air temperature measurements are accurate to ± 0.1 K in most circumstances (Harrison
 232 & Burt, 2021). ISD data come with a set of ten data quality flags, indicating various
 233 problems and levels of quality. We only use ISD data flagged as highest quality, i.e.,
 234 data must be flagged with either 1 (‘Passed all quality control checks’) or 5 (‘Passed
 235 all quality control checks, data originate from an NCEI data source’).

236 We chose ISD ground stations as our reference dataset for the following reasons:
 237 (1) it is not reanalysis, which assimilates AIRS and SNPP-CrIMSS information, as
 238 well as information from dynamical weather modeling; (2) ISD is among the most
 239 comprehensive ground station datasets available over land; (3) ISD NSAT estimates
 240 have low errors relative to remote sensing estimates.

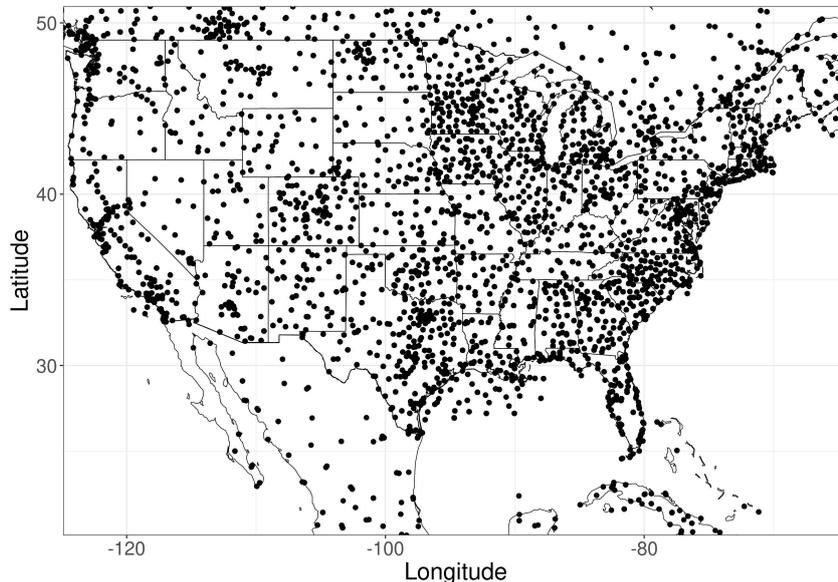


Figure 1: Spatial coverage of the ISD stations over North America. Note that ISD also includes stations elsewhere in the world.

241

2.3 Reanalysis NSAT data

242

243

244

245

246

247

248

249

We also compare the fused NSAT results to ECMWF Reanalysis 5 (ERA5)-Land reanalysis data. The ERA5 is the fifth-generation global atmospheric reanalysis from ECMWF, replacing the ERA-Interim reanalysis which stopped being produced on August 31, 2019. Newly reprocessed datasets along with recent instruments have been assimilated into the ERA5 that could not be ingested into the ERA-Interim (Hennermann & Berrisford, 2019). We note that some AIRS spectral channels under clear conditions are incorporated into ECMWF reanalysis (McNally et al., 2006), but that ISD data are not.

250

251

252

253

We use hourly ERA5-Land output which is a high-resolution version of the land component of the ERA5 reanalysis. ERA5-Land 2 m air temperature was chosen over the full ERA5 reanalysis for its finer spatial resolution of 0.1×0.1 degrees and hourly temporal resolution.

254

2.4 Bias and variance estimation

255

256

257

258

259

260

Biases and variances of input data sources are the key to high-quality data fusion. SSDF assumes input data are unbiased relative to some reference dataset, and weights them by the inverse of their respective variances. This minimizes output errors of the fused estimates relative to the reference dataset. Therefore, data must be bias-corrected before SSDF ingestion, and the quality of the final fused product depends on the quality of uncertainty estimates for the inputs.

261

262

263

264

265

266

267

268

269

270

271

272

273

To estimate bias and variance for satellite footprints, we create an ensemble of “matchups”: matched pairs of satellite and ISD station estimates that are close in space (less than 100 km apart) and time (less than an hour apart). For a given period, the matchups are sorted into 240 km (\sim two-degree) diameter hexagonal spatial bins based on satellite footprint location, with three-day time bins (day of interest, along with preceding and following days). We empirically tested different time bins (monthly, seven days, and three days) for aggregating matchups for determining bias and variance, and the three-day time bins minimized the mean standard deviation of a sample SSDF product over CONUS, while allowing for adequate sample size. This binning is the basis for quantifying bias and variance for all satellite footprints in a given space-time cell. We randomly select 1% of the ISD matchup pairs to withhold for validation (we do not withhold entire ISD stations). We chose a relatively small amount to withhold in order to maximize the information content for the SSDF product.

274

To obtain the matchups we apply the following steps.

275

276

277

278

279

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

280

281

282

283

284

Note that these steps will result in a one-to-one match between an ISD observation and a single AIRS footprint. Some ISD observations may have no corresponding AIRS match, in which case no matchup is returned. We next tessellate a fixed hexagonal spatial grid over CONUS and find the biases and variances using matchups aggregated over 3 days within each grid cell, as follows:

285

286

287

- I. To compute a bias on day d and mode j (day or night) and in hexagonal grid cell i , we find the set of all valid (i.e., non-null) AIRS-ISD matchups from Steps 1 to 3 above such that,

- 288 (a) the AIRS data come from mode j ,
 289 (b) the AIRS footprint belongs within the grid cell i ,
 290 (c) the ISD date is in $(d - 1, d, d + 1)$.
- 291 II. The bias and variance for day d , mode j , and grid cell i are then computed
 292 using the set of paired ISD-AIRS matchups.

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

298 After the bias field is estimated for a given dataset relative to the ISD reference
 299 dataset, every datum in that dataset is then bias-corrected. After the variance field
 300 is estimated for a given dataset, every datum in that dataset is assigned a variance
 301 estimate which is then used in the SSDF algorithm to weight the datum.

302 2.5 Data fusion methodology

303 SSDF is an algorithm for fusing multiple remote sensing datasets by leveraging
 304 spatial dependence in the data, also known as kriging or optimal interpolation (Cressie,
 305 1993). Remote sensing data from different instruments in general are heterogeneous.
 306 By this we mean that the input remote sensing data sets may have different spatial
 307 footprints, sampling patterns, and measurement error characteristics. SSDF accounts
 308 for these heterogeneities by using a spatial statistical model that expresses the relation-
 309 ships between the quantity of interest at a particular location, and all the observations
 310 at all locations from all data sources.

311 We note that the main requirement of SSDF is that the different instruments in
 312 question (e.g., AIRS and CrIMSS) must be observing the same geophysical quantity of
 313 interest (e.g., NSAT). We assume that after bias correction, the retrievals from both
 314 instruments are unbiased relative to the reference dataset. We also assume that we
 315 have standard deviation estimates that characterize the relative informational content
 316 between the instruments.

317 One of the challenges encountered when applying spatial interpolation via tradi-
 318 tional kriging to remote sensing data is the massive data sizes involved. In traditional
 319 kriging, the computational complexity of the algorithm is $O(N^3)$ due to the need to
 320 invert an $N \times N$ covariance matrix \mathbf{C} , where N is the number of data points. This
 321 inversion makes traditional kriging infeasible for datasets with N on the order of tens
 322 of thousands of data points or larger. To account for this, we use a scalable vari-
 323 ant of kriging that employs a dimension-reduction technique (Spatial Random Effects
 324 modeling) to parameterize the matrix \mathbf{C} as a rank- r update to a diagonal matrix,
 325 where $r \ll N$. This allows us to invert the covariance matrix \mathbf{C} analytically us-
 326 ing the Sherman-Morrison-Woodbury formula with computational complexity $O(Nr^2)$
 327 (Cressie & Johannesson, 2008). SSDF is essentially an extension of Fixed-Ranked Krig-
 328 ing (FRK) for combining multiple datasets. Indeed, SSDF works by concatenating all
 329 the datasets into a meta-dataset (with each data point encoded with a value, location,
 330 and variance estimate) and then applying the FRK algorithm. Therefore, SSDF can
 331 easily generalize to more datasets than two, and it can also be applied to a single
 332 dataset (a sub-case needed for the AIRS-only part of the multi-instrument record,
 333 from 2002-2012), without mathematical modification.

334 A second challenge with traditional kriging is handling arbitrary spatial foot-
 335 prints of the input datasets and those of the output grid. Gotway and Young (2002)
 336 identified this “change of support” problem of inferring a spatial process at one res-

337 olution from data at another resolution. However, their solution is computationally
 338 intensive, requiring integration over footprints and making it difficult to do parameter
 339 estimation for general non-linear covariance classes. In SSDF the SRE model is linear,
 340 which makes change of support and the associated parameter estimation straightforward
 341 (Nguyen et al., 2012).

342 As a scalable variant of Gaussian process prediction (Cressie, 1993), SSDF pro-
 343 vides two other advantages over other non-statistical data fusion approaches such as
 344 binning or non-parametric methods such as machine learning. First, the standard er-
 345 rors are optimized because SSDF minimizes errors relative to the unknown process;
 346 SSDF estimates are therefore “best linear unbiased estimates.” Within the class of
 347 linear estimators, this method produces the smallest prediction errors. In addition,
 348 SSDF provides a statistically principled method for estimating uncertainties. Mini-
 349 mizing errors and quantifying uncertainties allows SSDF to create more accurate and
 350 usable data products from input datasets.

351 For the full mathematical formulation of SSDF, see Appendix B.

352 **2.6 Dataset preparation for assessment**

353 We assess our SSDF fused product using a randomly chosen reserved 1% of the
 354 ISD dataset. The assessment focuses on the product which is produced by the end-to-
 355 end workflow, including bias and variance estimation, bias correction, and the SSDF
 356 procedure. We match up SSDF, AIRS, CrIMSS, and ERA5 estimates to withheld ISD
 357 data using a 100 km and 1 hour matchup criterion (see Section 2.4 for more detail).
 358 These matchup datasets generally differ in their coverage; for instance, a fused estimate
 359 might be matched to an ISD observation at a location where there are no nearby AIRS
 360 or CrIMSS estimates. Therefore, to mitigate the effect of biases due to differing spatial
 361 and temporal coverage in these matchup pairs, we require that fused estimates are also
 362 close to (within the same matchup distance and time) of at least one datum from the
 363 comparison dataset. This matchup procedure generates multiple paired datasets: ISD-
 364 AIRS, ISD-CrIMSS, ISD-SSDF, and ISD-ERA5, allowing comparison, for example, of
 365 pairs of datasets such as AIRS and SSDF (AIRS) (i.e., a subset of the fused points
 366 matched up to AIRS points) which have the same number of samples, each of which is
 367 collocated in space and time within the matchup criterion. To put this another way,
 368 the reason we have separate plot traces for SSDF(AIRS) and SSDF(CrIMSS) is to
 369 allow an apples-to-apples comparison despite differing spatial coverage of the AIRS,
 370 CrIMSS, ERA5, and SSDF datasets.

371 The choices of a 1% test ISD dataset and this matchup scheme results in over
 372 4000 AIRS-SSDF sample pairs and over 13,000 CrIMSS-SSDF sample pairs for 2013,
 373 a typical year.

374 **3 Results**

375 **3.1 SSDF product overview**

376 We produced fused NSAT using two satellite input datasets over North America
 377 between 25 N and 50 N. We chose to fuse the AIRS and SNPP-CLIMCAPS products
 378 because the orbits of these satellites have similar overpass times of approximately 1:30
 379 and 13:30 local solar time, and the records extend back to at least 2013. We note
 380 that although we initially restrict our product to CONUS, the two input L2 retrievals
 381 provide global coverage, and that we plan to extend our fused product to global land
 382 surfaces in the future in regions with adequate reference (ISD) data coverage. We
 383 produce two products, a main product from both AIRS and SNPP-CrIMSS which
 384 runs from November 28 2012 through 2020 and which we will denote SSDF-AC; and

385 a long-record product with just AIRS, which runs from August 31 2002 through 2020
 386 and which we will denote SSDF-A. These two product lines were created identically,
 387 with the only difference being that the list of input data tuples (bias-corrected NSAT,
 388 latitude, longitude, and variance) fed to the SSDF algorithm consisted of tuples from
 389 either two remote sensing datasets or just one. Between 2013 and 2020 there were 32
 390 days and 30 nights with no AIRS data, and 29 days and 24 nights with no SNPP-
 391 CLIMCAPS data. Because outages happened not to occur for both input datasets on
 392 the same day or night over this period, the SSDF-AC product was created from only
 393 the single dataset when necessary, thus creating a continuous record. The SSDF-A
 394 record has 74 missing daily files due to AIRS outages, often due to single event upsets
 395 (for a list of AIRS outages, see <https://airs.jpl.nasa.gov/data/outages/>). In
 396 what follows, if not otherwise specified, “SSDF” refers to SSDF-AC.

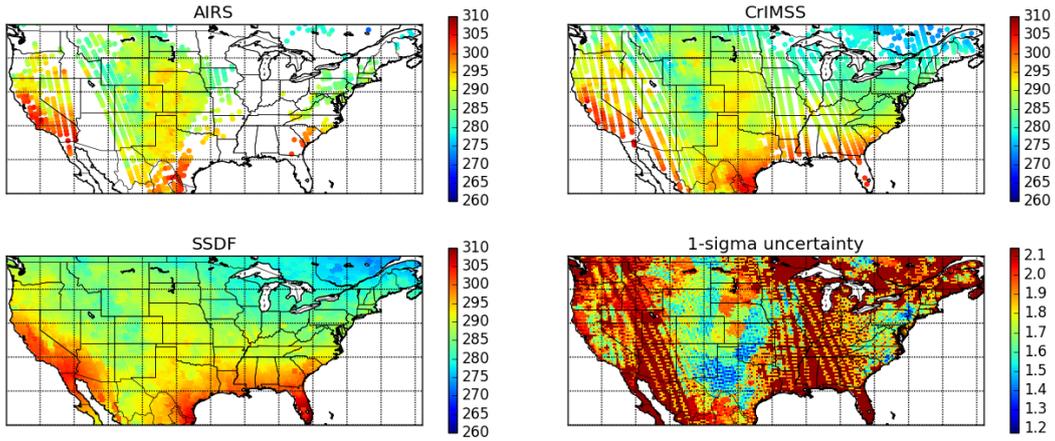


Figure 2: Sample data fusion satellite NSAT inputs, SSDF fused NSAT results, and uncertainty estimates for 2015 October 31, day. The top two plots show maps of the input satellite NSAT data ingested into the SSDF product (restricted to CONUS and neighboring regions), with AIRS on the left and SNPP-CrIMSS 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 degrees K.

397 Figures 2 and 3 provide maps representing one arbitrarily chosen day and night
 398 of the SSDF-AC product. For both the day and night cases, the top two plots show
 399 maps of the input satellite data ingested into the SSDF product, with AIRS on the left
 400 and SNPP-CrIMSS on the right; the bottom left plot shows the SSDF fusion results;
 401 and the bottom right plot shows the uncertainty estimates on the SSDF fusion results
 402 at the 1-sigma level. These sample maps demonstrate how our SSDF method fills
 403 in missing data in the input datasets by exploiting spatial correlations to provide a
 404 complete gap-filled, gridded product. Note that the estimated uncertainties are higher
 405 in regions that contain no observations, contain observations from only a single input
 406 dataset, or in which the two input datasets have relatively poor agreement.

407 3.2 Comparison of bias, standard deviation, and RMSE

408 We now turn to validation against withheld ISD reference data to quantify im-
 409 provement in the SSDF products. We emphasize that the ISD data used for validation
 410 were not the same as the ISD data used to estimate bias and variance in the course
 411 of creating the SSDF products, as we split the ISD matchup data into ‘training’ and
 412 ‘testing’ sets. We examine bias, standard deviation, and RMSE, calculated from the

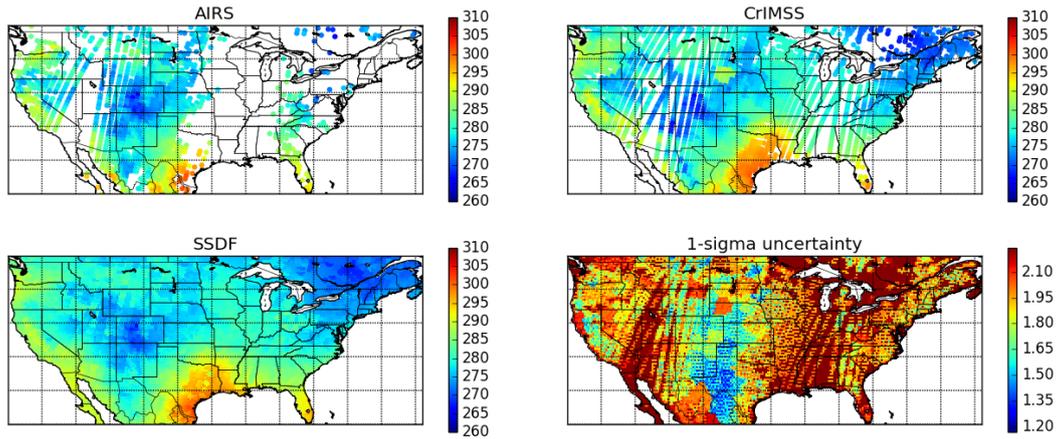


Figure 3: Same as Figure 2 but for night. All units are degrees K.

413 withheld matchups, of AIRS, CrIMSS, ERA5-Land, and the corresponding matched
 414 SSDF data. In what follows, we often analyze daytime and nighttime separately, as
 415 daytime and nighttime biases can differ significantly.

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

421 Overall, in the mean over CONUS and over the entire time period, our procedure
 422 (bias correction and data fusion) provides a reduction in the magnitude of daytime
 423 bias of 1.7 K and 0.5 K relative to AIRS and CrIMSS, respectively. At night, the fused
 424 product is essentially unbiased in the mean over the domain (relative to the reference
 425 dataset) and provides a reduction in the magnitude of bias of 1.5 K and 0.2 K relative
 426 to AIRS and CrIMSS, respectively.

427 AIRS shows a strong cold bias in daytime over the mountainous West, which is
 428 also present in CrIMSS, although less severe. AIRS shows a near-constant warm bias
 429 over the entire Eastern CONUS at night, while CrIMSS shows a sharp warm bias over
 430 small regions of the mountainous West at night. Our procedure mitigates these biases
 431 (through the bias-correction procedure described above) and produces estimates with
 432 lower biases than either of its input satellite data sets over the domain.

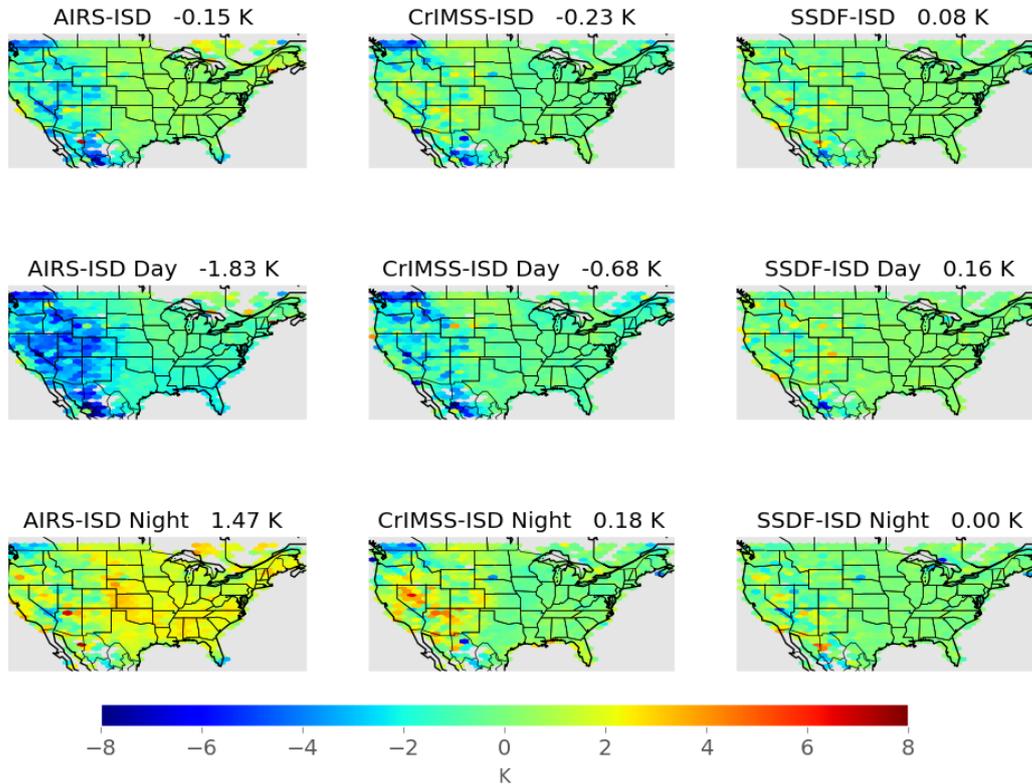


Figure 4: Maps of bias (retrieval - ISD) over the product period of 2013-2020, created against the withheld ISD test data, for AIRS (first column), CrIMSS-CLIMCAPS (second column) and SSDF (third column), for both day and night together (top row), for day only (second row) and for night only (third row). 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.

433 Figures 5 and 6 show maps of standard deviation and RMSE for AIRS, CrIMSS
 434 and SSDF, for the 2013-2020 period, and for daytime only and nighttime only. Stan-
 435 dard deviation and RMSE tell a similar story to that of bias. Overall, in the mean
 436 over CONUS and over the entire time period, SSDF provides a reduction in RMSE of
 437 35% and 15% compared to AIRS and CrIMSS, respectively.

438 CrIMSS has high RMSE over the mountainous West in both day and night,
 439 but low RMSE over the eastern two-thirds of the continent. Similarly, AIRS has
 440 relatively high RMSE over the entire domain, but especially over the mountainous
 441 West. Mountainous regions pose particular challenges for remote sensing of surface
 442 quantities, and of NSAT in particular, which can vary greatly depending on e.g., north-
 443 facing versus south-facing mountain surfaces. Furthermore, variations in topographic
 444 features between ISD stations and their matched remote sensing retrievals can lead to
 445 random errors, increasing RMSE and variance estimates. However, the fused NSAT
 446 product shows a clear decrease in bias over all regions, including in the mountainous
 447 western CONUS, although there is potential for improvement in the SSDF product
 448 over the West.

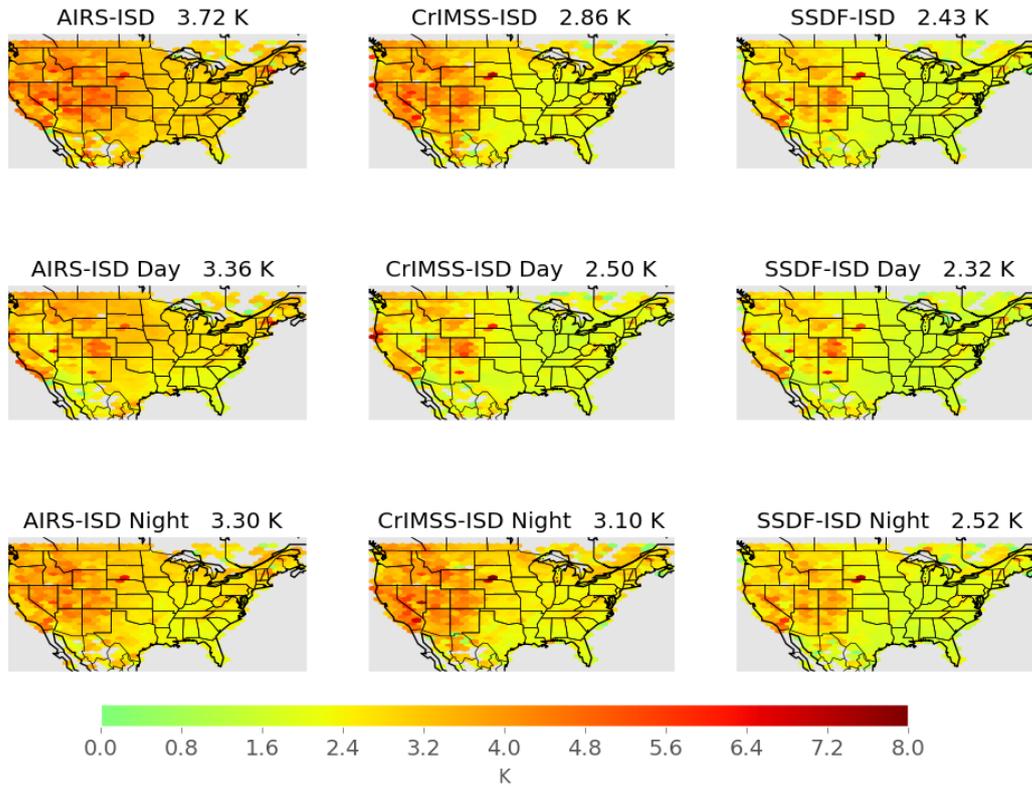


Figure 5: Standard deviation maps. The nine panels are similar to those in Figure 4 but for standard deviation.

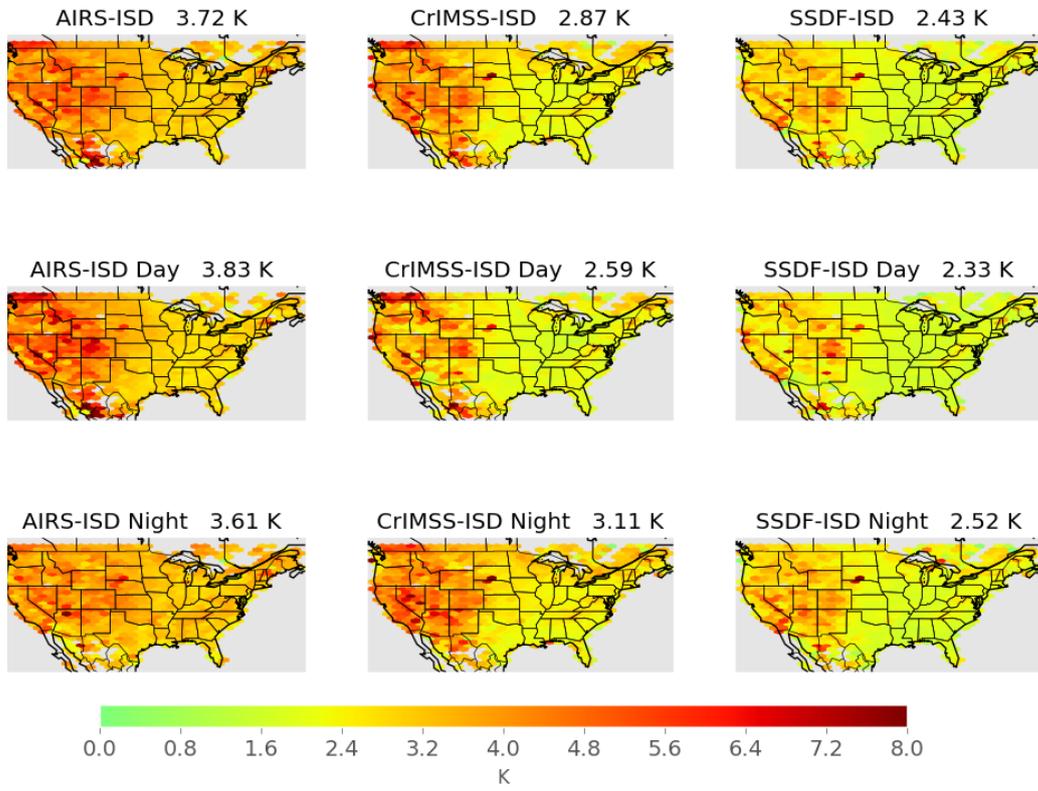


Figure 6: RMSE maps. The nine panels are similar to those in Figure 4 but for RMSE.

449 We repeated this analysis over CONUS and the 2013-2020 period for the SSDF-
 450 A product. We found similar improvements in bias, standard deviation, and RMSE.
 451 The mean bias of SSDF-A over the entire domain was -0.08 K for daytime only, and
 452 -0.03 K for nighttime only. The overall RMSE was 2.52 K, 4% higher than the overall
 453 RMSE of the SSDF-AC product.

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

468 Next, we examine the seasonality of bias and RMSE. Figure 8 shows the mean
 469 bias (retrieval/reanalysis - ISD) by month split into day/night to examine seasonality.
 470 There is a significant cold bias during the day for AIRS and CrIMSS that switches
 471 to a warm bias at night. During the day, AIRS has a smaller bias during winter
 472 months (Dec/Jan/Feb) and a larger bias during summer months (Jun/Jul/Aug). This
 473 is switched during nighttime where a larger warm bias is observed during winter and
 474 a smaller warm bias is observed during summer. These AIRS biases are of course also
 475 apparent in Figure 7. The SSDF product is relatively unbiased for both day and night.
 476 The SSDF bias magnitude is slightly larger during the day than night. From May to
 477 December, the SSDF product has a smaller bias at night than does ERA5-Land while
 478 during the day the reanalysis and the SSDF mean biases are of similar magnitude.

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

484 We next examine relative performance in hot and cold extremes. Figure 10 shows
 485 the mean bias (retrieval/reanalysis - ISD) by ISD percentile of the ISD matchups.
 486 The error bars indicate the standard error of the mean at the 95 percent confidence
 487 level. The lighter shade of every color is the matched SSDF corresponding to the
 488 comparison dataset. All retrievals and reanalysis do best in the mean state (25th to
 489 75th percentile). At the extremes, each of the datasets being compared to ISD have
 490 warm biases for low values (1st through the 15th percentile) and cold biases for high
 491 values (85th through the 99th); in other words, all of the datasets understate cold or
 492 warm extremes represented in the ISD. This is perhaps to be expected, as the ISD
 493 dataset consists of point measurements which capture fine-scale extremes, whereas the
 494 satellite datasets represent spatial means over scales ranging from about ~ 50 km at
 495 nadir to ~ 150 km at the edge of scan.

496 The SSDF product captures the extremes better than both the AIRS and CrIMSS
 497 inputs. However, the reanalysis generally does best, having the smallest bias regardless
 498 of percentile, and is better at capturing the extremes.

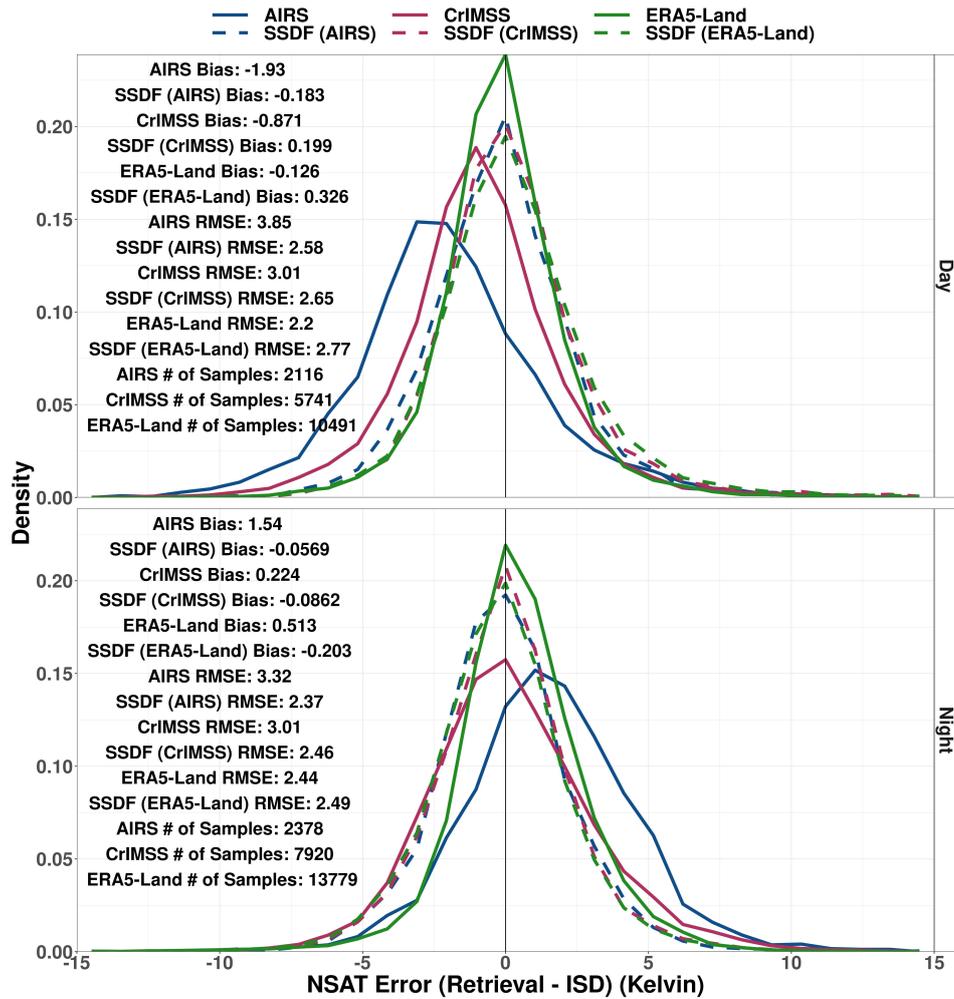


Figure 7: Histograms of errors for day (top) and night (bottom) for 2013 over CONUS, for AIRS (blue), CrIMSS (red) and ERA5-Land (green). The dashed line is the SSDF-AC subset matched to the other datasets. Mean statistics of bias, RMSE, and the number of samples are provided.

499 We next examine performance at extremely high elevations. Figure 11 shows
 500 mean biases (retrieval/reanalysis – ISD) aggregated by ISD elevation. At around 2500
 501 meters, mean biases increase with elevation in the SSDF product, AIRS, CrIMSS, and
 502 reanalysis. Daytime mean biases at these high elevations are larger in SSDF, although
 503 we note that the sample size is small. At night, SSDF shows lower mean biases than
 504 AIRS, CrIMSS, or ERA5-Land at high elevations.

505 In order to increase the sample size for high-elevation cases, Figure 12 shows
 506 the mean biases aggregated by ISD elevation for elevations higher than 2000 meters
 507 over the period 2012-2020. During the day, the SSDF bias exceeds AIRS and CrIMSS,
 508 consistent with Figure 11. We hypothesize that this excess bias in SSDF for a very small
 509 number of data points at very high elevations is caused by the bulk-binning method
 510 for bias estimation. As Figure 11 shows, both remote sensing datasets exhibit a cold
 511 bias during the daytime at lower elevations. Because the two-degree hexagonal bins for
 512 bias estimation are dominated by lower elevations (as the problematic high elevations

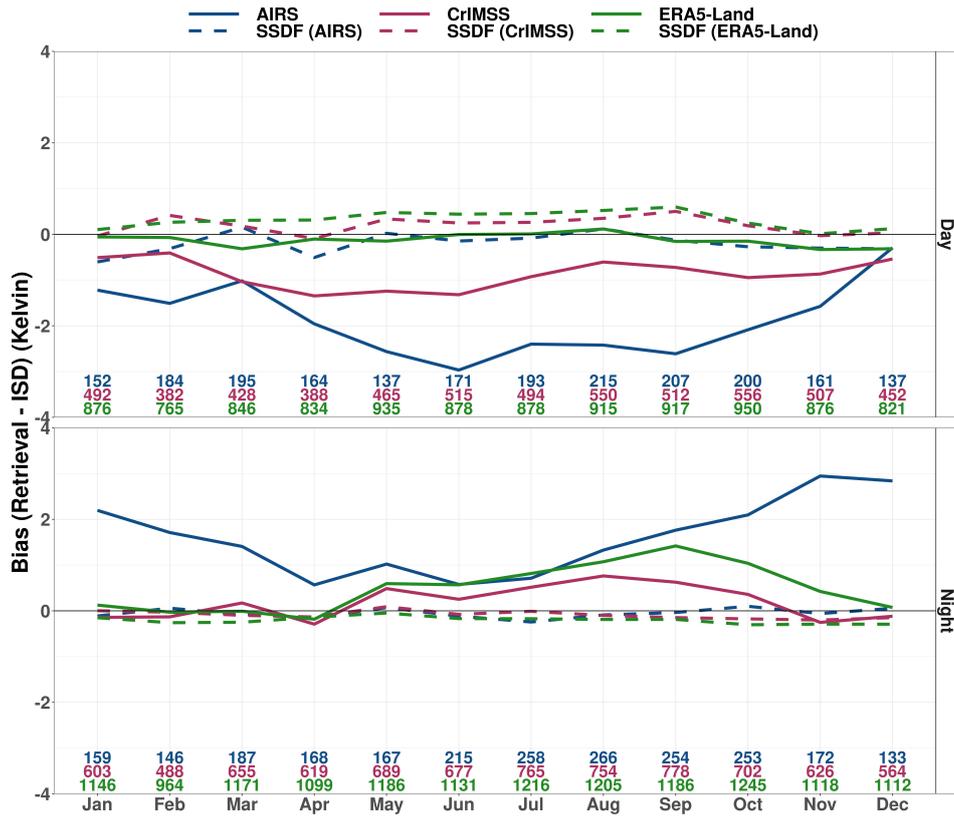


Figure 8: Mean bias as a function of month for day (top) and night (bottom) for 2013 over CONUS. Numbers at the bottom indicate the number of data points, and are color-coded according to dataset.

513 are high mountain surfaces), and because both remote sensing dataset biases switch
 514 signs from cold bias to warm bias at approximately 2500 m, the cold bias correction
 515 calculated from the bulk bins ends up exacerbating the warm bias from the input
 516 datasets at the highest elevations. In a future version of SSDF, we will improve the
 517 bias estimation of the input datasets, which could mitigate or eliminate this bias at
 518 the 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.

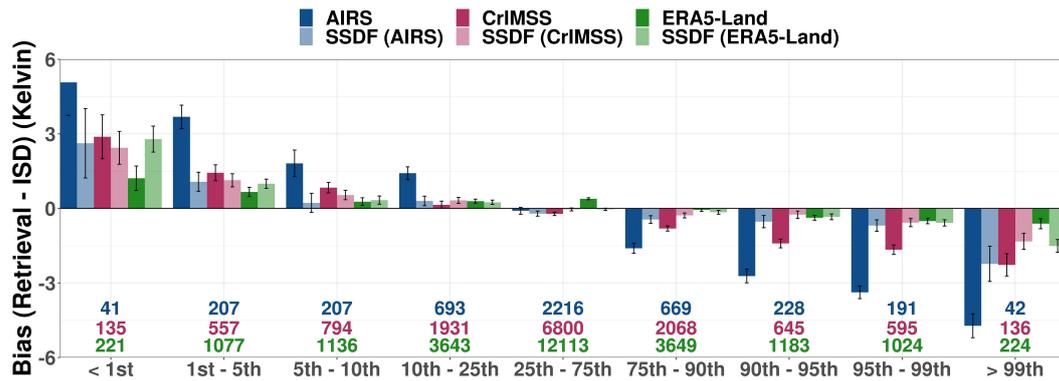


Figure 10: Mean biases as a function of ISD percentile for 2013 over CONUS. Numbers at the bottom indicate the number of data points, and are color-coded according to dataset.

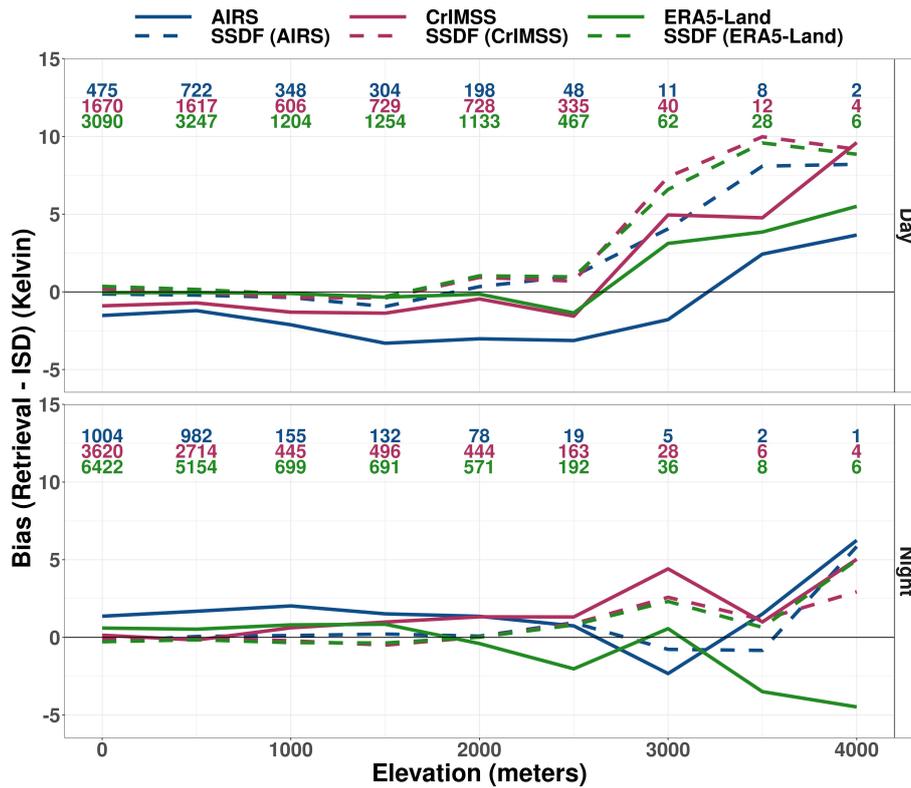


Figure 11: Mean biases as a function of ISD elevation for day (top) and night (bottom) for 2013 over CONUS. Numbers at the top indicate the number of data points, and are color-coded according to dataset.

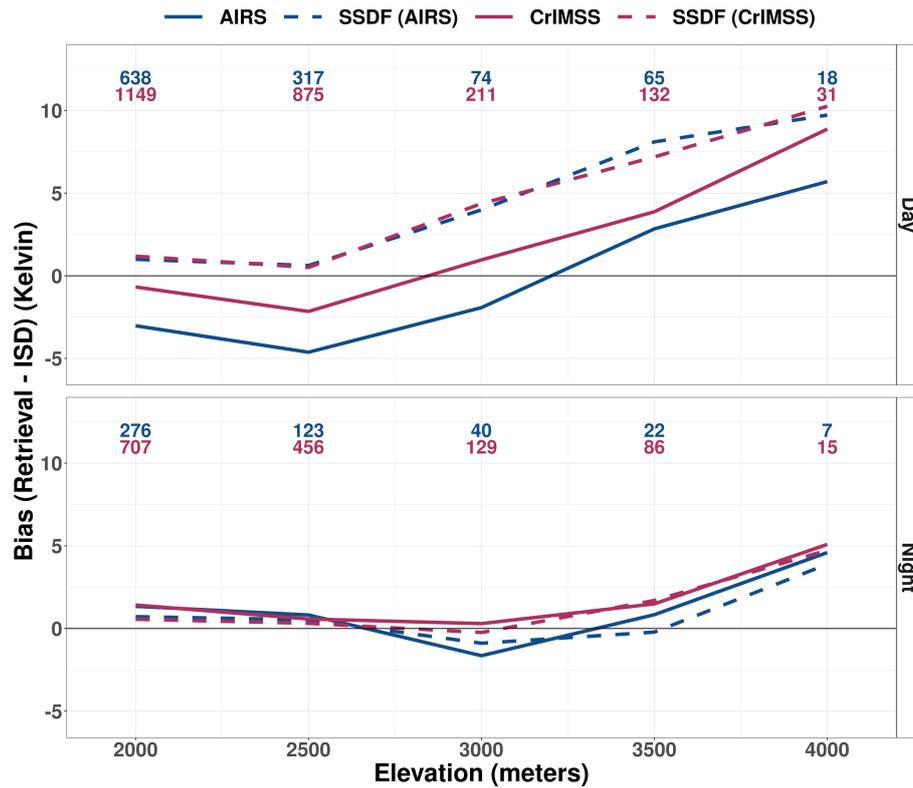


Figure 12: Mean biases as a function of ISD elevation for day (top) and night (bottom) over CONUS from 2012-2020 for AIRS, CrIMSS, and SSDF. Numbers at the top indicate the number of data points, and are color-coded according to dataset.

519

3.3 Comparison of uncertainty estimates

520

521

522

523

524

525

526

527

The SSDF algorithm provides a mean (prediction/estimate) and standard deviation (uncertainty) of the conditional distribution of 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. 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 .

528

529

In addition to each fused NSAT estimate, \hat{Y}_i , the algorithm also provides the conditional standard deviation of the predictive distribution, denoted $\hat{\sigma}_{\hat{Y},i}$.

530

531

532

533

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.

534

535

536

537

538

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 estimated “true” state given the observed radiances for a single footprint and is an output of the optimal estimation (OE) approach used in CLIMCAPS.

539

540

541

542

543

544

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.

545

546

547

548

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 assumed “true” state (as approximated by the reference dataset). For example, the error for SSDF is $e_{\hat{y},i} = \hat{Y}_i - Y_i$, where Y_i is the ISD validation for colocation i . If the predictive distribution is assumed to be Gaussian, the empirical coverage of intervals of the form

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

545

546

547

548

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 $c = 2$.

549

550

551

552

553

554

555

556

557

558

Figure 14 shows scatterplots of the joint distribution of the uncertainty estimate (x-axis) and the observed error (retrieval-*ISD*). There are cases for AIRS and CrIMSS where the uncertainty estimate underestimates the error relative to the *ISD* reference dataset; over 15% of the time for both datasets and for day and night, the 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 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.

559

560

561

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

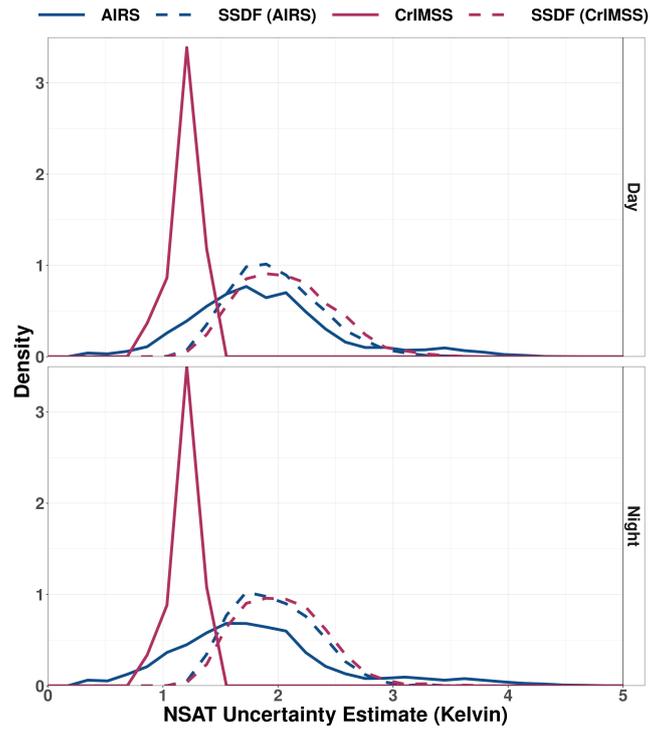


Figure 13: Histograms of uncertainty estimates for day (top) and night (bottom) for 2013 over CONUS.

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

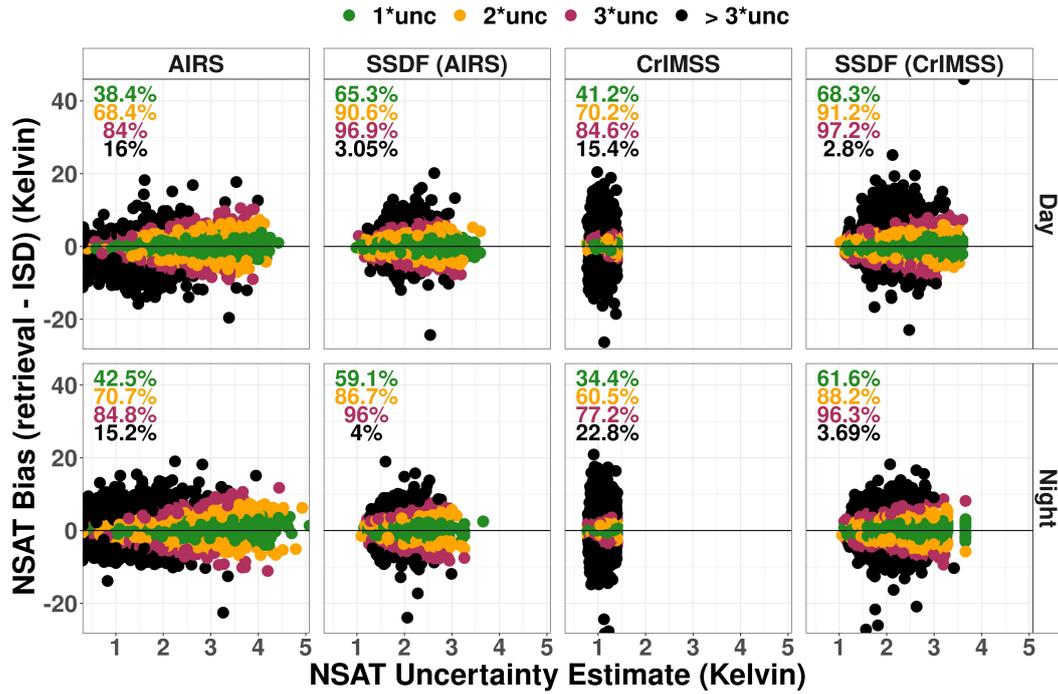


Figure 14: Observed errors (retrieval - ISD) versus uncertainty estimates for day (top) and night (bottom) for 2013 over CONUS. The colors show whether the range of each observed error was within the uncertainty bound, as described in the text: 1×uncertainty (green, should cover the error about 68% of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

565

3.4 Empirical distribution consistency

566

567

568

569

570

571

572

573

574

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.

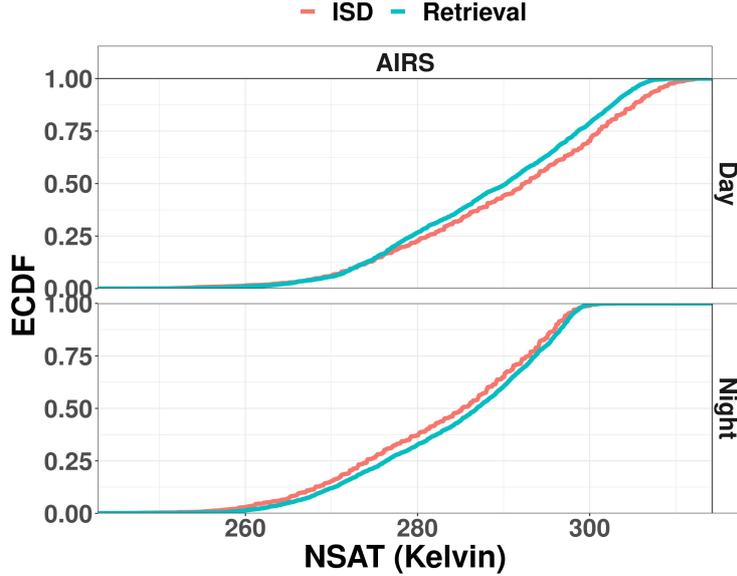


Figure 15: ECDF for AIRS (blue) and ISD (pink) for day (top) and night (bottom) for 2013 over CONUS.

575

576

577

578

579

580

581

582

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

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

583

584

585

586

where δ is the traditional two-sample KS statistic. The KS statistic is the maximum difference in the two ECDFs being compared. Thus, the test statistic γ_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.

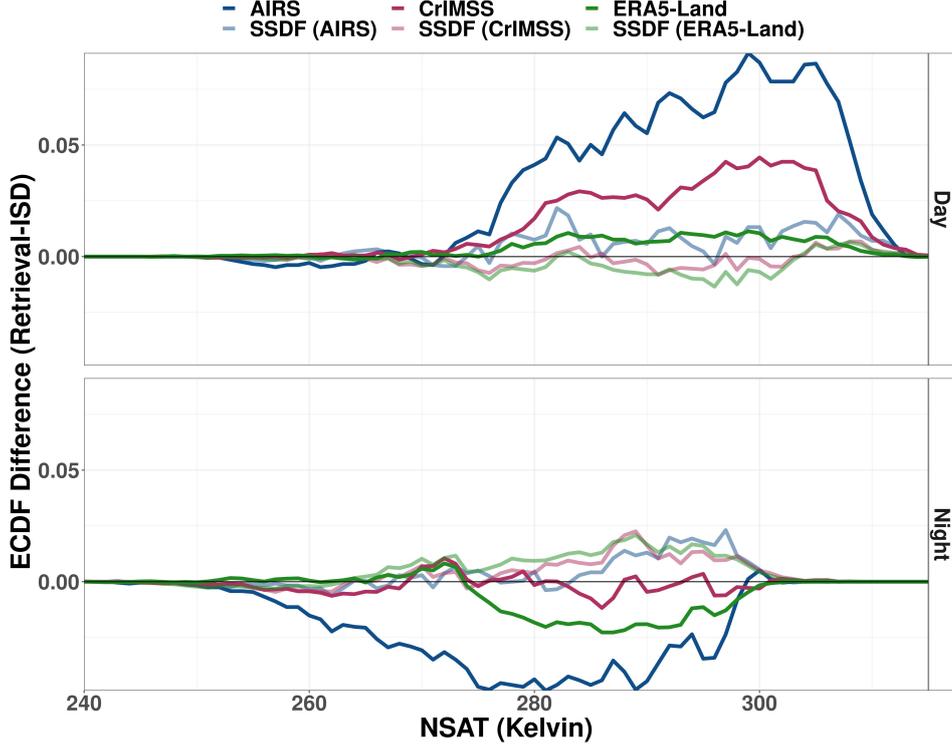


Figure 16: The ECDF difference between the retrieval/reanalysis and the ISD color coded for day (top) and night (bottom) for 2013 over CONUS.

587 The distribution of the test statistic under the null hypothesis can be estab-
 588 lished through a resampling procedure. The procedure should preserve the inherent
 589 dependence of the matched triples, but the assignment of the two comparison groups
 590 can be shuffled randomly. A null distribution is generated by repeating these steps
 591 $m = 1, \dots, M$ times:

- 592 1. Define shuffled data vectors $\mathbf{W}_{m,1}$ and $\mathbf{W}_{m,2}$.
- 593 2. For each validation matchup ($i = 1, \dots, n_k$), assign $W_{i,m,1} = \hat{Y}_i$ and $W_{i,m,2} =$
 594 $Z_{k,i}$ with probability 0.5; otherwise assign $W_{i,m,1} = Z_{k,i}$ and $W_{i,m,2} = \hat{Y}_i$. This
 595 effectively shuffles the labels for SSDF and the comparison product for each
 596 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 γ_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

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

597 where I_γ is an indicator function.

598 The p -values for each of the resampling tests of SSDF versus other products are
 599 displayed as text in Figure 17. All tests, except the night comparison of SSDF and
 600 CrIMSS, yield p -values of 0, indicating a significant difference in consistency with the
 601 ISD reference distribution. These results can also be seen visually as the observed test
 602 statistics γ_k , shown as vertical lines, lie well outside the corresponding null distribu-
 603 tions. The tests indicate SSDF is more consistent with ISD than AIRS for both day
 604 and night conditions, as well as a favorable result for SSDF versus CrIMSS for day and
 605 versus ERA5-Land at night. The positive test statistic for SSDF versus ERA5-Land
 606 during the day indicates the reanalysis is more consistent with ISD in this case.

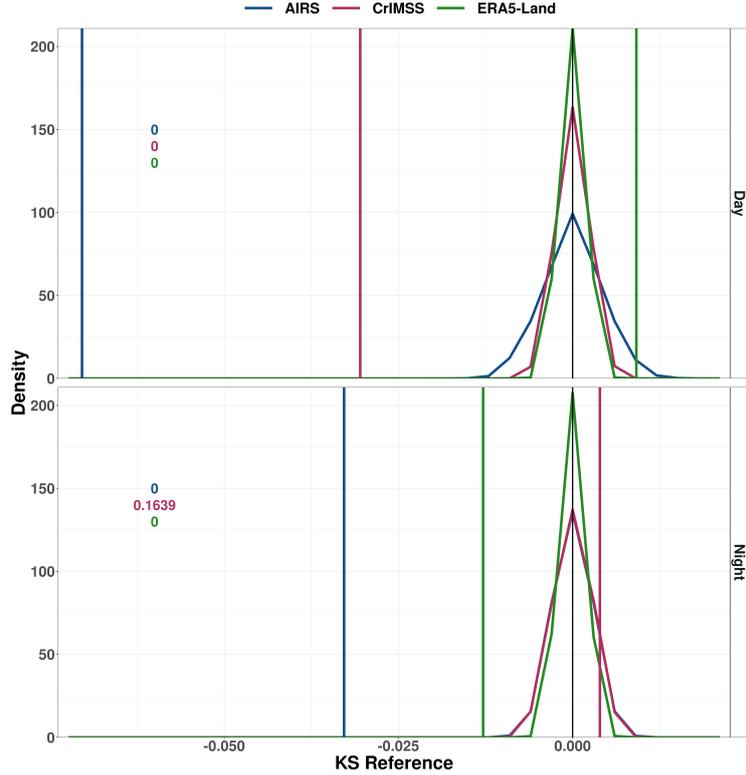


Figure 17: Histogram of the KS statistic for AIRS (blue), CrIMSS (maroon) and ERA5-Land (green), for day (top) and night (bottom) for 2013 over CONUS. The corresponding p -value is color-coded on the left side.

607 3.5 Stationarity

608 Long-term stationarity is a key characteristic for creating long, stable, multi-
 609 instrument Earth science data records. To assess long-term bias stationarity, we cal-
 610 culated mean annual biases over CONUS relative to the withheld ISD data for the two
 611 input datasets and SSDF. Figure 18 shows the annual mean bias for both the input
 612 datasets, as well as for SSDF-AC and SSDF-A. Shading shows two standard deviations
 613 of these annual bias estimates. We include full years only.

614 SSDF reduces the mean magnitude, the variance, and the trend in these annual
 615 bias time series, with the biases estimated relative to the ISD reference dataset. For
 616 AIRS and SSDF-A matched to AIRS from 2003-2020, the overall means of the annual

617 bias time series were -0.10 K and -0.035 K and the standard deviations of the annual
 618 bias time series were 0.17 K and 0.035 K.

619 We estimated trends and trend uncertainties using the nonparametric technique
 620 called Thiel Sens Slope (Sen, 1968) which is based on the medians. We used the Mann-
 621 Kendall test to assess statistical significance (Mann, 1945; Kendall, 1948). Trends
 622 for AIRS and SSDF-A were -0.01 K/yr (p -value 0.08) and -0.003 K/yr (p -value 6e-
 623 8), respectively, over the 2003-2020 period. The AIRS trend was less statistically
 624 significant due to the high standard deviation in the time series.

625 For CrIMSS and SSDF-AC from 2013-2020, the overall means of the annual bias
 626 time series were -0.23 K and 0.076 K and the standard deviations of the annual bias
 627 time series were 0.059 K and 0.024 K respectively. Trends were 0.009 K/yr and -0.0007
 628 K/yr, respectively; neither trend is statistically significant, with p -values of 0.6 and
 629 0.8, respectively.

630 The annual mean biases also reveal a shift of about 0.1 K between the SSDF-AC
 631 and SSDF-A products. This shift is small compared to the biases in the input remote
 632 sensing datasets, but it is undesirable. We hypothesize that it could be an artifact
 633 of the bulk-binning bias estimation procedure, and subsequent bias correction, due to
 634 differing systematic error characteristics in the two input datasets. Future versions of
 635 SSDF will use simulation-based uncertainty quantification methods to estimate input
 636 dataset biases (e.g., Hobbs et al., 2017; Braverman et al., 2021), which could miti-
 637 gate this difference in the mean bias between SSDF products created from different
 638 combinations of input datasets.

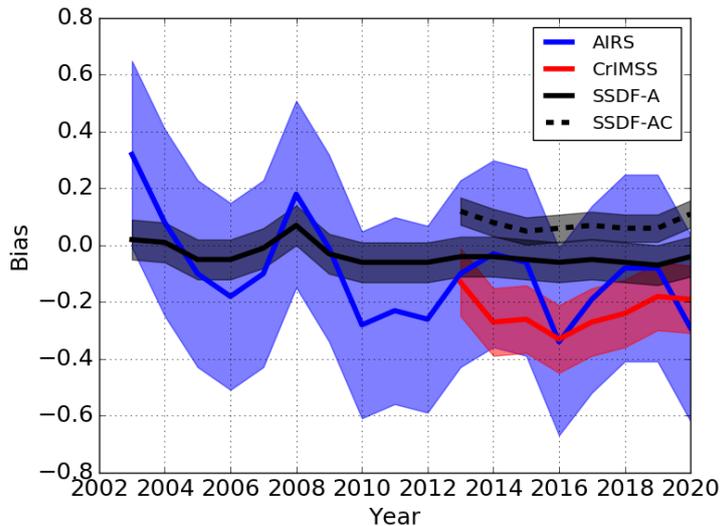


Figure 18: Annual mean bias for each year of the data record, for the SSDF product and each of the two remote sensing input products, relative to the withheld ISD data. Shading shows two standard deviations of these annual bias estimates. SSDF-A refers to the AIRS-only SSDF product; SSDF-AC refers to the SSDF product created from both the AIRS and SNPP-CLIMCAPS input datasets.

639 Figure 19 shows the histogram of the SSDF uncertainty estimates for 2011 (black)
 640 and 2013 (red). The mean uncertainty is provided as text. The histograms are com-
 641 parable, although the SSDF-AC product in 2013 has mean uncertainties that are 4%

642 lower on average than the SSDF-A product in 2011. This is to be expected as the
 643 additional information from CrIMSS provides greater certainty for SSDF.

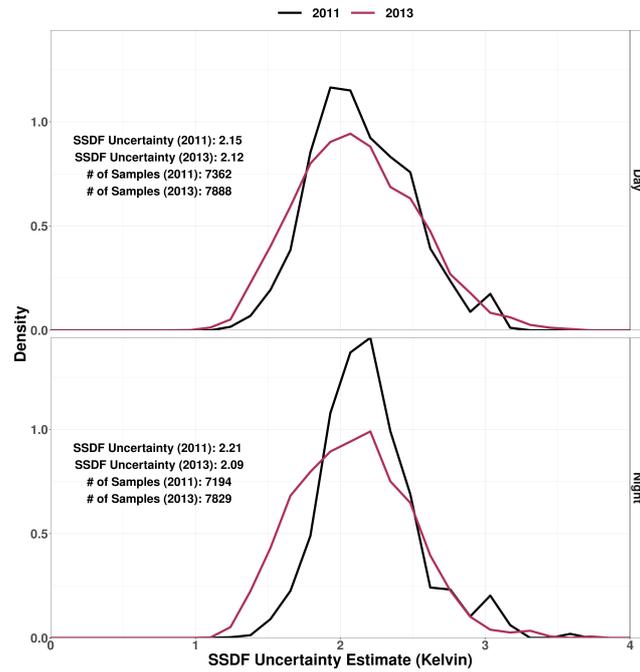


Figure 19: SSDF uncertainty histogram for 2011 (black) and 2013 (red) aggregated by day (top) and night (bottom). Summary statistics of mean SSDF uncertainty are provided as text on the upper left.

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 Aqua-AIRS V7 and SNPP-CrIMSS CLIMCAPS V2 L2 NSAT datasets. 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 fuse them into an improved dataset.

The fused NSAT product could be used for applications over CONUS that require NSAT data and that would benefit from the improvements we have demonstrated here from a detailed validation using withheld ISD data as a reference dataset. The SSDF method generates a fused gridded product that has no missing data; has improved accuracy and precision relative to the input satellite datasets; and includes uncertainty estimates that are more consistent with the observed errors relative to the ISD reference. The NSAT SSDF pilot product is comparable in precision and accuracy to the state-of-the-art ERA5-Land reanalysis, but unlike reanalysis it does not involve dynamical weather modeling, only spatial covariance modeling. Furthermore, unlike reanalysis it could in the future support a near-real-time version for operational applications.

SSDF is a general method and can be applied to one or more L2 datasets, so long as each dataset estimates the same observable. For example, fusion of Aqua-AIRS and SNPP-CrIMSS estimates of NSAT works because both satellites estimate NSAT at approximately 1:30 and 13:30 local solar time. However, it would not make sense to directly fuse NSAT estimates from Infrared Atmospheric Sounding Interferometer (IASI) instruments on the MetOp satellites with the Aqua and SNPP datasets, as the MetOp satellites pass over at approximately 9:30 and 21:30 local solar time, when NSAT is at different points of the diurnal cycle. On the other hand, the details of instruments used to make the input datasets, and their spatial footprints and sampling, are immaterial. For example, it would be possible to fuse NSAT derived from the Visible Infrared Imaging Radiometer Suite (VIIRS) land surface temperature (LST) product via (for example) regression modeling (Good, 2015), since such a LST-derived NSAT product would also sample at approximately 1:30 and 13:30 local solar time. SSDF could be applied across a wide range of observables estimated as L2 satellite datasets, such as atmospheric composition, water vapor profiles, or vapor pressure deficit (the difference between the water vapour pressure and the saturation water vapour pressure). Bias and variance estimates of the input datasets are required, and we emphasize that the quality of the SSDF product depends on the quality of those error estimates.

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 the two issues our validation has uncovered, namely (1) increased bias at a small number of data points at elevations in excess of 2500 m, and (2) a ~ 0.1 K shift in annual mean bias between the SSDF-AC and SSDF-A (AIRS-only) versions.

We also plan to create an NSAT SSDF product over global land areas, create a high spatial resolution NSAT SSDF product by including high spatial resolution input NSAT datasets in the fusion, and apply the SSDF method to other hyperspectral surface products, starting with near-surface specific humidity.

Open Research

The fused SSDF NSAT datasets described in this paper are available from the NASA GES DISC repository at <https://doi.org/10.5067/CPXNAPA2WSQ8> (SSDF-AC) and <https://doi.org/10.5067/8AE9Y5TSXFX4> (SSDF-A).

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.

SNPP-CrIMSS-CLIMCAPS V2 is available from the NASA GES DISC repository (Barnet, 2019). Near surface temperature (surf_air_temp), the corresponding QC flag (surf_air_temp_qc), and the corresponding error estimate (surf_air_temp_err) were obtained from the NSR product.

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

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

Acknowledgments

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

Appendix A Matchups and bias estimation

In this appendix, 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 τ_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,

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

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

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

747 where $M(d, j, \mathbf{s})$ is the number of time points acquired by the ISD station at \mathbf{s} on
 748 day d in mode j , L is the total number of hexagonal grid cells, and we write \mathbf{u}_{ms}^* to
 749 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 $Z_{dji}^A(\mathbf{u})$, 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 $Z_{dji}^{k*}(\mathbf{u})$ as follow:

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

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

751 Appendix B SSDF methodology

752 Consider a discretized domain where $\{Y(\mathbf{s}) : \mathbf{s} \in D\}$ is a hidden, real-valued
 753 spatial observable. The domain of interest is $\cup\{A_i \subset \mathbb{R}^d : i = 1, \dots, N_D\}$, which is
 754 made up of N_D fine-scale, non-overlapping, areal regions $\{A_i\}$ with locations $D \equiv$
 755 $\{\mathbf{p}_i \in A_i : i = 1, \dots, N_D\}$. Nguyen et al. (2012) call these fine-scale regions Basic
 756 Areal Units (BAUs), and they represent the smallest resolution at which we will make
 757 estimates with the model.

758 For a given day and mode (d and j using the notation of the previous subsection),
 759 denote the vector of NSAT data at all locations by \mathbf{Z}^k , where $k = 1$ for AIRS and
 760 $k = 2$ 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 observable is averaged over the areal region plus an independent error term. That is,

$$Z^k(B) = \frac{1}{|D \cap B|} \left\{ \sum_{\mathbf{s} \in D \cap B} Y(\mathbf{s}) \right\} + \epsilon^k(B); \quad B \subset \mathfrak{R}^d. \quad (\text{B1})$$

761 where $Y(\cdot)$ is a geophysical observable (here, NSAT) that is common to both datasets,
 762 and $\epsilon^k(\cdot)$ is an independent but non-identically distributed Gaussian random variable.
 763 That is, we assume that the q -th error in the k -th dataset is distributed as $\epsilon_q^k \sim$
 764 $N(b_q^k, v_q^k)$. In general, b_q^k is not zero, however, in our case b_q^k is assumed to be zero
 765 because we performed bias correction as described in the previous subsection, and
 766 v_q^k are calculated from the hexagonal-cell-specific mean and variance estimates (see
 767 Appendix A for details).

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

$$\hat{Y}(B_0) = \mathbf{a}'_1 \mathbf{Z}_1 + \mathbf{a}'_2 \mathbf{Z}_2, \quad (\text{B2})$$

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

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

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

$$\mathbf{1} = \mathbf{a}'_1 \mathbf{1}_{N_1} + \mathbf{a}'_2 \mathbf{1}_{N_2}, \quad (\text{B3})$$

775 where $\mathbf{1}_{N_k}$ is an N_k -dimensional vector of ones. The solution to the minimization
 776 problem in (B3) can be found via the method of Lagrange multipliers; but it requires
 777 knowledge of the spatial covariance structure $C(B_i, B_j)$, which can be expanded in
 778 terms of the BAU covariances:

$$C(B_i, B_j) = \frac{1}{|D \cap B_i| |D \cap B_j|} \sum_{\mathbf{u} \in D \cap B_i} \sum_{\mathbf{v} \in D \cap B_j} C(\mathbf{u}, \mathbf{v}). \quad (\text{B4})$$

Typically, the covariance structure in kriging-based approaches is estimated from the data, but the formulation in Equation B4 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 observable, here NSAT, can be written as the linear mixed model,

$$Y(\mathbf{s}) = \mathbf{t}(\mathbf{s})' \boldsymbol{\alpha} + \mathbf{S}(\mathbf{s})' \boldsymbol{\eta} + \xi(\mathbf{s}). \quad (\text{B5})$$

779 where $\mathbf{t}(\cdot) \equiv (t_1(\cdot), \dots, t_p(\cdot))'$ is a vector of p known covariates, such as geographical
 780 coordinates or other physical variables. The vector of linear coefficients, $\boldsymbol{\alpha}$, is unknown
 781 and will be estimated from the data. The middle term captures the spatial dependence
 782 as the product of an r -dimensional vector of known spatial basis functions, $\mathbf{S}(\mathbf{s})$, and an
 783 r -dimensional Gaussian random variable, $\boldsymbol{\eta}$. Here, we assume that with $\boldsymbol{\eta} \sim N(\mathbf{0}, \mathbf{K})$.
 784 Similar to the implementation in Nguyen et al. (2012), we implement these using
 785 multi-resolution bisquare basis functions centered at different resolutions of the Inverse

786 Snyder Equal-Area Projection Aperture 3 Hexagon (ISEA3H) type within the Discrete
 787 Global Grid (DGGRID) software (specifically, resolutions 2, 3, and 5 of ISEA3H, for
 788 details see Sahr, 2019). The last term, $\xi(\cdot)$, describes the BAU-scale variability of the
 789 process. We assume that $\xi(\cdot)$ is an independent Gaussian process with mean zero and
 790 variance σ_ξ^2 .

791 The SME model in Equation B5 has useful change-of-support properties, which
 792 makes computation of the spatial covariance function straightforward. In particular,
 793 Nguyen et al. (2012) shows that

$$\text{cov}(Z(B_i), Z(B_j)) = \mathbf{S}(B_i)' \mathbf{K} \mathbf{S}(B_j) + \sigma_\xi^2 \frac{|D \cap B_i \cap B_j|}{|D \cap B_i| |D \cap B_j|} + v_i^k I(i = j), \quad (\text{B6})$$

where

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

794 Notice that Equation B6 allows us to express the covariance between spatial averages
 795 *explicitly* in terms of the spatial dependence parameter \mathbf{K} . This allows for straightfor-
 796 ward estimation of it from footprint data.

797 Another advantage of the SME model is its scalability. For a general covariance
 798 structure, solving for \mathbf{a}_1 and \mathbf{a}_2 requires inverting a $(N_1 + N_2) \times (N_1 + N_2)$ covariance
 799 matrix, which has computational complexity $O((N_1 + N_2)^3)$. For large datasets such
 800 as AIRS and CrIMSS where the data size is on the order of tens of thousands, this
 801 matrix inversion is computationally infeasible. However, the model in Equation B5
 802 implies the following full covariance matrix:

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

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

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

808 Note that the inversion above, and hence the calculation of the coefficients \mathbf{a}_1 and
 809 \mathbf{a}_2 for the fused estimate, is very fast because it only requires inversion of the *sparse*
 810 $(N_1 + N_2) \times (N_1 + N_2)$ matrix \mathbf{U} , which is typically very sparse, and inversion of \mathbf{K}
 811 and $(\mathbf{K}^{-1} + \mathbf{S}' \mathbf{U}^{-1} \mathbf{S})$, both of which are $r \times r$ matrices ($r \ll N_1 + N_2$).

812 References

- 813 AIRS Project. (2019). *Aqua/AIRS L2 Standard Physical Retrieval (AIRS-*
 814 *only) V7.0*. Goddard Earth Sciences Data and Information Services Cen-
 815 *ter (GES DISC)*. Greenbelt, MD, USA. (Accessed: 2019-2021) doi:
 816 10.5067/VP1M6OG1X7M1
- 817 AIRS Project. (2020). *About the data*. Retrieved 2021-12-26, from [https://airs](https://airs.jpl.nasa.gov/data/about-the-data/granules/)
 818 [.jpl.nasa.gov/data/about-the-data/granules/](https://airs.jpl.nasa.gov/data/about-the-data/granules/)
- 819 AIRS team. (n.d.). *About the products*. Retrieved from [https://airs.jpl.nasa](https://airs.jpl.nasa.gov/data/products/retrieval-systems/)
 820 [.gov/data/products/retrieval-systems/](https://airs.jpl.nasa.gov/data/products/retrieval-systems/) (accessed: June 3, 2022)
- 821 Barnet, C. (2019). *Sounder SIPS: Suomi NPP CrIMSS Level 2 CLIMCAPS normal*
 822 *spectral resolution: Cloud cleared radiances v2*. Goddard Earth Sciences Data
 823 *and Information Services Center (GES DISC)*. Greenbelt, MD, USA. (Ac-
 824 *cessed: 2019-2021) doi: 10.5067/CNG0ST72533Z*

- 825 Blackwell, W. (2005). A neural-network technique for the retrieval of atmospheric
 826 temperature and moisture profiles from high spectral resolution sounding data.
 827 *Geoscience and Remote Sensing, IEEE Transactions on*, 43(11), 2535-2546.
 828 doi: 10.1109/TGRS.2005.855071
- 829 Braverman, A., Hobbs, J., Teixeira, J., & Gunson, M. (2021). Post hoc uncertainty
 830 quantification for remote sensing observing systems. *SIAM/ASA Journal on*
 831 *Uncertainty Quantification*, 9(3), 1064–1093.
- 832 Chahine, M. T., Pagano, T. S., Aumann, H. H., Atlas, R., Barnet, C., Blaisdell, J.,
 833 ... Zhou, L. (2006). AIRS: Improving weather forecasting and providing new
 834 data on greenhouse gases. *Bulletin of the American Meteorological Society*,
 835 87(7), 911–926.
- 836 Chen, Y., Liang, S., Ma, H., Li, B., He, T., & Wang, Q. (2021). An all-sky 1 km
 837 daily land surface air temperature product over mainland china for 2003–2019
 838 from modis and ancillary data. *Earth System Science Data*, 13(8), 4241–4261.
- 839 Cressie, N. (1993). *Statistics for spatial data*. John Wiley & Sons.
- 840 Cressie, N., & Johannesson, G. (2008). Fixed rank kriging for very large spatial data
 841 sets. *Journal of the Royal Statistical Society: Series B (Statistical Methodol-*
 842 *ogy)*, 70(1), 209–226.
- 843 Ferguson, C. R., & Wood, E. F. (2010). An evaluation of satellite remote sensing
 844 data products for land surface hydrology: Atmospheric infrared sounder. *Jour-*
 845 *nal of Hydrometeorology*, 11(6), 1234–1262.
- 846 Ghamisi, P., Rasti, B., Yokoya, N., Wang, Q., Hofle, B., Bruzzone, L., ... others
 847 (2019). Multisource and multitemporal data fusion in remote sensing: A
 848 comprehensive review of the state of the art. *IEEE Geoscience and Remote*
 849 *Sensing Magazine*, 7(1), 6–39.
- 850 Good, E. (2015). Daily minimum and maximum surface air temperatures from
 851 geostationary satellite data. *Journal of Geophysical Research: Atmospheres*,
 852 120(6), 2306–2324.
- 853 Gotway, C. A., & Young, L. J. (2002). Combining incompatible spatial data.
 854 *Journal of the American Statistical Association*, 97(458), 632-648. doi:
 855 10.1198/016214502760047140
- 856 Hammerling, D. M., Michalak, A. M., & Kawa, S. R. (2012). Mapping of CO₂ at
 857 high spatiotemporal resolution using satellite observations: Global distribu-
 858 tions from OCO-2. *Journal of Geophysical Research: Atmospheres*, 117(D6).
- 859 Harrison, G., & Burt, S. D. (2021). Quantifying uncertainties in climate data: mea-
 860 surement limitations of naturally ventilated thermometer screens. *Environmen-*
 861 *tal Research Communications*.
- 862 Henderson, H., & Searle, S. (1981). On deriving the inverse of a sum of matrices.
 863 *SIAM Review*, 23, 53-60.
- 864 Hennermann, K., & Berrisford, P. (2019). *ERA5 data documentation*, ECMWF.
- 865 Hobbs, J., Braverman, A., Cressie, N., Granat, R., & Gunson, M. (2017).
 866 Simulation-based uncertainty quantification for estimating atmospheric CO₂
 867 from satellite data. *SIAM/ASA Journal on Uncertainty Quantification*, 5(1),
 868 956–985.
- 869 Kendall, M. G. (1948). Rank correlation methods.
- 870 Mann, H. B. (1945). Nonparametric tests against trend. *Econometrica: Journal of*
 871 *the econometric society*, 245–259.
- 872 McNally, A. P., Watts, P. D., A. Smith, J., Engelen, R., Kelly, G. A., Thépaut,
 873 J. N., & Matricardi, M. (2006). The assimilation of AIRS radiance data at
 874 ECMWF. *Quarterly Journal of the Royal Meteorological Society*, 132(616),
 875 935–957. Retrieved from <http://dx.doi.org/10.1256/qj.04.171> doi:
 876 10.1256/qj.04.171
- 877 Nguyen, H., Cressie, N., & Braverman, A. (2012). Spatial statistical data fusion
 878 for remote sensing applications. *Journal of the American Statistical Asso-*
 879 *ciation*, 107(499), 1004-1018. Retrieved from <https://doi.org/10.1080/>

- 880 01621459.2012.694717 doi: 10.1080/01621459.2012.694717
- 881 Nguyen, H., Katzfuss, M., Cressie, N., & Braverman, A. (2014). Spatio-temporal
882 data fusion for very large remote sensing datasets. *Technometrics*, *56*(2), 174–
883 185.
- 884 Olsen, E. T., Fishbein, E., Manning, E., & Maddy, E. (2017). AIRS/AMSU/HSB
885 Version 6 L2 product levels, layers and trapezoids. *Jet Propulsion Laboratory,*
886 *Pasadena, CA, USA*.
- 887 Rayner, N. A., Auchmann, R., Bessembinder, J., Brönnimann, S., Brugnara, Y.,
888 Capponi, F., ... others (2020). The eustace project: delivering global, daily
889 information on surface air temperature. *Bulletin of the American Meteorologi-*
890 *cal Society*, *101*(11), E1924–E1947.
- 891 Rohde, R. A., & Hausfather, Z. (2020). The berkeley earth land/ocean temperature
892 record. *Earth System Science Data*, *12*(4), 3469–3479.
- 893 Sahr, K. (2019). *Dggrid version 7.0: User documentation for discrete global grid*
894 *software*.
- 895 Sahr, K., White, D., & Kimerling, A. J. (2003). Geodesic discrete global grid sys-
896 tems. *Cartography and Geographic Information Science*, *30*(2), 121–134.
- 897 Sen, P. K. (1968). Estimates of the regression coefficient based on kendall’s tau.
898 *Journal of the American statistical association*, *63*(324), 1379–1389.
- 899 Smith, A., Lott, N., & Vose, R. (2011). The integrated surface database: Recent de-
900 velopments and partnerships. *Bulletin of the American Meteorological Society*,
901 *92*(6), 704–708.
- 902 Smith, N., & Barnett, C. D. (2019). Uncertainty characterization and propagation
903 in the Community Long-Term Infrared Microwave Combined Atmospheric
904 Product System (CLIMCAPS). *Remote Sensing*, *11*(10), 1227.
- 905 Smith, N., & Barnett, C. D. (2020). CLIMCAPS observing capability for temper-
906 ature, moisture, and trace gases from AIRS/AMSU and CrIS/ATMS. *Atmo-*
907 *spheric Measurement Techniques*, *13*(8), 4437–4459.
- 908 Sun, J., McColl, K. A., Wang, Y., Rigden, A. J., Lu, H., Yang, K., ... Santanello Jr,
909 J. A. (2021). Global evaluation of terrestrial near-surface air temperature
910 and specific humidity retrievals from the atmospheric infrared sounder (airs).
911 *Remote Sensing of Environment*, *252*, 112146.
- 912 Susskind, J., Blaisdell, J. M., & Iredell, L. (2014). Improved methodology for
913 surface and atmospheric soundings, error estimates, and quality control pro-
914 cedures: the atmospheric infrared sounder science team version-6 retrieval
915 algorithm. *Journal of Applied Remote Sensing*, *8*(1), 084994. Retrieved from
916 <http://dx.doi.org/10.1117/1.JRS.8.084994> doi: 10.1117/1.JRS.8.084994
- 917 Thrastarson, H. T., Manning, E., Kahn, B., Fetzer, E., Yue, Q., Wong, S., ... others
918 (2020). AIRS/AMSU/HSB Version 7 Level 2 product user guide. *Jet Propul-*
919 *sion Laboratory, California Institute of Technology: Pasadena, CA, USA*,
920 83–92.
- 921 Wilks, D. S. (2006). *Statistical methods in the atmospheric sciences* (Second ed.).
922 Burlington, MA: Academic Press.
- 923 Yue, Q., Lambrigtsen, B., et al. (2017). AIRS V6 test report supplement: Perform-
924 ance of AIRS+AMSU vs. AIRS-only retrievals. *Jet Propulsion Laboratory,*
925 *California Institute of Technology: Pasadena, CA, USA*. Retrieved from
926 [https://docserver.gesdisc.eosdis.nasa.gov/repository/Mission/](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)
927 [AIRS/3.3.ScienceDataProductDocumentation/3.3.5.ProductQuality/](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)
928 [V6_Test_Report_Supplement_Performance_of_AIRS+AMSU_vs_AIRS-Only](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)
929 [_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)
- 930 Yue, Q., Lambrigtsen, B., et al. (2020). AIRS V7 L2 performance test and vali-
931 dation report. *Jet Propulsion Laboratory, California Institute of Technology:*
932 *Pasadena, CA, USA*. Retrieved from [https://docserver.gesdisc.eosdis](https://docserver.gesdisc.eosdis.nasa.gov/public/project/AIRS/V7_L2_Performance_Test_and_Validation_report.pdf)
933 [.nasa.gov/public/project/AIRS/V7_L2_Performance_Test_and_Validation](https://docserver.gesdisc.eosdis.nasa.gov/public/project/AIRS/V7_L2_Performance_Test_and_Validation_report.pdf)
934 [_report.pdf](https://docserver.gesdisc.eosdis.nasa.gov/public/project/AIRS/V7_L2_Performance_Test_and_Validation_report.pdf)

935 Yue, Q., Lambrigtsen, B., et al. (2021). Version 2 CLIMCAPS-Aqua re-
936 trieval product performance test report. *Jet Propulsion Laboratory, Cal-*
937 *ifornia Institute of Technology: Pasadena, CA, USA.* Retrieved from
938 [https://docserver.gesdisc.eosdis.nasa.gov/public/project/Sounder/](https://docserver.gesdisc.eosdis.nasa.gov/public/project/Sounder/CLIMCAPS.V2.Test.Report.Aqua.pdf)
939 [CLIMCAPS.V2.Test.Report.Aqua.pdf](https://docserver.gesdisc.eosdis.nasa.gov/public/project/Sounder/CLIMCAPS.V2.Test.Report.Aqua.pdf)