

Structures of Error Covariance in Global Navigation Satellite System Reflectometry

by

Charles Evans Powell

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Climate and Space Sciences and Engineering)
in the University of Michigan
2024

Doctoral Committee:

Professor Christopher S. Ruf, Chair
Professor Christiane Jablonowski
Associate Research Scientist Darren S. McKague
Professor Emeritus Fawwaz T. Ulaby

Charles Evans Powell

cepowell@umich.edu

ORCID iD: 0000-0003-4757-782X

© Charles Evans Powell 2024

Dedicated to the engines of ingenuity:

Patient teachers and curious students.

Acknowledgements

This work was supported in part by the National Aeronautics and Space Administration Science Mission Directorate under Contract NNL13AQ00C with the University of Michigan.

This work has not been reviewed or endorsed by any U.S. Government Department or Agency, including the U.S. Department of Commerce, National Oceanic and Atmospheric Administration, and National Aeronautics and Space Administration. Any opinions therein are the author's own and not representative of the U.S. Government.

Several chapters of this work are derived substantially from published or submitted manuscripts. Copyright may be owned by authors or publishers, and this work makes no new claim to copyright of previously published work. Copyrighted works have been reproduced with permission or under appropriate licensing terms. Copyright of unpublished work may be subject to transfer without notice.

Several chapters of this work were initially written with coauthors. I am grateful to their generosity and collaboration over the last five years, and their contributions significantly improved the quality and scope of the work. Any remaining errors are my own.

Table of Contents

Dedication.....	ii
Acknowledgements.....	iii
List of Tables	viii
List of Figures.....	ix
Abstract.....	xiv
Chapter 1 Introduction	1
1.1 Motivations	1
1.2 The CYGNSS Observatory.....	2
1.2.1 Sampling Characteristics and Constellation Architecture	3
1.2.2 The Delay-Doppler Map.....	5
1.2.3 The CYGNSS Observatory.....	8
1.3 GNSS-Reflectometry for Ocean Wind Sensing.....	9
1.4 Data Assimilation Basics	11
1.4.1 What Happens When R is Poorly Specified	12
1.5 Chapter Summary	14
1.6 Abstract.....	15
Chapter 2 Improving CYGNSS Duty Cycle.....	16
2.1 Introduction (Abridged).....	16
2.2 Theory (Abridged)	17
2.3 Methods.....	20

2.4 Estimated Beta Angle and Thermal Variation for Selected Days.....	21
2.4.1 Characterizing the Impact of Change in Cadence.....	23
2.4.2 Optimal Blackbody Sampling Cadence	28
2.5 Performance by Orbit Sector	29
2.6 Discussion.....	30
2.6.1 Characterizing the Impact to Duty Cycle.....	30
2.6.2 Characterizing the Impact to the Level 1A Error Budget	31
2.6.3 Comparing Expected Results to Empirical Data	33
2.7 Summary of CYGNSS Performance Before and After Blackbody Sequence Modification.....	34
2.8 Summary and Conclusions	34
Chapter 3 Exploring Representation Error	36
3.1 Abstract.....	36
3.2 Introduction.....	37
3.3 Methodology.....	41
3.3.1 Formulation.....	41
3.3.2 Choice of Data and Study Selection.	43
3.3.3 Temporal Decorrelation	44
3.3.4 Spatial Decorrelation	47
3.4 Discussion.....	50
3.4.1 Shortcomings	50
3.4.2 Utility and Future Work.....	51
3.5 Data Availability Statement.....	53
Chapter 4 An Instrument Error Correlation Model for GNSS-R.....	54
4.1 Abstract.....	54
4.2 Introduction (Abridged).....	55

4.3 Materials and Methods (Abridged).....	56
4.3.1 The Bottom-up Correlated Error Model	57
4.3.2 Model Assumptions	58
4.4 Covariance Matrix Construction.....	59
4.4.1 Calibrated Nadir Power <i>P_g</i> Error Model	63
4.4.2 Zenith Power <i>P_Z</i> Error Model.....	67
4.4.3 Error Model for Antenna Gain Patterns <i>GR</i> and <i>GZ</i>	68
4.4.4 Zenith-Specular Ratio ζ Error Model	73
4.5 Verification Techniques.....	77
4.5.1 Curating Matchup Observations	77
4.5.2 Generating Model NBRCS	79
4.5.3 Estimating Total Correlated Error	80
4.5.4 Model Tuning.....	82
4.6 Results.....	84
4.6.1 Bulk Behavior	84
4.6.2 Single-Track Comparisons.....	87
4.6.3 Dynamic Correlated Error Estimation and Impact of Tuning	89
4.7 Discussion.....	92
4.7.1 Limitations	92
4.7.2 Impact of Tuning Parameters.....	93
4.8 Conclusions.....	95
Chapter 5 Error Covariance in Observation Space and Geophysical Partitions.....	97
5.1 Abstract.....	97
5.2 Introduction (abridged).....	98
5.3 Error Model Construction (abridged)	98

5.3.1 Definitions.....	98
5.3.2 Observation Error Covariance Model.....	101
5.3.3 Correlated Instrument Errors in Observation Space	103
5.3.4 Evaluation of Representation-Retrieval Errors	105
5.3.5 Data Selection	106
5.4 Partition of Representation-Retrieval Error	107
5.4.1 Selection of Correctable Variables	108
5.4.2 Generating Windspeed Correction.....	110
5.5 Results and Discussion	114
5.5.1 Validation.....	116
5.5.2 Limitations	118
Data Availability Statement.....	119
Chapter 6 Conclusions	120
6.1 Contributions	120
6.2 Future Work	122
6.3 Lessons learned for future GNSS-R missions	123
References.....	125

List of Tables

Table 2-1. Exemplar temperature variability for different orbit beta angles.	21
Table 4-1. The magnitudes of 1-sigma errors for each term in Equation (1.2). The shaded rows indicate that the absolute magnitudes of these terms are negligible compared to the dominant error terms and are neglected in the construction of the error model in this work.	57
Table 4-2. The magnitudes of 1-sigma errors for each term in Equation B3. The shaded rows indicate that the absolute magnitude is negligible compared to the dominant error terms and is neglected in the construction of the error model in this work.	64
Table 4-3 –The chosen magnitudes of model tuning parameters for <i>Rmod</i>	87
Table 5-1 – Abbreviated list of ERA5 model variables matched up with the CYGNSS windspeed observations. These variables are sorted by the absolute magnitude of correlation with windspeed error <i>eobs</i> . Also displayed is the correlation of the model matchup with CYGNSS observations and ERA5 model windspeed. Variables in <i>italics</i> were passed over in analysis because they are too strongly correlated with ERA5 windspeed. Variables in bold are used in this analysis because they are strongly correlated with windspeed error but are sufficiently decorrelated from ERA5 windspeed.	109
Table 5-2 – List of ERA5 model variables matched up with the MWP-corrected windspeed. These variables are sorted by the absolute magnitude of correlation with MWP-corrected windspeed error <i>eobs</i> , λ . Also displayed is the correlation of the model matchup with MWP-corrected CYGNSS windspeed and MWP-corrected ERA5 windspeed. Variables in <i>italics</i> were passed over in analysis in the first step for being too strongly correlated with uncorrected ERA5 windspeed. Variables in bold are used in this analysis. Note that the MWP correction increases the WSS correlation with ERA5 windspeed, because MWP and WSS are related and the correction itself introduces additional windspeed dependence. For the purposes of this analysis, we assert that this does not violate the original rationale for excluding other parameters, because we know that WSS is not a robust proxy for ERA5 winds.	112
Table 5-3 – List of regression parameters for the corrections applied in this section. Note that the regression coefficients <i>ci</i> are listed in descending order, where <i>ci</i> take the form $cnxn + cn - 1xn - 1 + \dots$	114

List of Figures

Figure 1.1. The sampling characteristics of CYGNSS from Ruf et al. (2022). (a) 90 min sampling widow for eight CYGNSS satellites. (b) 24 hours of CYGNSS observations. (c) the probability density (blue) and cumulative density (orange) of revisit time for a given observation..... 5

Figure 1.2 – A diagram representing the geometry of the CYGNSS GNSS-R observation. A GPS satellite travelling at velocity vt emits a signal that is received both directly and via specular reflection by a CYGNSS receiver travelling at velocity vr . The CYGNSS receiver also receives a direct signal from the GPS transmitter to enable geolocation and DDM processing. The specular point is mapped by isodelays (purple) and iso-Dopplers (blue). Different points P and Q are mapped to the same delay and Doppler bin. Figure from Ulaby and Long (2014)..... 7

Figure 1.3 – A diagram illustrating the mapping between physical space and delay-Doppler space. The left panel shows a physical measurement space with isodelays (green) and a Doppler hyperbola. The right panel illustrates the mapping to delay and Doppler. Points A, B, C, D map to bins A', B', C', and D', respectively. Figure from Zavorotny et al. (2014). 7

Figure 1.4 - Exploded view (left) and underside view (right) of CYGNSS. Courtesy Keith Smith, SouthWest Research Institute, from Ruf et al. (2022). 8

Figure 2.1 – Example of temperature influence on receiver gain. Above, FM4’s port receiver gain is plotted (blue) with the temperature at a nearby probe (red) for a full day starting at 0Z 12 APR 2019. Note that gain is inversely related to local temperature, and that these swings are periodic, consistent with CYGNSS’s 95-minute orbital period..... 19

Figure 2.2 – Orbit beta angle calculated for FM04 for 2019. The black, magenta, and orange vertical lines denote the low, medium, and high beta angle days respectively. 20

Figure 2.3 – Histograms of gain error at $n = \{5, 15, 25, 35, 45\}$ minutes. The blue and red distributions represent data from the port and starboard antennas, respectively. At low n , the gain error approaches a delta function. As sampling cadence time increases, the distribution widens, illustrating the degradation from the original one-minute sampling cadence. 24

Figure 2.4 – Standard deviations of growth in error as a function of sampling cadence for NBRCS (top), FDS windspeed (middle), and YSLF windspeed (bottom) at low beta angle. Red traces indicate values for the starboard antenna, and blue traces for the port side. 25

Figure 2.5 – PCT1 of error as a function of sampling cadence period for NBRCS (top), FDS windspeed (middle), and YSLF windspeed (bottom) at low beta angle. Red traces indicate values for the starboard antenna, and blue traces for the port side. 25

Figure 2.6 – PCT1 of NBRCS error as a function of sampling cadence period at various beta angles. At lower beta, the performance degradation is more pronounced. 27

Figure 2.7 – Standard deviation of YSLF windspeed error as a function of sampling cadence period at various orbit sectors. The errors due to increased blackbody sampling rates are greatest immediately after crossing the terminator. This effect is largest at longer sampling periods. 29

Figure 3.1 – (a) The temporal decorrelation behavior for a year of observations is demonstrated for various meteorological parameters from NDBC buoy station 51004. (b) The representativity error as calculated from (1) is shown. Note that the magnitude is dependent on the unit of measure. 46

Figure 3.2 – The zonal and meridional transects for the Pacific basin centered at Station 51004 are highlighted. The zonal transect consists of 210 grid cells at 0.625 deg spacing. The meridional transect consists of 254 grid cells at 0.5 deg spacing. The background field is of an initialization of MERRA-2’s global windspeed output to emphasize that the spatial patterns of these parameters are quite different meridionally than they are zonally. 48

Figure 3.3 – The spatial decorrelation behavior and error statistics are partitioned into meridional and zonal components. (a) The meridional decorrelation behavior is shown for the specified environmental parameters. (b) The meridional representativity error is shown, note that the y-axis units are parameter-dependent. (c) The zonal decorrelation behavior is shown for the specified environmental parameters. (d) The zonal representativity error is shown, note that the y-axis units are parameter-dependent. 49

Figure 4.1 - The filtering kernel used to smooth nadir and zenith antenna gain patterns. This kernel imposes correlated error structure onto the antenna gain patterns. The coordinate system should be read as distance in the relevant antenna reference frame. Therefore, if two observations are nearby in the antenna pattern, they will have strongly correlated errors. However, if two observations are far apart in the pattern, the correlated structure decays. 72

Figure 4.2 – This figure illustrates a calculated zenith–specular ratio ζ as a function of observation incidence angle θ_{inc} at a fixed GPS antenna azimuth for GPS PRN 2. The blue trace is interpolated from raw observations over a two-year period at each of the elevation gridpoints in the GPS antenna pattern for a single azimuthal cut of PRN 2. The red trace is a generated smoothed zenith–specular ratio ζ that would be similar to the ones used in the operational LUTs using a 4th-order power series fit. Note that at large incidence angles, i.e., grazing observations, there is a great deal of uncertainty in ζ because there are few valid observations in those regions. In practice, only data at incidence angles < 60 degrees constrain error in ζ 75

Figure 4.3 – This figure depicts an error correlation matrix derived from Equation (E4) for GPS PRN 2, which is used to produce $Ecorr\zeta$. The matrix is a function of the observation incidence angle for sample points i and j . Note that the mapping from coordinates in GPS elevation angle θ_z and θ_s to Earth scattering incidence angle θ_{inc} is coarse and produces the checkerboard-like pattern near the diagonal. A single error in the measurement of GT in the GPS antenna pattern is will highly correlate within a range of incidence angles in θ_{inc} , as mapped. At high incidence angles, errors are strongly correlated as the power series fit is likely to be wrong in the same direction. 76

Figure 4.4 – Comparison of FM1 and FM5 observations for 11 SEP 2019. (a) Rendering of the ground samples captured during the 24 h period in the near-overlap condition; (b) A single track of samples for both FM1 and FM5 that has been matched sample-for-sample to facilitate one-to-one comparisons of the observations. The samples from each observatory are no more than 0.5 deg apart in distance and were acquired approximately 3 s apart in time. 78

Figure 4.5 – A display of model (blue) and observed (green) single-differenced matchup data for a single CYGNSS track, with the double-differenced data (red) overlaid. Note that both the model and observed single differences are similarly high early in the track, suggesting that the differences in the observations may be because the samples are observing fundamentally different targets, captured in $SDmod$. The double-difference accounts for these types of errors, and for the entire track, the double-difference is quite stable. The sparsity of data is caused by quality control parameters, such as when the CYGNSS observatory is performing its onboard calibration procedure, which occurs once every minute. Because the two matched observatories do not have synchronized calibration clocks, the matched dataset typically flags out two of these cadences for every minute of sampling. 81

Figure 4.6 – Comparison of the bulk error autocorrelation behavior for all matched tracks. The single difference of modeled σ_{modo} (blue), the single difference of observed σ_o (green), the double-difference DD (purple), the untuned bulk model error correlation \mathcal{R}_{mod} (dotted red), and the final tuned bulk model error correlation \mathcal{R}_{mod} (orange) are shown. The shading indicates the estimated 1-sigma standard deviation of the population decorrelation behavior for each individual trace. The lags are computed in seconds, which correspond to single samples at 1 Hz sampling for the CYGNSS observatory. Note that DD generally follows the $SDobs$ trace, suggesting that single-differencing for this particular use case may be a reasonable representation of bulk error correlation behavior. 85

Figure 4.7 – Comparison of the single-track error autocorrelation behavior for three selected tracks. The autocorrelation of the NBRCS double-difference (solid lines), untuned modeled error correlation ρ_{mod} , and tuned modeled error autocorrelation ρ_{mod} (dashed lines) for each of the three tracks are shown. Each track is painted a different color. This figure demonstrates the wide variability of single-track autocorrelations of double-differences, arising from the variability of the observable, the limited amount of data in a single track, and the challenges in quality controlling sufficient observation data. It is possible that the correlated error does vary this much from track to track. In contrast, both the tuned and untuned modeled error autocorrelation ρ_{mod} are much more stable from track to track. 88

Figure 4.8 – Generated untuned and tuned *Rmod* for two nearby tracks captured by the same CYGNSS receiver, each with three different representations. (a) Untuned *Rmod* as represented by sample index. The two tracks are concatenated in a vector, and the modeled error correlation is calculated for each combination of indices. The green vertical line represents an exemplar index ($i = 129$) along the track, where the correlated error is estimated for every other observation in the neighborhood. (b) Untuned *Rmod* for the same exemplar index as represented by the physical location of sample acquisition. (c) Untuned *Rmod* for the same exemplar index represented by time of acquisition. The traces for the two different tracks are plotted in different colors. The error correlation along the same track is in blue, while the error correlation for the adjacent track is in red. (d) Tuned *Rmod* as represented by sample index for the same two tracks as in (a). (e) Tuned *Rmod* for the same tracks represented in the physical location. (f) Tuned *Rmod* for the same exemplar index represented by the time of acquisition. Note that for both the untuned and tuned cases, the correlated error for the adjacent track is non-zero but generally very small compared to the error correlation along the track. 89

Figure 4.9 – Generated untuned and tuned *Rmod* for two nearby tracks captured by two different CYGNSS receivers during the overlap period, each with three different representations. (a) Untuned *Rmod* as represented by sample index. The two tracks are concatenated in a vector, and the modeled error correlation is calculated for each combination of indices. The green vertical line represents an exemplar index ($i = 129$) along the track where the correlated error is estimated for every other observation in the neighborhood. The correlated error model allows for correlated errors across receivers. (b) Untuned *Rmod* for the same exemplar index as represented by the physical location of sample acquisition. (c) Untuned *Rmod* for the same exemplar index represented by the time of acquisition. The traces for the two different tracks are plotted in different colors. The error correlation along the same track is in blue, while the error correlation for the adjacent track is in red. (d) Tuned *Rmod* as represented by the sample index for the same two tracks as in (a). (e) Tuned *Rmod* for the same tracks represented by the physical location. (f) Tuned *Rmod* for the same exemplar index represented by the time of acquisition. Note that the tuning has virtually eliminated cross-track error correlation. 91

Figure 5.1 – The inverse GMF (red) and its partial derivative with respect to NBRCS (blue) for a fixed incidence angle $\theta = 30$ deg. Note that as NBRCS approaches 0, the retrieved windspeed increases substantially, as does its sensitivity to errors in NBRCS. At higher NBRCS, the sensitivity to errors approaches 0. The FDS GMF is valid for windspeeds < 25 m/s. 104

Figure 5.2 – Contribution to error covariance by propagating instrument error sources into observation space using the GMF. The colors correspond to different windspeeds: 5 m/s (red), 7 m/s (blue), and 15 m/s (green). The line styles correspond to observation incidence angle: 10 degrees (solid), 30 degrees (dashed), 45 degrees (dotted). 105

Figure 5.3 – (a) A 2-dimensional histogram of MWP versus ERA5 windspeed. The color shading of the bins represent the frequency of occurrence to illustrate the shape of the overall distribution of windspeed dependence on MWP. A linear regression is computed (magenta) through the region of highest count density ($N > 1e4$) to minimize the effects of outliers. This regression represents the windspeed dependence of MWP. The MWP anomaly is calculated as

distance from the regression for a given windspeed. (b) A 2-dimensional histogram with windspeed error versus MWP anomaly (note the change in axes). A linear regression is computed (red) to estimate the impact of MWP anomaly on WS error. 111

Figure 5.4 – (a) A 2-dimensional histogram with WSS versus MWP-corrected ERA5 windspeed. The color shading of the bins represents the frequency of occurrence to illustrate the shape of the overall distribution of (MWP-corrected) windspeed dependence on WSS. A linear regression is computed (magenta) through the region of highest count density ($N > 2e4$) to minimize the effects of outliers. This regression represents the windspeed dependence of WSS. The WSS anomaly is calculated as distance from the regression for a given windspeed. (b) A 2-dimensional histogram with MWP-corrected windspeed error versus WSS anomaly (note the change in axes). A quadratic regression is computed (red) to estimate the impact of WSS anomaly on windspeed error. 113

Figure 5.5 – A stacked area chart illustrating the error covariance partitions from this work. The blue region represents the observation error covariance from *Rinst* for a reference track at global average windspeed. The orange region represents the observation error covariance due to MWP. The yellow region represents the observation error covariance due to WSS. The three colored regions are defined in the model as independent and are plotted by stacking the contributions upon one another. The grey region is defined as undiagnosed representation-retrieval error in this model, and is constructed as the difference between *Ro* and the sum of the three defined error covariance contributions. *Ro* is traced in black and the corrected error covariance *Rcorr* is dotted in red, demonstrating the across-the-board improvement in observation error covariance due to the corrections. 116

Figure 5.6 – A stacked area chart illustrating the validation test for the error partitions. The orange region represents the observation error covariance due to MWP. The yellow region represents the observation error covariance due to WSS. The error covariance due to the correction *Ro - corr* is plotted in purple. The overall shape and magnitude between the combined MWP and WSS error covariance and the calculated *Ro - corr* are reasonably consistent, which suggests that the retrieval improvement is the result of the correction of these specific error sources. 118

Abstract

Observations represent the largest cost driver in the environmental prediction enterprise. This dissertation aims to enhance the return on investment for satellite observatories by maximizing both the quantity and quality of data collected. Taking NASA's Cyclone Global Navigation Satellite System (CYGNSS) as a case study, this work proposes several novel contributions to extract the most value in observation data.

CYGNSS uses a technique known as Global Navigation Satellite System – Reflectometry (GNSS-R), which is an opportunistic bistatic radar measurement of the Earth's surface using GNSS as a signal-of-opportunity. This technique enables several surface sensing products for both land and ocean parameters. This dissertation primarily focuses on CYGNSS's ocean surface windspeed measurement, which drove the design of the mission architecture and performance objectives.

This thesis is structured into four distinct, but related lines of effort. The first is an evaluation of the CYGNSS onboard calibration system that substantially increased the amount of usable science data. Pre-launch design decisions and conservative estimates of the thermal loading led to a suboptimal calibration sequence that significantly impacted science duty cycle. This work proposed a longer calibration cadence with minimal data quality degradation, increasing the science duty cycle from 90% to 98%.

The second line of effort produces a mechanism to price the cost of representativity error in satellite observations that are not exactly simultaneous and collocated. If two measurements

are nearby in space and time, but not *exactly* simultaneous or collocated, it begs the question whether the two measurements are representative of the same target. Applying a technique from optimal interpolation, this work produces a simple metric that can be employed by satellite observation planners for future missions, which may feature proliferated constellations of disaggregated sensors.

The third line of effort builds and validates a physical model of correlated error structure in GNSS-R. In addition to building up a “bottom-up” error inventory for GNSS-R, this construction is also generalizable to other observation types. This is especially useful for numerical weather prediction, as data assimilation schemes frequently discard large amounts of observations because the correlated error structure is not well-defined. This is particularly problematic for GNSS-R due to its unique sampling characteristics.

The fourth and final line of effort builds upon the instrument correlated error model and evaluates the CYGNSS windspeed observation error covariance in observation space. This analysis also produced a method to partition certain sources of representation-retrieval error and suggests new corrections to the CYGNSS windspeed product that improves retrieval RMSE by 11%.

This work focuses on not just enhancements to CYGNSS, but also analysis to explore the structure of certain sources of error or impairment. Taken together, this dissertation both increases the quantity of useful CYGNSS data and makes CYGNSS data more useful.

Chapter 1 Introduction

This chapter consolidates and harmonizes introductions from (Powell, Ruf, and Russel 2022; Powell, Ruf, Gleason, et al. 2024; Powell, Ruf, McKague, et al. 2024) to unify notation and reduce repetition in subsequent chapters.

1.1 Motivations

Building Earth observation satellites is an expensive endeavor. The United States National Oceanic and Atmospheric Administration (NOAA) and National Aeronautics and Space Administration (NASA) each spend billions of dollars every year to develop, assemble, and launch environmental observation satellites (Lipiec 2023; Morgan 2024). These assets can range from the size of shoeboxes to school buses, and when combined with the terrestrial infrastructure for command and control, data processing, and product distribution, represents a capitalized infrastructure worth tens of billions of dollars.

The value of the data is harder to price. Space-based environmental observation infrastructure is critical to predicting the weather, monitoring a changing climate, and delivering warnings that save lives and protect property. Weather and climate data underpin entire sectors of the economy, such as agriculture, transportation, and hospitality (World Meteorological Organization 2015). Some valuations place the impact of weather variability on U.S. gross domestic product as high as 3.4% (Lazo et al. 2011).

The central premise of this dissertation is to improve both sides of this cost-benefit analysis. Given the intrinsic value of weather and climate information to society, and further the

extensive costs associated with designing, deploying, and operating environmental observation infrastructure, it is wise to extract the most value from the data.

Increasing return on investment for environmental satellites can take many forms. One way is to increase the amount of usable data. Another is to increase the overall utility of the data. This thesis performs both tasks. Taking NASA's Cyclone Global Navigation Satellite System (CYGNSS) as its primary case study, this work demonstrates four separate investigations that increases both the quantity and utility of satellite observations.

1.2 The CYGNSS Observatory

The CYGNSS mission was selected in 2012 as the first NASA Earth Ventures flight project (Harrington 2012). Comprised of eight small satellites in equatorial Low Earth Orbit (LEO), each approximately 29 kg in mass and the size of a microwave oven, CYGNSS was designed and optimized to measure surface windspeeds in tropical oceans (Ruf et al. 2016). As NASA's first small Earth science satellite mission, CYGNSS represented a new class of Earth observation capabilities: a low-cost constellation of smaller satellites enabled by a miniaturized sensing payload that takes advantage of existing signals of opportunity using a technique known as Global Navigation Satellite System – Reflectometry (GNSS-R) (Zavorotny et al. 2014). GNSS-R exploits the presence of known and well-characterized transmitters such as the Global Position System (GPS) to enable bistatic radar measurements of reflected surface quantities. CYGNSS measures GPS L1 signals in the L-band at 1575 MHz, which is especially useful for tropical cyclone observations, as thick cloud decks and precipitation are largely transparent in the L-band.

CYGNSS has been collecting GNSS-R observations since shortly after its launch in 2016, supporting a variety of land and sea applications, including ocean surface winds (Clarizia

and Ruf 2016b; Mayers and Ruf 2019), ocean surface microplastic detection (Evans and Ruf 2022), ocean altimetry (Carreno-Luengo et al. 2017; Li et al. 2020), soil moisture (C. C. Chew and Small 2018), inland flood inundation (Chew, Reager, and Small 2018), methane emissions (Gerlein-Safdi et al. 2021), and freeze-thaw state (Carreno-Luengo and Ruf 2022).

1.2.1 Sampling Characteristics and Constellation Architecture

CYGNSS's sampling is qualitatively distinct from most other observations. As an opportunistic measurement, the CYGNSS instrument has no control over the sample location and does not “image” in the traditional sense. The georegistration of the observation is instead a function of geometry and time: as GPS in medium-Earth orbit (MEO) and CYGNSS in LEO propagate in their respective orbits and the Earth rotates underneath, CYGNSS collects the GPS specular reflections from Earth's surface. This results in streaks or “tracks” of observations as a series of samples that share a GPS transmitter and CYGNSS receiver in short succession. Each of the 8 CYGNSS satellites captures four specular points at any given time, so the constellation is collecting 32 simultaneous tracks of observations across the Earth surface. Because CYGNSS and GPS are not synchronized in orbit, the sampling characteristics is best described by a probability distribution.

At +/- 35 degrees inclination, the CYGNSS constellation was designed to enable fast revisit windspeed observations of developing tropical cyclones for initialization in weather forecasts. Since CYGNSS satellites are all in the same orbital plane, tracks of data will “train” and follow in quick succession. This result is a median 2.8 h, mean 7.2 h revisit for a given sample (Ruf et al. 2022). Images are constructed from overlaying observations from multiple satellites within some predefined neighborhood in space and time. Figure 1.1 illustrates this effect with the global observations over 90 minutes (a), 24 hours (b), and the probability

distribution of revisit time (c). These sampling characteristics are fundamental to the challenges of making CYGNSS observations useful for weather models, described in depth in Chapters 2 and 3.

The satellites have no onboard propulsion, and are controlled via differential-drag techniques (Bussy-Virat et al. 2019). This maneuver was performed upon deployment to phase the CYGNSS satellites such that they roughly follow each other by 15 minutes in the 95 minute orbit. This has enabled certain experiments and calibration activities that depend on orbit phase, such as those described in Chapter 4.

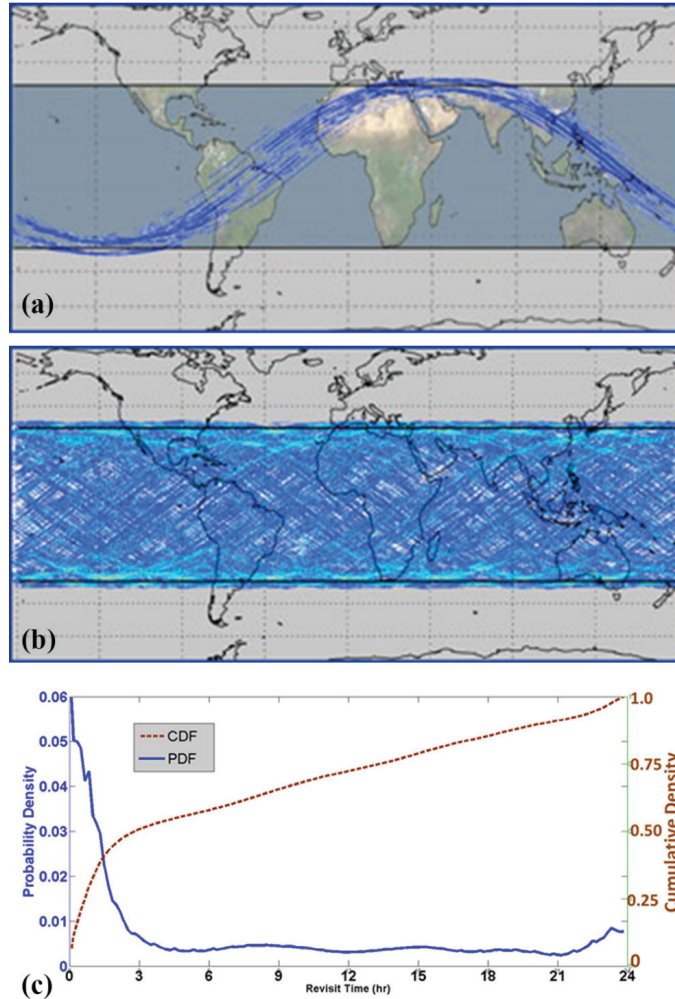


Figure 1.1. The sampling characteristics of CYGNSS from Ruf et al. (2022). (a) 90 min sampling window for eight CYGNSS satellites. (b) 24 hours of CYGNSS observations. (c) the probability density (blue) and cumulative density (orange) of revisit time for a given observation.

1.2.2 The Delay-Doppler Map

CYGNSS's primary payload is the Delay-Doppler Mapping Instrument (DDMI). The DDMI collects and processes the science data for the mission by using both the direct and reflected GPS signal. The direct signal enables correlation with the specific GPS pseudorandom number (PRN), which enables the DDMI to process the constellation almanac to identify GPS ephemerides. From there, the DDMI can calculate the anticipated delay and Doppler shift for a

specular-reflected signal “search” for reflection in delay-Doppler space. A diagram of this observing geometry is shown in Figure 1.2. The isodelay lines represent lines of constant delay, which is measured in the time domain to account for the extra path length of the reflection. The iso-Doppler lines represent lines of constant Doppler shift, accounting for the differences in relative velocities of the transmitter, receiver, and reflection surface.

As suggested by the name, the native coordinate system for the DDMI is delay-Doppler space, producing what’s known as Delay-Doppler Maps (DDMs). DDMs map non-uniquely to physical space, as shown in Figure 1.3. A line of non-ambiguity exists along the axis of the specular point in the Doppler domain. The DDMI performs most of the calculation of the reflection geometry and specular point search on-orbit, and only the resolved specular point data is downlinked for ground processing at a resolution of 17 delay bins x 11 Doppler bins. Each delay bin is approximately $\frac{1}{4}$ chip of the L1 code, and each Doppler bin is approximately 500 Hz (Ruf et al. 2022). The physical area represented in the DDM is approximately 10,000 km² (Ruf et al. 2022), but varies depending on the geometry of the observation, including incidence angle and altitude (Clarizia and Ruf 2016a). In practice, most CYGNSS derived products only use the information contained in the bins representing the specular point.




Figure 1.2 – A diagram representing the geometry of the CYGNSS GNSS-R observation. A GPS satellite travelling at velocity v_t emits a signal that is received both directly and via specular reflection by a CYGNSS receiver travelling at velocity v_r . The CYGNSS receiver also receives a direct signal from the GPS transmitter to enable geolocation and DDM processing. The specular point is mapped by isodelays (purple) and iso-Dopplers (blue). Different points P and Q are mapped to the same delay and Doppler bin. Figure from Ulaby and Long (2014).

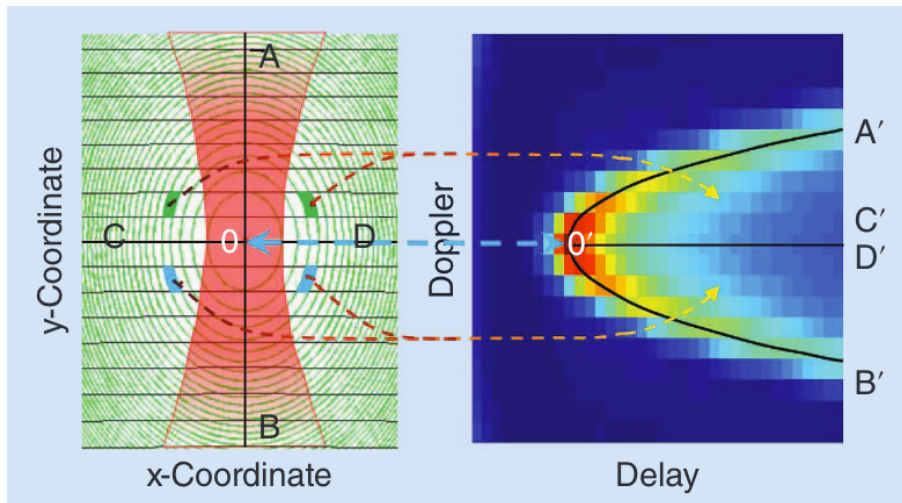


Figure 1.3 – A diagram illustrating the mapping between physical space and delay-Doppler space. The left panel shows a physical measurement space with isodelays (green) and a Doppler hyperbola. The right panel illustrates the mapping to delay and Doppler. Points A, B, C, D map to bins A', B', C', and D', respectively. Figure from Zavorotny et al. (2014).

1.2.3 The CYGNSS Observatory

The exterior of each CYGNSS satellite has two nadir science antennas, which measure the reflected GPS signal, and one zenith antenna which measures the direct signal. Each satellite deploys two solar panels after separation from the launch vehicle, which are used for both power and the differential drag maneuvers. There are also communication antennas and various launch vehicle interfaces on the satellite exterior. The receiver and DDMI electronics are packed in the interior volume of the satellite. There is a high degree of thermal isolation between the exterior of the satellite and the interior volume. The edges of the solar panels can experience thermal variations upwards of 100 deg C within an orbit, while the interior volume may only vary by 10 deg C. Chapter 2 explores the thermal behavior of CYGNSS in depth as it relates to observation quality and quantity. Figure 1.4 displays the component locations and physical construction of CYGNSS.

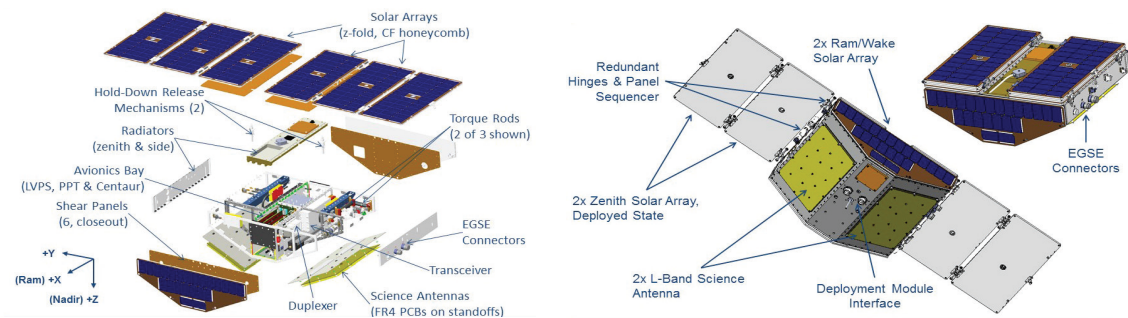


Figure 1.4 - Exploded view (left) and underside view (right) of CYGNSS. Courtesy Keith Smith, SouthWest Research Institute, from Ruf et al. (2022).

1.3 GNSS-Reflectometry for Ocean Wind Sensing

The bulk of this work focuses on the CYGNSS ocean surface wind product. The GNSS-R technique uses the reflected power from bistatic surface forward scattering as a proxy for mean ocean surface roughness, which is related to local windspeeds. Mean ocean surface roughness is estimated from the main instrument observable, the normalized bistatic radar cross section (NBRCS). The NBRCS is determined from the scattered power received at the CYGNSS satellite by the radar-range equation (Gleason et al. 2016):

$$\sigma^o = \frac{P_g (4\pi)^3 L^{atm} I}{P^T G^T \lambda^2 G^R L^R A} \quad (1.1)$$

where σ^o is NBRCS in normalized units [m²/m²], P_g is the received power; L^{atm} is the atmospheric attenuation along the propagation path from the GPS satellite to the specular point and to CYGNSS; I is a term to account for instrument losses, L^R is the range loss of the reflected transmission via the Earth specular point; P^T is the transmit power of the GPS satellite; G^T is the antenna gain of the GPS satellite in the direction of the specular point; λ is the wavelength of the GPS L1 signal; G^R is the gain of the CYGNSS nadir receive antenna in the direction of the specular point; and A is the effective scattering area. Note that P_g and A vary as functions of the time delay and Doppler shift of the received signal. As a result, NBRCS also depends on delay and Doppler. σ^o is calculated from the average of σ over a 3-by-5 bin window centered at the specular point and is divided by the effective area A corresponding to the surface area bounded by the 3-by-5 window.

After the launch of the CYGNSS mission, it became clear that uncertainty in the GPS effective isotropic radiated power (EIRP) would significantly impact calibration quality, (Gleason et al. 2019; Wang, Ruf, McKague, et al. 2021; Wang et al. 2020). The EIRP is simply

the product $P^T G^T$ from Equation (1.1). Errors are present in both terms: not only is there general uncertainty in the actual (versus published) transmit power of GPS satellites, but newer generations of the constellation operate with a flexible power mode and dynamically change transmit power levels across their respective orbits (Steigenberger, Thölert, and Montenbruck 2019). Further, only a subset of the GPS antenna patterns was published (Marquis and Reigh 2015), and they were not sufficiently detailed to constrain the error in EIRP for CYGNSS calibration.

To remedy this, Wang et al. developed a novel calibration technique that uses the onboard zenith CYGNSS antenna and empirically derived relationships from measured GPS antenna patterns to estimate GPS EIRP in the direction of the Earth’s specular point at the time of observation (Wang, Ruf, Gleason, et al. 2021). This new technique has been adopted in the CYGNSS Level 1 calibration algorithm since version 3.0 and modifies the radar equation to become the following (Gleason 2020):

$$\sigma^o = \frac{P^g 4\pi L^{atm} L^Z G^Z}{P^Z \zeta^E G^R L^R \bar{A}} \quad (1.2)$$

where all the terms are the same as in Equation (1.1), with the additions of L^Z , the zenith transmission range loss from the GPS source to the CYGNSS zenith antenna (written as a loss in the numerator to distinguish from the L^R range losses); the CYGNSS zenith antenna gain G^Z in the direction of the GPS satellite; the power received from the CYGNSS zenith receiver P^Z ; and the specular–zenith ratio ζ^E for the GPS EIRP as described in (Wang, Ruf, Gleason, et al. 2021).

NBRCS is directly related to the mean square slope (MSS) of ocean wave spectra via (Ruf et al. 2016; Clarizia and Ruf 2016b):

$$\sigma^o(\theta) = \frac{|\Re(\theta)|^2}{MSS} \quad (1.3)$$

where θ is the angle of incidence of the scattering geometry and $\mathfrak{R}(\theta)$ is the Fresnel reflection coefficient for the ocean surface at the specular point. CYGNSS retrieves wind speed from σ^0 via an empirical Geophysical Model Function (GMF) (Clarizia and Ruf 2016b; Ruf and Balasubramaniam 2019) that is processed operationally via a stored look-up table (LUT).

Stronger received signals P_g indicate smoother, more specular scattering of the ocean surface; weaker signals indicate rougher ocean surfaces associated with higher windspeed. CYGNSS’s electronics measures received power in raw digital counts, and calibrated received signal power P_g in units of Watts is calculated by:

$$P_g = \frac{C - C_N}{G} \quad (1.4)$$

where G is receiver gain in Watts/count, C is the raw counts measured at delay-Doppler bins where a scattered signal is present, and C_N is the mean raw counts for background noise without any scattered signal present. In practice, C_N is an average over many delay-Doppler bins at delay coordinates shorter than that of the specular point. Shorter delays correspond to radar reflections from the atmosphere above the ocean surface, for which there is no appreciable scattering at L-Band.

1.4 Data Assimilation Basics

Weather models are fundamentally initial value problems. The mechanics of the atmosphere are generally well-known. A host of observations from terrestrial-, sea-, air-, and space-based platforms are combined to create an initial condition of model atmosphere, known as the analysis. The model then “plays the tape forward” and propagates the virtual atmosphere of the analysis into forecast increments with known behavior about the dynamics of atmospheric

phenomena. Data assimilation (DA) is the step of a numerical weather prediction (NWP) system that combines the observations into the forecast models. A common DA technique poses the problem as a variational equation that minimizes a cost function (Lorenc et al. 2000):

$$\operatorname{argmin}_{\mathbf{x}} \mathbf{J}(\mathbf{x}) = (\mathbf{x} - \mathbf{x}_b)^T \mathbf{B}^{-1}(\mathbf{x} - \mathbf{x}_b) + (\mathbf{y} - H[\mathbf{x}])^T \mathbf{R}^{-1}(\mathbf{y} - H[\mathbf{x}]) \quad (1.5)$$

where $\mathbf{J}(\mathbf{x})$ represents the cost function, \mathbf{x} is the variational argument, \mathbf{x}_b is the background model state of dimension n , \mathbf{B} is the n -by- n background error covariance matrix, \mathbf{y} is the observation vector of dimension p , H is the forward model operator for transforming from model space to observation space, and \mathbf{R} is the p -by- p observation error covariance matrix. In practice, a great deal of focus has been placed on estimating \mathbf{B} in the weather community (Bannister 2008a; 2008b) because an incorrectly specified \mathbf{B} can have significant impacts on the performance of the model. Famously, introducing perfect observations to a perfect model can still lead to forecast degradations if \mathbf{B} is poorly constructed (Morss and Emanuel 2002). Chapters 4 and 5 of this thesis focus on \mathbf{R} , which has historically received less extensive focus in the DA community.

1.4.1 What Happens When \mathbf{R} is Poorly Specified

DA schemes often assume that each observation error is completely uncorrelated (Stewart, Dance, and Nichols 2013; Fowler, Dance, and Waller 2018; Liu and Rabier 2002) and, as such, treat \mathbf{R} operationally as a diagonal matrix. This is performed for three main reasons: to simplify the assumptions needed for minimizing the cost function in Equation (1.5) and reduce the computational resources required to solve the forecast problem; to reduce the data handling and storage requirements for observation data at low latencies; and because frequently the information required to specify correlated observation errors was unavailable or difficult to ascertain. To achieve this, satellite observations are frequently thinned or ‘super-obbed’ to ensure

that samples are independent and thereby contain uncorrelated error (Bauer et al. 2011; Gao et al. 2019; Hoffman 2018).

Regardless of the rationale, the under-specification of correlated observation error in \mathbf{R} leads to the suboptimal use of the information content in observations (Rainwater, Bishop, and Campbell 2015). This is especially true for CYGNSS because, given the unique sampling characteristics of the mission, it negates some of the observatory's unique value proposition of high time and spatial resolution of tropical cyclone cores. While thinned and super-obbed CYGNSS measurements have been shown to improve tropical cyclone forecast skill (Zhang et al. 2017), this practice still reduces the overall value the observatory can provide to the forecasting community.

For observation errors to be represented as independent and uncorrelated, \mathbf{R} would only have nonzero elements along the main diagonal. In contrast, physically representative \mathbf{R} matrices would specify nonzero off-diagonal elements. This work suggests a mechanism to estimate these off-diagonal elements as a function of space and time for GNSS-R.

Investigations of correlated observation error have been explored for a number of different observatories including interchannel correlations in infrared (Stewart et al. 2014; Bormann et al. 2016) and microwave sounders (Bormann and Bauer 2010; Campbell et al. 2017). In other studies, spatial error correlations have been explored for Doppler radars (Waller et al. 2016a), geostationary imagers (Waller et al. 2016b), and atmospheric motion vectors (Cordoba et al. 2017). There are several ways to estimate the observation error correlation statistics. One of the most common techniques is by estimating \mathbf{R} is from observation-minus-background and observation-minus-analysis residuals from assimilated forecasts (Desroziers et al. 2005), which is relatively simple to compute for data that is routinely assimilated into

forecasts. Other techniques include applying a “three-cornered hat” (Sjoberg, Anthes, and Rieckh 2021; Semane et al. 2022), which requires collocation of three independent datasets, but this technique generally assumes the three datasets have uncorrelated errors. Recently, “bottom-up” approaches of constructing error inventories have been introduced to provide estimates of correlated error structure from first principles (Smith et al. 2021; Yang et al. 2023).

1.5 Chapter Summary

Chapter 2 explores how a single assumption in the CYGNSS mission development led to suboptimal observation duty cycles. This chapter improves the science duty cycle by analyzing observation quality degradation associated with imperfect knowledge of thermally-dependent gain variations, and as a result, suggests specifying a more appropriate blackbody calibration cadence. The content of this chapter is largely drawn from (Powell, Ruf, and Russel 2022).

Chapter 3 explores the nature of representativity error and suggests a mechanism for pricing the amount of error incurred by imperfect matchups in space and time. The content of this chapter is largely drawn from (Powell, Ruf, Gleason, et al. 2024).

Chapter 4 derives an instrument observation error covariance model from first principles for CYGNSS and validates the model with a unique orbital condition in which two CYGNSS assets were in a near-overlap observation geometry. The content of this chapter is largely drawn from (Powell, Ruf, McKague, et al. 2024).

Chapter 5 propagates the instrument error covariance modelled in Chapter 4 into observation space and suggests a method for partitioning CYGNSS representation-retrieval error into geophysical components. This chapter also identifies new opportunities to correct the CYGNSS windspeed retrieval with ancillary datasets. The content of this chapter is largely drawn from

(Powell and Ruf 2024).

Chapter 6 summarizes this work and suggests future avenues of investigation. This chapter presents analysis suggesting that CYGNSS had been initially too conservative in its calibration strategy, sampling its onboard blackbody once per minute, and recommends decreasing the calibration frequency to once every ten minutes to improve the science data duty cycle. This chapter is substantially derived from a work published in *IEEE Transactions on Geoscience and Remote Sensing* under the title “An Improved Blackbody Calibration Sequence for CYGNSS” (Powell, Ruf, and Russel 2022). The research is presented as published, with deletions of introductory material and theory (covered in Chapter 1), updates to formatting, and minor notation changes.

1.6 Abstract

An improved blackbody calibration procedure is developed, implemented, and tested for the Cyclone Global Navigation Satellite System (CYGNSS). Previously, CYGNSS calibrated its receivers once every minute to account for temperature-induced gain fluctuations. The time spent making calibration measurements limited the duty cycle of windspeed measurements to approximately 90%. Analysis presented here shows that the one-minute cadence was overly conservative and can be increased to once every 10 minutes with minimal impact to data quality, thereby improving the windspeed duty cycle to 98%. A permanent change to the blackbody cadence was made for the complete 8 satellite constellation during July 27, 2021 – August 3, 2021, and subsequent analysis verifies that the new cadence improves duty cycle without impacting science data quality, as expected.

Chapter 2 Improving CYGNSS Duty Cycle

This chapter presents analysis suggesting that CYGNSS had been initially too conservative in its calibration strategy, sampling its onboard blackbody once per minute, and recommends decreasing the calibration frequency to once every ten minutes to improve the science data duty cycle. This chapter is substantially derived from a work published in *IEEE Transactions on Geoscience and Remote Sensing* under the title “An Improved Blackbody Calibration Sequence for CYGNSS” (Powell, Ruf, and Russel 2022). The research is presented as published, with deletions of introductory material and theory (covered in Chapter 1), updates to formatting, and minor notation changes.

2.1 Introduction (Abridged)

CYGNSS’s utility is impaired by a suboptimal sampling cadence related to the satellite’s onboard calibration system. The sequencing rate for sampling the onboard blackbody calibration target had been once every minute since launch, and each time this event occurred, the calibration would last approximately four seconds, inclusive of the time required to switch loads from the antenna source to the onboard blackbody target. As a result, CYGNSS occasionally missed sampling certain high-value targets.

Based on the analysis presented below, the flight software was updated in 2021 to modify the sampling cadence to once every 10 minutes, with a slightly increased dwell time of approximately six seconds. This improves the windspeed retrieval duty cycle from 90% to 98% and shows no detectable deterioration in performance.

2.2 Theory (Abridged)

CYGNSS satellites are in an equatorial low-Earth orbit with a period of approximately 95 minutes, and experience significant variations in their thermal loads associated with their orbit. As each satellite crosses the terminator into sunlight, the body heats unevenly, causing thermal gradients across the structure. As the spacecraft enters the nightside, it cools suddenly and unevenly.

As the body of the structure changes temperature, so does the thermal environment of the CYGNSS receiver electronics. These dynamics induce variations in the gain of CYGNSS's receiver. Because the quality of windspeed retrieval depends on the consistency of receiver power, each satellite's nadir instruments have built-in onboard calibration equipment to account for thermal gain variations in the receiver amplifiers.

CYGNSS has two nadir science antennas, facing the satellite's port and starboard directions, respectively. The receiver connected to each nadir science antenna has its own blackbody calibration target. A thermistor adjacent to the target monitors its temperature in real time to determine the power in the blackbody thermal emission. During a blackbody calibration sequence, the input to the receiver is redirected from the science antenna to the blackbody load and the power emitted by the blackbody is recorded in raw counts for later processing. This sequence takes approximately four seconds because the calibration clock and the science data clock are not synchronized, and the switching of loads is not perfectly instantaneous. To ensure that a full reading of the blackbody is measured without contamination from the load switching process, samples immediately preceding and succeeding calibration are flagged out.

During data processing, the nearest calibration samples before and after a science sample are linearly interpolated to the time of the science sample to estimate gain variations between

calibration looks. At launch, this sequence was scheduled to occur once every minute, such that every science sample was within approximately 30 seconds of a receiver calibration event. This calibration sequence sets receiver gain for calculating received power as given by

$$G = \frac{C_B}{P_B + P_r} \quad (2.1)$$

where C_B is the counts measured while looking at the blackbody load, P_B is the power from the blackbody as estimated from the thermocouple on the low noise amplifier, and P_r is the receiver noise estimated from a temperature-dependent noise floor parameterization determined pre-launch. All three values vary as the spacecraft temperature fluctuates. The one-minute cadence was estimated to bound the 1-sigma calibration error for the received power to 0.13 dB (Gleason et al. 2019).

Fig. 2.1 shows how the spacecraft gain varies with temperature. The blue trace is the calibrated receiver gain in dB, the red trace is variation in low-noise amplifier temperature in degrees Celsius. The gain varies approximately 0.4-0.6 dB peak-to-peak, whereas the absolute temperature of the low-noise amplifier varies approximately 8 degrees Celsius peak-to-peak.

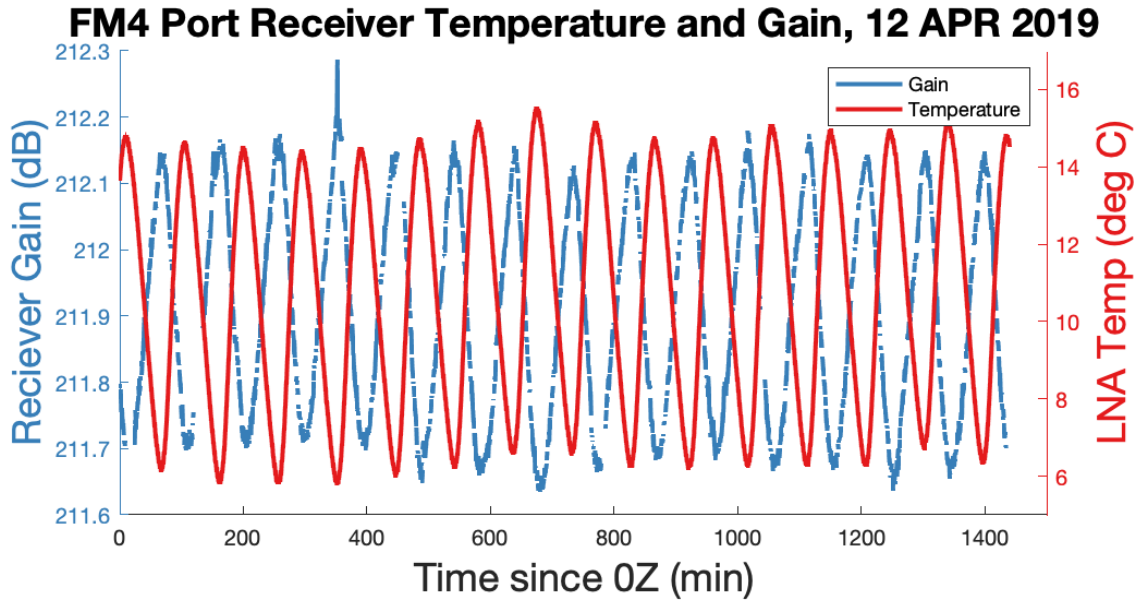


Figure 2.1 – Example of temperature influence on receiver gain. Above, FM4’s port receiver gain is plotted (blue) with the temperature at a nearby probe (red) for a full day starting at 0Z 12 APR 2019. Note that gain is inversely related to local temperature, and that these swings are periodic, consistent with CYGNSS’s 95-minute orbital period.

CYGNSS’s thermal environment also varies due to the orientation of its orbit plane relative to the sun. The relative time spent in sunlight or shade varies throughout the year as the orbital plane precesses about the Earth. This precession is conveniently characterized by the time dependent nature of the angle between the orbit plane and a line from the Earth to the sun, referred to as the orbit beta angle (see Fig. 2.2). In general, CYGNSS is exposed to the greatest temperature variations during low magnitudes of orbit beta angle, which implies that the spacecraft will spend the maximum amount of time in shade (Grey et al. 2020). Additionally, during periods at the highest beta angles, at approximately 50 degrees in magnitude and greater, the spacecraft are commanded into a roll configuration to maintain adequate solar radiance on the solar array panels.

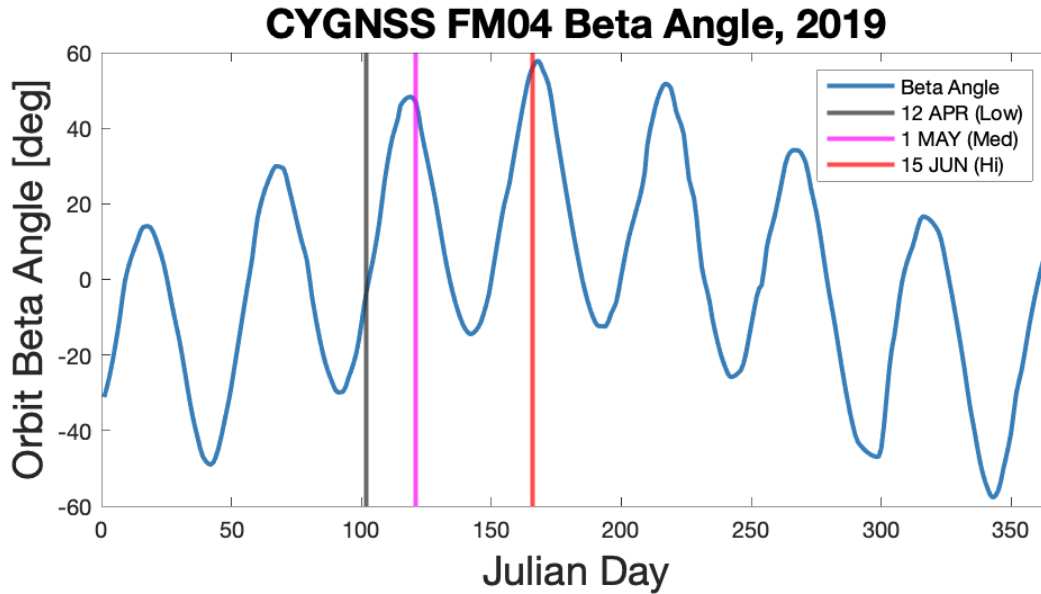


Figure 2.2 – Orbit beta angle calculated for FM04 for 2019. The black, magenta, and orange vertical lines denote the low, medium, and high beta angle days respectively.

2.3 Methods

Decreasing the calibration frequency risks increasing error caused by inaccurate compensation for time-varying receiver gain. Science data taken between blackbody calibration measurements are calibrated using linearly interpolated values of the nearest calibration data before and after the science measurement. If the temperature of the spacecraft receiver electronics changes in a non-linear way between calibrations, it can result in significant errors due to improper corrections for gain variation. An on-orbit experiment was performed to explore how two fundamental CYGNSS data products, NBRCS and retrieved windspeed, degrade due to an increase in the time between blackbody samples.

Three days of data from one satellite (FM4) were used, with days selected representing different characteristic beta angles to explore how different thermal cycling may impact the calibration sequence, as shown in Table 2-1.

2.4 Estimated Beta Angle and Thermal Variation for Selected Days

Table 2-1. Exemplar temperature variability for different orbit beta angles.

Date	Orbit Beta Angle	Peak-to-peak Temperature Swing
12 APR 2019	-3 deg ("Low")	9.8 deg C
1 MAY 2019	46 deg ("Med")	6.6 deg C
15 JUN 2019	56 deg ("High")	5.6 deg C

For all three days, science samples were collected and processed with normal one-minute blackbody sampling. This generated baseline measurements: the fully-developed sea (FDS) windspeed retrieval product u_{FDS_1} , the young-seas limited-fetch (YSLF) windspeed retrieval product u_{YSLF_1} , and the NBRCS σ_1^o (CYGNSS 2020b). CYGNSS uses two separate geophysical model functions for windspeed retrieval. Both retrievals are empirical fittings, but the YSLF differs from the FDS in its sensitivity to longwave swell, which tends to be under-developed in high-wind, dynamic weather where wind direction is frequently changing, such as during tropical cyclones (C. S. Ruf and Balasubramaniam 2019; C. S. Ruf, Gleason, and McKague 2019). The YSLF product is tuned to respond to higher wind speed conditions where measurement sensitivity is diminished, and so will be more sensitive to possible degradation in calibration quality that results from the change in blackbody cadence.

These data were then reprocessed using every n th blackbody sample, where n is an integer that ranged from 2 through 45, representing sampling the blackbody once every n minutes. The maximum value considered corresponds to performing a blackbody calibration approximately once every half orbit. Between the blackbody samples used, the raw counts recorded by the blackbody were linearly interpolated. An additional case was considered in

which the daily mean value of all blackbody samples was used to calibrate the entire day of science data. This case represents calibration without regard for short term gain variations.

Degradation in calibration accuracy is assessed by examining the difference between science samples calibrated using every n th blackbody sample vs. using blackbody samples every minute. These differences are given by

$$\Delta u_n = u_n - u_1 \quad (2.2)$$

where u_n is the windspeed produced at a calibration period of n minutes, u_1 is the windspeed produced at the original calibration sequence, and Δu_n is the difference in windspeed due to increased calibration period of n minutes. If u_1 is assumed to be accurately calibrated, Δu represents the error that results from less-frequent blackbody calibration. Similar values can be produced for NBRCS and gain:

$$\Delta \sigma_n^o = \sigma_n^o - \sigma_1^o \quad (2.3)$$

$$\Delta G_n = G_n - G_1 \quad (2.4)$$

where σ_n^o is the NBRCS at n -minute calibrations, σ_1^o is the NBRCS produced during one-minute calibrations, and $\Delta \sigma_n^o$ is the NBRCS error induced by sampling at longer periods; and G_n is the receiver gain at n -minute calibrations, G_1 is the original gain at one minute sampling, and ΔG_n is the error due to a longer calibration sequence at period n .

If the use of a particular (longer) sampling rate produces a measurable change in the derived CYGNSS data products relative to the 1-minute baseline measurements, that is an indication of a degradation in calibration accuracy.

To explore a “worst case” scenario, an additional comparison is made assuming the satellite only calibrates the blackbody once per day using the daily average value. This provides a sense for how NBRCS and windspeed would degrade if the calibration was performed at timescales much slower than the orbit-induced changes in the thermal loading of the spacecraft.

2.4.1 Characterizing the Impact of Change in Cadence

Changing the period of calibration impacts the quality of the gain estimate in (4). If the non-linear component of the thermal environment changes at timescales faster than the calibration period, the gain estimated by the calibration process will no longer be representative of the true gain.

Fig. 2.3 depicts histograms of normalized gain changes due to increased blackbody sampling cadences. For every sample, both the original gain and the relative fraction of its change is calculated. At blackbody sampling rates close to one minute, there is very little difference from the original gain, and distribution of samples approach a delta function. At longer




Figure 2.3 – Histograms of gain error at $n = \{5, 15, 25, 35, 45\}$ minutes. The blue and red distributions represent data from the port and starboard antennas, respectively. At low n , the gain error approaches a delta function. As sampling cadence time increases, the distribution widens, illustrating the degradation from the original one-minute sampling cadence.

sampling intervals, this distribution widens, illustrating how gain estimates degrade as the sampling cadence increases. At approximately two samples per orbit, or $n = 45$, the gain variations can be up to +/- 5% of the baseline.

The results in Fig. 2.3 are shown for the impact on gain, but this behavior is observed across all derived products. To estimate the growth in error as a function of sampling cadence, the standard deviation is calculated for each distribution. Fig. 2.4 shows how the standard deviation grows as the sampling period increases.




Figure 2.4 – Standard deviations of growth in error as a function of sampling cadence for NBRCS (top), FDS windspeed (middle), and YSLF windspeed (bottom) at low beta angle. Red traces indicate values for the starboard antenna, and blue traces for the port side.

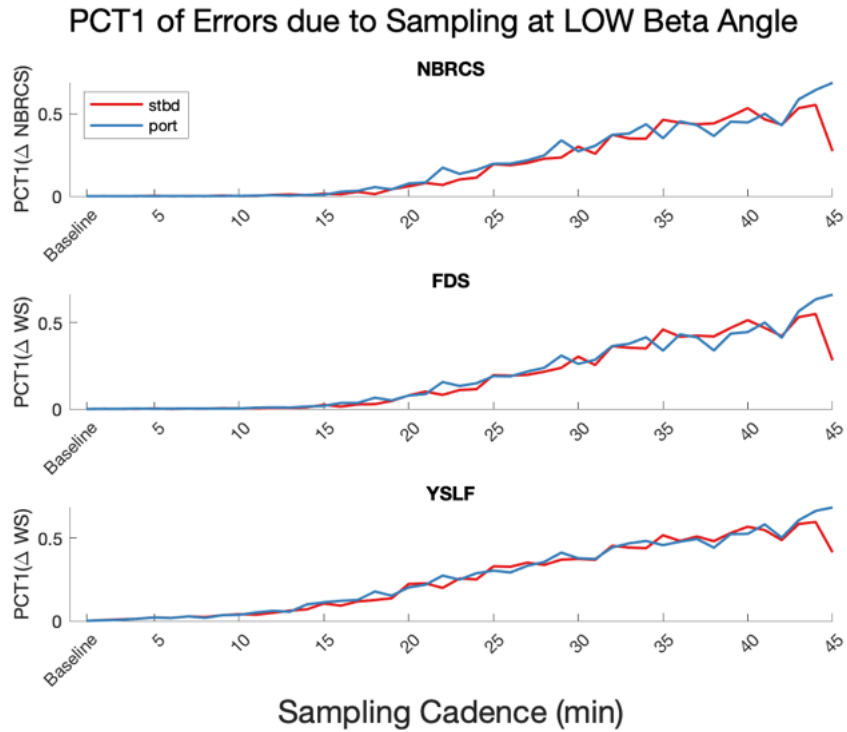


Figure 2.5 – PCT1 of error as a function of sampling cadence period for NBRCS (top), FDS windspeed (middle), and YSLF windspeed (bottom) at low beta angle. Red traces indicate values for the starboard antenna, and blue traces for the port side.

Sampling every 2 minutes, the NBRCS error standard deviation is approximately 0.25 m²/m². This grows to an error of just over 1 m²/m² at 45 min cadence. The error characteristics for sampling once per day resembles the 45-minute scenario, as the sampling at every half-orbit maximizes error due to unrepresentative temperature-dependent gain corrections. The standard deviation of windspeed error is approximately three times larger for the YSLF retrieval than the FDS retrieval. At the worst possible case, sampling once every day, the standard deviation of error on the windspeed product approaches 0.6 m/s. For context, the requirement for CYGNSS retrieval accuracy is a root-mean-square (RMS) error of less than 2 m/s at windspeeds lower than 20 m/s, and RMS error of less than 10% at winds greater than 20 m/s (Clarizia and Ruf 2016b). The error inventory for the mission suggests that the total L1 error will be 0.82 dB at low wind speeds and 0.70 dB at high winds (Gleason et al. 2016). Assuming a low wind NBRCS of 100 and a high wind NBRCS of 16, the once-daily sampling degrades the measurement by 0.04 dB and 0.28 dB respectively. Therefore, even at this cadence, the magnitude of the standard deviation of error is still small enough to meet system requirements.

The standard deviation can, however, be a misleading metric, as it describes the behavior of the total distribution, as opposed to errors in high-value samples. Changes in NBRCS are much more significant at low values, because NBRCS is inversely related to windspeed. In addition, large errors in windspeed retrievals in high-wind areas such as hurricanes can be averaged out by the much more frequent and smaller errors over calm seas.

For this reason, a new metric is developed to appropriately capture how modifications in calibration cadence can impact the CYGNSS science products. This metric is named PCT1, which stands for the percent of samples with error magnitudes that are 1% or greater than its




Figure 2.6 – PCT1 of NBRCS error as a function of sampling cadence period at various beta angles. At lower beta, the performance degradation is more pronounced.

value when using the original one-minute calibration cadence. The PCT1 metric ranges from 0 to 1, with 0 indicating that no samples have errors greater than 1%, and 1 indicating that every sample has an error which is at least 1%. The PCT1 for the 3 data products is shown in Fig. 2.5. As the cadence period increases, the PCT1 rises, with an inflection point between 10 and 15 minutes. At that point, the increase in blackbody cadence significantly impacts the overall population of data. As expected, the YSLF windspeed retrieval is the most sensitive product to changes in the blackbody sampling rate.

To see how this relationship varies across different orbit beta angles and resulting thermal conditions, the PCT1 for NBRCS is plotted across different orbit beta angles in Figure 2.6. At

lower beta angles, the errors are more pronounced, as the spacecraft spends more time in the shaded portion of the orbit, resulting in greater temperature variability.

2.4.2 Optimal Blackbody Sampling Cadence

Determining an optimal sampling cadence requires exploring the tradeoff between the potential benefit of increased science duty cycle over the cost of potentially poorer instrument performance.

For this analysis, the worst possible conditions are employed to provide an upper bound on potential impacts to downstream products. Because the low beta angle orbit has greatest temperature variability, it leads to increased performance degradation, and because the YSLF windspeed product is more sensitive to calibration errors than its FDS counterpart, the choice of blackbody sampling rate should be driven by the low beta YSLF error growth.

We chose an arbitrary YSLF PCT1 value of 0.05 as the optimal cutoff threshold, which means that the number of samples with greater than 1% error due to increased sampling period could be no more than 5% of the total population. The largest sampling rate below the PCT1 cutoff for the YSLF data product is once every 10 minutes as shown in Fig. 2.5. At 11 minutes, the PCT1 values exceed 0.05 during the low beta angle day.

STD(Δ YSLF WS) by Time Since Terminator (Low Beta, Port)

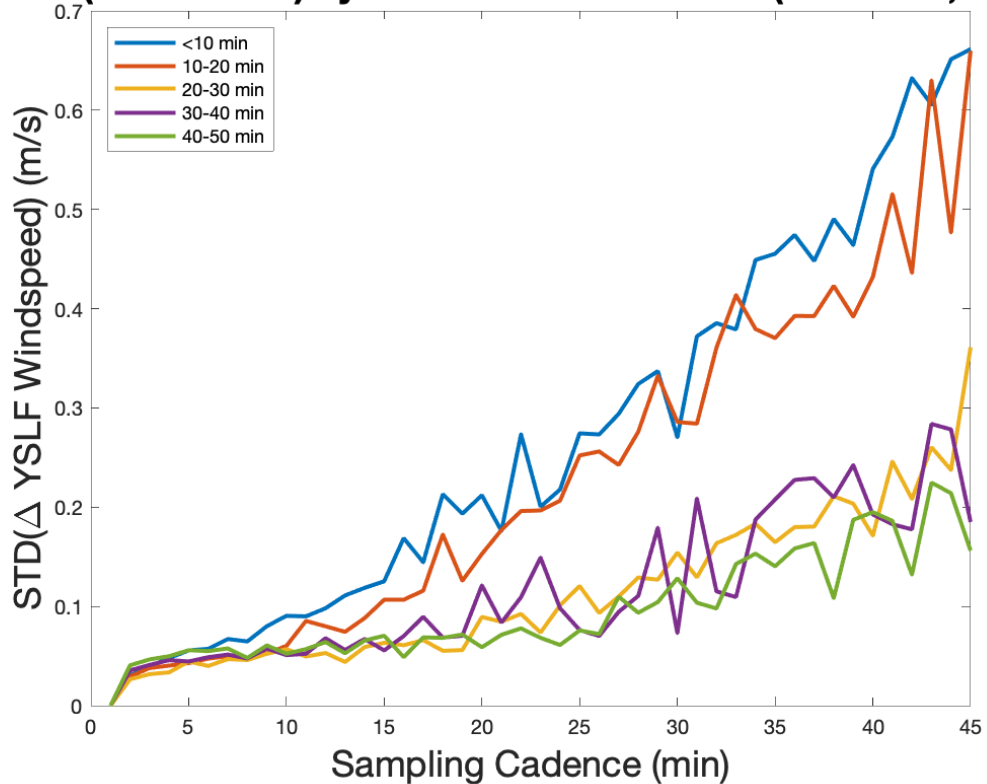


Figure 2.7 – Standard deviation of YSLF windspeed error as a function of sampling cadence period at various orbit sectors. The errors due to increased blackbody sampling rates are greatest immediately after crossing the terminator. This effect is largest at longer sampling periods.

2.5 Performance by Orbit Sector

The temperature gradients on CYGNSS are the largest as the spacecraft crosses the terminator. This implies that the performance degradation due to increased sampling period should be greatest as the spacecraft crosses the terminator. To validate this theory, we calculate the standard deviation of the windspeed anomaly as a function of sampling cadence as before, and further segregate results by orbital phase.

The terminator crossing is approximated by utilizing the temperature traces of the low-noise amplifiers. The spacecraft is assumed to cross the terminator five minutes prior to each peak and valley of the temperature trace. After the terminator crossings are established, the time since crossing the terminator is segregated into five 10-minute bins.

Fig. 2.7 shows the standard deviation of the YSLF windspeed anomaly for the low beta case, separated by time since terminator crossing. The performance impact of increasing the sampling cadence is clearly dependent on orbital location, with the worst performance immediately following the terminator crossing, and moderate impacts as the spacecraft enters a steady-state thermal environment.

2.6 Discussion

As a result of this analysis, the CYGNSS operations team elected to transition science operations to the 10-minute blackbody sampling cadence in August 2021.

2.6.1 Characterizing the Impact to Duty Cycle

To evaluate the utility of a refined sampling cadence, the preferred figure of merit is the duty cycle of high-quality ocean observations in which windspeed retrievals are possible, as given by

$$D_{fom} = \frac{N_{good}}{N_{good} + N_{BB}} \quad (2.5)$$

where N_{good} is the number of good-quality ocean observations capable of otherwise retrieving a windspeed retrieval, and N_{BB} is the number of samples impacted by the blackbody calibration sequence.

The original 1-minute blackbody cadence took on average four seconds, inclusive of the time it took for the spacecraft to switch loads (1 s), sample the blackbody target (2 s), and switch back to the science antenna (1 s). This would suggest a nominal duty cycle of approximately $56/60 = 93.33\%$ as computed in (8) without considering other factors.

For the 10-minute sequence, the length of time observing the blackbody target was increased to four seconds to reduce variability due to sudden receiver excursions. When considering the two seconds required for changing the antenna load to the target, an average blackbody measurement takes approximately six seconds. Theoretically, the duty cycle should improve to $594/600 = 99\%$.

CYGNSS L1 data files have several quality flags associated with each sample, and N_{good} is determined by an “overall quality” indicator, which is the sum of several other flags with logical OR operations as described in the product data dictionary (CYGNSS 2020a). In practice, only about 60% of samples collected over the ocean are useful due to other quality control flags independent of the blackbody calibration. Therefore, any improvement in the raw duty cycle with a new blackbody sequence will have outsized impact in D_{fom} . When accounting for this difference, the expected D_{fom} in the naïve 1-minute cadence where 6.67% data are lost is no longer 93%, but approximately 89% as $6.67\% / 0.6 = 11\%$. Similarly, the expected D_{fom} for the improved blackbody sequence is 98%.

N_{BB} is composed of logical ORs of flags associated with the blackbody sample itself plus flags that indicate the instrument is reconfiguring in preparation for or immediately after a blackbody sample, adjusted for samples that occur only over the ocean. By implementing this modification, the D_{fom} increases from approximately 90% to 98%.

2.6.2 Characterizing the Impact to the Level 1A Error Budget

CYGNSS’s Level 1 error budget described in the Algorithm Theoretical Basis Document is composed of the root sum of squares of individual error terms in the Level 1A equation (Gleason 2014). Combining (1.4) and (2.1) shows the full Level 1A equation with each source term

$$P_g = \frac{(C - C_N)(P_r + P_B)}{C_B} \quad (2.6)$$

The individual error terms can be approximated by taking the partial derivative with respect to each individual term

$$E(q_i) = \left| \frac{\partial P_g}{\partial q_i} \right| \Delta q_i \quad (2.7)$$

where q_i represents an individual input parameter. The total error from the root sum of squares of individual error terms can be expressed as

$$E_{L1A} = \left[\sum_i [E(q_i)]^2 \right]^{\frac{1}{2}} \quad (2.8)$$

In this experiment, we modified the input parameter C_B and held all other inputs constant.

Therefore, we can calculate the specific contribution of error due to the blackbody sequence by plugging (2.7) into (2.6) and setting $q_i = C_B$:

$$E(C_B) = \frac{(C - C_N)(P_r + P_B)}{C_B^2} \Delta C_B \quad (2.9)$$

The latest Level 1A error budget estimates that the 1-sigma $E(C_B)$ is 0.05 dB (Gleason et al. 2019). With the change in blackbody cadence from 1 minute to 10 minutes, the one-sigma uncertainty becomes approximately 0.07 dB. The total error now follows by computing the root sum of squares in (2.8) with the other values specified in Gleason et. al (2019) (Gleason et al. 2019). Prior to the blackbody calibration, the rolled-up Level 1A error budget was 0.225 dB, and after the change, it has increased to 0.232 dB.

2.6.3 Comparing Expected Results to Empirical Data

To demonstrate that windspeed performance has not significantly deteriorated, a statistical comparison is performed with reanalysis data using the CYGNSS Climate Data Record Version 1.1. This product matches up MERRA-2 windspeed reanalysis with CYGNSS observations, and generates an expected “modeled” NBRCS for a given ocean condition with the appropriate spectral corrections for CYGNSS. The performance of the CYGNSS retrieval is compared as the difference in measured NBRCS versus modeled NBRCS for each sample before and after the calibration sequence:

$$(\Delta\sigma^o)_{1\ min} = (\sigma_{obs}^o - \sigma_{mod}^o)_{1\ min} \quad (2.10a)$$

$$(\Delta\sigma^o)_{10\ min} = (\sigma_{obs}^o - \sigma_{mod}^o)_{10\ min} \quad (2.10b)$$

To account for variability in global windspeed distributions, 24 prior samples at the 1-minute cadence were selected in 2020 to account for any seasonal or sub seasonal variability in CYGNSS performance. Because year-round data is not available for the new sampling cadence, a week of data was collected to represent current behavior.

With this data, Student’s T-test was performed with both the standard deviation and mean of $\Delta\sigma^o$ to evaluate if there were any statistically significant changes in the performance of CYGNSS. Each day of the 1-minute data serves as a realization of the prior distribution, and the week of 10-minute data is a realization used in the significance test, where the parameters are shown in (2.11a-d). To minimize the sensitivity to outliers, only values between the 5th and 95th percentiles of each dataset were utilized.

$$\mu_1 = \text{mean}(\Delta\sigma^o)_{1 \text{ min}} ; \mu_2 = \text{mean}(\Delta\sigma^o)_{10 \text{ min}} \quad (2.11a)$$

$$s_1 = \text{std}(\Delta\sigma^o)_{1 \text{ min}} ; s_2 = \text{std}(\Delta\sigma^o)_{10 \text{ min}} \quad (2.11b)$$

$$H_0: \mu_1 = \mu_2, H_1: \mu_1 \neq \mu_2 \quad (2.11c)$$

$$H_0: s_1 = s_2, H_1: s_1 \neq s_2; \quad (2.11d)$$

2.7 Summary of CYGNSS Performance Before and After Blackbody Sequence

Modification

Table 2-2. Performance of CYGNSS before and after blackbody cadence update.

Parameter	Confidence interval of 1-minute cadence	Value at 10-minute cadence	Significance (p-value)
$\text{mean}(\Delta\sigma^o)$	11.91 - 51.92	14.77	0.3866
$\text{std}(\Delta\sigma^o)$	53.56 - 61.59	54.87	0.1755

As shown in Table 2-2, there are no statistically significant differences in the overall performance of CYGNSS, both in terms of the average absolute performance of the difference between observations and expected NBRCS values, as well as any appreciable increase in variability. The errors inherent in windspeed model performance are several orders of magnitude greater than any expected deterioration that could be attributed to elongated blackbody sequence.

2.8 Summary and Conclusions

This work explores an optimal blackbody sampling cadence for the CYGNSS constellation. Fundamentally, the calculus to modify sampling cadence weighs the potential deterioration of data quality as a result of less precise thermal gain variation knowledge against the potential utility of increased science collection coverage.

With this analysis, we show that the original design generally overestimated the effects of thermal gain variations on the end derived products and the calibration sequencing oversampled at the expense of science operations.

The most dynamic thermal environment for CYGNSS spacecraft occurs as they cross the terminator, and the most significant variations occur when the orbit is at its lowest beta angle, when the spacecraft has the longest opportunity to cool in Earth's shadow. Under those conditions, any changes to the blackbody cadence will have the most significant impact on the derived products. However, even considering CYGNSS's most sensitive product, the YSLF windspeed, the blackbody sampling rate can be increased 10-fold to once every 10 minutes without degrading more than 5% of the data population by more than 1%.

With this modification, CYGNSS is able to improve its retrieval duty cycle from 90% to 98%, significantly improving the availability of data previously lost to excessive blackbody calibration.

Chapter 3 Exploring Representation Error

This chapter presents analysis exploring the nature of representativity error for situations in which two samples are nearby but not exactly collocated or simultaneous. This chapter is substantially derived from a work published in *Bulletin of the American Meteorological Society* under the title “Sampled Together: Assessing the Value of Simultaneous Collocated Measurements for Optimal Satellite Configurations” (Powell, Ruf, Gleason, et al. 2024). The research is presented as published, with minor updates to formatting.

© American Meteorological Society. Used with permission.

3.1 Abstract

This work applies a quantitative metric in order to capture the relative representativeness of non-simultaneous or non-co-located observations and quantify how these observations decorrelate in both space and time. This methodology allows for the effective determination of thresholding decisions for representative matchup conditions, and is especially useful for informing future network designs and architectures.

Future weather and climate satellite missions must consider a range of architectural trades to meet observing requirements. Frequently, fundamental decisions such as the number of observatories, the instruments manifested, and orbit parameters are determined based upon assumptions about the characteristic temporal and spatial scales of variability of the target observation. With the introduced methodology, representativity errors due to separations in space

and time can be quantified without prior knowledge of instrument performance, and errors driven by constellation design can be estimated without model ingest or analysis.

3.2 Introduction

It is often desirable to measure the Earth system from two or more different instruments at the same place and time. Simultaneous co-located measurements, otherwise known as matchups, frequently form the basis for calibration activities, science investigations, and operational retrievals. Satellite platforms offer a unique opportunity to capture co-located and simultaneous observations for extended periods of time. Many operational missions manifest multiple sensors onto the same satellite platform. Operational sea surface altimetry missions generally employ a radar altimeter as the primary mechanism to determine sea surface height, but due to uncertainties due to tropospheric delay, the radar is supplemented by a microwave radiometer to measure integrated atmospheric refractivity due to water vapor in the observing column to meet accuracy requirements (Donlon et al. 2021). Weather satellites utilize similar techniques to meet requirements. The NOAA-NASA Joint Polar Satellite System employs co-aligned microwave and infrared sounders to retrieve atmospheric profiles. To ensure operational consistency, the scanning mechanisms between the infrared and microwave instruments are synchronized such that they share the same field of regard across the scan (Kim et al. 2014).

The same logic also extends to formation flying, where instruments are not manifested on the same platform, but rather on multiple platforms in nearby, coordinated orbits. In the early 2000s, NASA populated its A-Train constellation satellites, named for its compact assemblage of several Earth science missions in the afternoon sun-synchronous polar orbit. When CloudSat and CALIPSO joined the A-Train in 2006, five separate satellites would fly in formation over the

same ground track within roughly a 15 minute window (Stephens et al. 2002; Schoeberl 2002). The quick succession of satellites and near-simultaneous observations were critical to several science goals of the constellation.

The decision to co-manifest instruments on a single satellite platform usually involves various trade studies to evaluate the relative costs, risks, and performance benefits of the design. While sharing two or more instruments on the same satellite platform is often the most intuitive way to achieve simultaneity and co-location, it can increase the system complexity, as well as the volume, mass, and power budgets of the spacecraft. These budgets are known to drive overall mission cost and execution risk. At the other end of the spectrum, recent advances in miniaturized sensors, small satellite platforms, and low-cost launch services have enabled constellations of proliferated sensors. These new capabilities enable constellation designs previously considered untenable or uneconomical. There are inherent challenges and risks with proliferated constellations, including cross calibration, formation maneuvers, and operating complexity.

The transformations in the space industry, including the development of new business models for collecting observations from space, combined with growing demand for enhanced weather and climate services, are fostering new conditions for government agencies to consider as they embark on the next generation of Earth observing architectures. NOAA is currently formulating a new architecture for low Earth orbit, the Near Earth Orbit Network (NEON), and is considering constellations that look very different from its legacy missions (Werner 2023). NEON Series One satellites will contain microwave and infrared sounders, while the manifest for Series Two is still undefined, and may contain instruments such as visible imagers or ozone profiles. The orbital planes and constellation size for NEON is still undetermined, but it is expected that the

disaggregation of instruments will enable more observations at similar or reduced costs. This program marks a shift from the current-generation of polar-orbiting weather satellites, the Joint Polar Satellite System, which hosts five instruments on a single truck-sized platform. NASA recently commissioned a study from the National Academies of Science, Engineering, and Medicine to assess the utility of hosting a number of Earth science payloads on a single large commercial platform (National Academies of Sciences, Engineering, and Medicine 2023). One of the advantages of using multiple smaller satellite platforms for Earth observations is that on-orbit resources are limited, and often the constellation risk posture can be improved by disaggregating and distributing sensors across multiple platforms. This advantage can be traded off against the detrimental effects of representativity error.

A number of tools have been developed to facilitate constellation design, usually with defined objective functions, such as minimizing revisit time, maximizing coverage, or balancing cost and utility (Marcuccio et al. 2019; Williams, Crossley, and Lang 2001; St. Germain, Gallagher, and Maier 2018; Nag et al. 2015). These considerations feature prominently in the development of Earth observing missions and frequently trade off with one another. However, often the most important questions for selecting an optimal architecture – the definition of threshold and objective revisit requirements for a given observation – are decided somewhat arbitrarily.

One of the primary metrics used to measure the utility of future observing systems, particularly for weather satellites, is a technique known as Observing System Simulation Experiments (OSSEs) (Arnold and Dey 1986). OSSEs essentially use a high-resolution model of the atmosphere and Earth system as ground truth, which are then “measured” by realistic simulated observations that are assimilated into a known forecasting model (Hoffman and Atlas

2016). These experiments can provide significant insight into the forecast impact of certain types of future observations but can be labor intensive and computationally expensive to run (to illustrate, cf. Li et al. 2019; 2018; Christophersen et al. 2021). Some newer techniques, such as employing non-cycling data assimilation for OSSEs, can minimize the computational burden (Privé et al. 2023). Other approaches which are less resource intensive, such as ensemble data assimilation (EDA) studies, measure a simulated observing system's impact on forecasts by characterizing the spread between ensemble members (Tan et al. 2007). Nevertheless, these studies are frequently individually tuned to the observing system being developed, and the utility is restricted to estimating improvements for forecasting users. Because of the number of variables inherent in these models, and the lack of standardization between techniques, it is difficult to make precise decisions about what utility is gained by changing the temporal or spatial coverage of future observation architectures. Further, these computational frameworks are not entirely applicable to non-forecasting applications of satellite observations, where a simple objective metric could guide requirements flow-down for observation representativity.

This work examines precisely how much representativity error is incurred when observations are separated in space and time, without *a priori* knowledge about the observing system or model at hand. As a result, it is a tool to empower observation planners with objective functions to make architecture trades. Further, when coupled with OSSEs, EDA studies, and other forecast skill metrics, it provides a more comprehensive assessment of a specific architecture selection.

3.3 Methodology

3.3.1 Formulation

The challenge of assimilating many different observations, which are often irregularly sampled in space and time, has been well documented since the advent of numerical weather prediction over a century ago (Richardson 1922). Gandin introduced a number of innovations in objective analysis and optimum interpolation of fields, which moved away from purely mathematical polynomial fittings and incorporated statistical arguments for how parameters of interest might decorrelate in space and time (Gandin 1965). These optimal interpolation methods were further generalized for data assimilation purposes by Rutherford (Rutherford 1972) who incorporated short-term forecast error information to produce a blend of observations and model rather than the forecast alone as a first guess. Implicit in this work is the need to inform the process with autocorrelation functions that represent how observations at a particular place and time relate to model-predicted values on a numerical grid. Bretherton et. al. (1976) extended the rapidly-advancing optimal interpolation assimilation methods used for initializing weather models to guide the development of an oceanographic field experiment. It is this context in which we adopt the general approach of Bretherton et al, not for an oceanographic experiment, but for satellite-based observations.

The basic premise is that given simple assumptions about the statistical behavior of any observable parameter, such as wind speed or sea surface temperature, one can measure the rate that this parameter decorrelates over space and time. The error from interpolation is then simply:

$$\epsilon_x(\tau_{t,s}) = \sigma_x \sqrt{1 - R_x(\tau_{t,s})} \quad (3.1)$$

where representativity error for parameter x is ϵ_x and is a function of either time lag τ_t or spatial lag τ_s , σ_x is the standard deviation of samples of x , and R_x is the autocorrelation of parameter x at lag $\tau_{t,s}$, where the subscript t represents the time lag, and the subscript s represents the spatial lag. This formulation makes general assumptions about the stationarity and isotropy of variance for the parameter of interest which are not necessarily true for the Earth system. There are many cases in which the statistical distribution and rate of decorrelation of weather and climate parameters are non-stationary, such as from seasonality or phase of teleconnections. Generally, the error in (3.1) can be calculated in any case where a representative decorrelation roll-off can be estimated. The correlation scale sizes and decorrelation behavior of various Earth parameters is extensively studied and readily accessible for many parameters. As a partial sampling, confer (Colosi and Barnett 1990) [identifying characteristic spatial and temporal scales in surface pressure, sea surface temperature, and air temperature over the Southern Hemisphere via drifting buoys]; (White 1995) [measuring decorrelation scales of temperature at various depths in global oceans to design an in-situ network to measure gyre-scale seasonal-to-interannual variability]; (Gille and Kelly 1996), [decorrelation scales of sea surface height of the Southern Ocean as measured from satellite altimeters]; (Kuragano and Kamachi 2000) [spatial and temporal scales of global ocean surface variability from the TOPEX/POSEIDON mission]; (Chu, Guihua, and Chen 2002) [decorrelation scales of temperature and salinity in the Japan Sea from in-situ ocean profiles]; (Delcroix et al. 2005) [temporal and spatial decorrelation in sea surface salinity in tropics]; (Romanou, Rossow, and Chou 2006) [decorrelation scales of latent and sensible heat fluxes in global oceans]; (Eden 2007) [eddy length scales of the North Atlantic derived from satellite altimetry]; (McLean 2010) [correlation scales of global oceans derived from Argo floats].

This analysis treats temporal and spatial decorrelation separately, even though they are coupled in the Earth system. The rationale is twofold. First, this maintains the simplicity and generality of the representativity errors. Second, this assumption makes this metric much more useful for the task of designing satellite architectures. Architecture trades frequently feature orbit and constellation decisions that optimize various sampling characteristics in both space and time, which are also coupled by orbital dynamics. Instead of trying to determine the dynamics of multiple coupled systems, this relaxation allows planners to set simple threshold and objective requirements in space and time for their trade studies.

3.3.2 Choice of Data and Study Selection.

The goal of this exercise is to demonstrate how we apply a known benchmark to determine satellite constellation architectures. In practice, care should be taken to ensure that the chosen benchmarks are, in fact, representative of the target phenomenology. In essence, (3.1) simply measures the error induced by the separation between samples in space and time in which natural variability may cause a different measurement to be obtained. The statistical meaning of (3.1) also suggests that any input data need to adequately resolve the target parameter both in scale and variability. The temporal and spatial resolutions should be considered as well as any processing performed that would otherwise impact the variability. There is no policy constraint on the source of data, provided that it is sufficiently representative of the dynamics at the appropriate scale and variability of a given phenomenon.

In certain cases, raw observation data may be advantageous over model or reanalysis data; in others, the reverse may be true. This claim warrants some explanation. We generally have a strong preference for using observation data when available and practical. But there are certain use cases in which model or reanalysis data may be more representative, easier to

manipulate, or simply readily available where observations are not. For instance, remote sensing data is commonly resampled from the native instrument resolution to a standard grid in both space and time, which can have a significant impact on the representativity of the data (Schutgens et al. 2016; Schutgens et al. 2017). Further, the statistical behavior of raw observation data may be sensitive to quality control parameters, which may not be intuitive to unfamiliar users. Some phenomena, such as global sub-mesoscale precipitation, is not well-captured by any global observing system, and a high-resolution model may be more appropriate than any blended observation dataset. Finally, in certain cases, the convenience of a gridded reanalysis may prove decisive if it can be shown to adequately capture the variability and scales of the target phenomena.

For this exercise, we assume that we are designing an ocean observing system, with particular interest in surface parameters such as winds, sea surface temperature, air temperature, and surface air pressure. Our selection of data is purely illustrative for the demonstration of this technique, and we assert representativity in both scale and variability. For observation planners, care should be taken to ensure that the source of “ground-truth” for target observables provide representative sampling at timescales and spatial resolution appropriate for the phenomena of interest and are otherwise well calibrated to ensure representative variability and dynamic range.

3.3.3 Temporal Decorrelation

For temporal decorrelation analyses, we chose a moored marine buoy from NOAA’s National Data Buoy Center (National Data Buoy Center 1971) due to its high temporal resolution, the quality and calibration of the sensors, and the continuity of observations from a single platform. For our analysis, we examined one year of observations from Station 51004, located at 17°32'17" N 152°13'48" W (approximately 205 nm southeast of Hilo, HI), which is a 3-meter foam buoy

with an updated SCOOP payload that reports meteorological parameters at 10-minute intervals (Kohler, LeBlanc, and Elliott 2015).

As discussed above, the choice of environmental parameter and source data may have a significant impact on the representativeness of decorrelation scales. For instance, parameters such as precipitation are likely to have much shorter characteristic decorrelation scale sizes than sea surface temperature. It is assumed that the temporal decorrelation scales from Station 51004 data are sufficiently representative to use as a benchmark for determining satellite constellation objectives in this hypothetical. We opted to use buoy data instead of reanalysis or other model outputs because model timesteps tend to be too coarse to capture the decorrelation behavior of our target parameters. Station 51004 was chosen due to its long duration of continuous data collection without equipment changes as well as its near-equatorial location which highlights highly variable surface weather patterns. We acknowledge that the temporal decorrelation of these parameters may vary by location, but assert that the hypothetical satellite architecture is driven by requirements for tropical latitudes. For observation planners, we suggest identifying ‘worst-case’ decorrelation behavior to drive requirements (i.e., stations that capture phenomena that have large variability and decorrelate quickly), which will vary depending on the target of interest.

The autocorrelation for observations was calculated assuming wide-sense stationarity for lags of 10 minutes. The decorrelation behavior and the resulting representativity error is shown in Fig. 3.1.




Figure 3.1 – (a) The temporal decorrelation behavior for a year of observations is demonstrated for various meteorological parameters from NDBC buoy station 51004. (b) The representativity error as calculated from (1) is shown. Note that the magnitude is dependent on the unit of measure.

Figure 3.1 is illustrative of several factors that observation planners should consider. The left panel shows the decorrelation behavior of several observed parameters as well as one derived parameter, the air-sea temperature difference, which is obtained by simple arithmetic subtraction of the near surface air temperature from the sea surface temperature observations. Some environmental parameters, such as sea surface temperature, are slowly varying on the timescale of hours to days. Other parameters, such as surface air pressure, exhibit strong diurnal behavior. The air-sea temperature difference decorrelates at a much faster rate than either air temperature or sea surface temperature alone, suggesting that these parameters are decoupled on timescales of less than a day. Figure 1b shows how these decorrelation behaviors factor into absolute

representativity error. For example, the representativity error of air temperature and of the air-sea temperature difference are nearly equivalent, despite the fact that the decorrelation behavior is vastly different. This is because the air-sea temperature difference has a much smaller dynamic range than the surface air temperature. At time lags of 5 hours, both exhibit approximately 0.4°C of representativity error, but that value is much more significant for the air-sea temperature difference, which has a mean of -0.76°C and a standard deviation of 0.5°C .

It is important to emphasize that this formulation addresses only the component of matchup error due to representativity. The overall matchup error between two parameters should also include measurement and retrieval errors associated with the individual parameters. The representativity error can therefore be considered the floor for matchup accuracy of a given observing system.

3.3.4 Spatial Decorrelation

A similar analysis can be applied in spatial dimensions. For a spatial dataset, we selected NASA's Modern-Era Retrospective analysis for Research and Applications version 2 (MERRA-2) hourly, non-averaged reanalysis (M2I1NXASM) (Gelaro et al. 2017; Global Modeling And Assimilation Office 2015). As before, we assume for this exercise that reanalysis data is sufficient to capture the spatial variability of our target parameters based on analysis from (Gille 2005), which suggest that at tropical latitudes, the variability of ocean surface wind stress from reanalysis is consistent with observations from scatterometers. Ocean surface windspeed has the fastest decorrelation roll-off in our hypothetical study, but as before, this assumption should be reconsidered for other phenomena that vary at scale sizes smaller than the resolvability of the model.

Spatial decorrelation of ocean surface observations is known to be anisotropic and location dependent. This study adopts a simple solution frequently employed by oceanographers, which is to individually evaluate the meridional and zonal components of the spatial decorrelation scales of an ocean observation (White 1995; Reynolds and Smith 1994). For consistency, we center the meridional and zonal transects at the location of Station 51004 as illustrated in Fig. 3.2. While these statistics are not stationary across the ocean, we assume for the purposes of this hypothetical that the statistics from Station 51004 are representative.




Figure 3.2 – The zonal and meridional transects for the Pacific basin centered at Station 51004 are highlighted. The zonal transect consists of 210 grid cells at 0.625 deg spacing. The meridional transect consists of 254 grid cells at 0.5 deg spacing. The background field is of an initialization of MERRA-2’s global windspeed output to emphasize that the spatial patterns of these parameters are quite different meridionally than they are zonally.

Figure 3.2 identifies which model grid cells are used to compute the zonal and meridional autocorrelation statistics. Because there are only 210 data points in the zonal transect and 254 data points in the meridional transect, the autocorrelation behavior derived from a single model run is somewhat noisy. To create a more representative decorrelation behavior, the autocorrelation is averaged across multiple model realizations over three months of data (at each hourly output for 92 days, or for 2,208 realizations).

Compared to the NDBC buoy data, the MERRA-2 hourly reanalysis produces slightly different observation parameters. For instance, the air temperature is the 10 meter temperature, whereas the buoy data is observed at 3 meters. Additionally we use the MERRA-2 skin temperature, which over most of the oceans is very similar to the sea surface temperature.

The average decorrelation behavior for the meridional and zonal transects centered at Station 51004, as well as the implied representativity error are shown in Fig. 3.3.




Figure 3.3 – The spatial decorrelation behavior and error statistics are partitioned into meridional and zonal components. (a) The meridional decorrelation behavior is shown for the specified environmental parameters. (b) The meridional representativity error is shown, note that the y-axis units are parameter-dependent. (c) The zonal decorrelation behavior is shown for the specified environmental parameters. (d) The zonal representativity error is shown, note that the y-axis units are parameter-dependent.

Figure 3.3 reveals much about the spatial decorrelation behavior of the selected parameters. First, the behavior is substantially different between zonal and meridional transects,

reflecting the spatial anisotropy in the decorrelation behavior in these parameters. Second, for the selected parameters, the spatial decorrelation rates suggest characteristic spatial scale sizes in the hundreds or even thousands of kilometers. Third, even when the decorrelation is comparable in zonal and meridional directions (i.e., the normalized autocorrelation is roughly equivalent for a given distance), the variability of the parameter may have significantly different magnitudes, changing the corresponding representativity error. For satellite missions, designers would want to base requirements decisions on the ‘worst-case’ scenario, which may include regimes, conditions, and locations where the spatial representativity error grows the fastest. In the case of our hypothetical, that suggests that the meridional error statistics would generally drive architecture decisions.

3.4 Discussion

This technique enables observation planners to quickly estimate the representativity error caused by separations in observation space and time, which has widespread utility in planning future satellite constellations. Given appropriate and statistically representative datasets of planned observation targets, planners can set quantitative observation objectives for optimizing satellite constellations.

3.4.1 Shortcomings

It is worth repeating that the representativity error indicated in this analysis is not inclusive of the dominant error sources of most observing systems, such as ambiguity in the retrieval or noise from the sensor. In cases where the additional representativity error incurred by having non-simultaneous and non- co-located measurements is small relative to the dominant error sources, it becomes reasonable to consider more flexible architectural approaches.

Further, there is danger in applying this technique with inappropriate input data that is not representative of the statistical behavior of the target parameter. Gridding observations can sometimes substantially impact the represented behavior. It can impose unwanted spectral filtering, changing the variability of the data, and can decrease the dynamic range of the data by smearing and averaging samples within gridded cells. Model and reanalysis data can be similarly non-representative of the target phenomena, with a whole host of other computational and mechanical artifacts that can change the variability and scaling of model parameters.

As a statistical approach, this metric necessarily averages between ‘calm’ low variability periods and infrequent but high-variability events, which may be insufficient to capture the variability of dynamic and consequential phenomena, e.g. tropical cyclones. One could, however, restrict the distribution of input data to use only ‘worst-case’ situations such as tropical cyclones as a design basis in order to estimate what spatial and temporal separation thresholds are needed to meet baseline representativity error requirements.

This technique also treats spatial and temporal dimensions separately, even though they are, in fact, coupled in the Earth system. This simplification is likely to impact conclusions drawn by a satellite observation planner in a conservative way. Decoupling spatial and temporal decorrelations will generally result in an overestimation of the error associated with spatial and temporal separations between samples. The decoupled results hence serve as a conservative upper bound on the error, and there may be opportunities to buy back more affordable satellite constellations with a coupled error estimate.

3.4.2 Utility and Future Work

We believe this metric should be considered by observation planners when considering architectural trade studies for new satellite constellations, and early in the mission lifecycle when

requirements are set. This simple metric can price when temporal resolution is too sparse for the assumption of simultaneity or when spatial coverage is too coarse for the assumption of collocation.

From our example, consider the design of a constellation to observe ocean surface windspeed from two separate observatories with the same orbit ground track. Given a requirement of representativity error no greater than 0.5 m/s, the above results from Figure 1 indicate that the satellites be staggered no more than 25 minutes apart in the orbit plane. Note that this corresponds to an average representativity error under typical conditions. However, if the target observable is ocean winds for tropical cyclone monitoring, during which surface windspeed exhibits a much larger dynamic range and steeper decorrelation roll-off, a 25 minute separation would likely be too large. In this case, observation planners can quickly estimate that more satellites would be needed, and estimate cost of meeting this requirement. From another angle, planners with fixed budgets could mitigate this by placing the two satellites closer together within the same plane, reducing the revisit time between the two, but at the expense of daily coverage.

This problem can be extended to thresholds established for matching up observations for opportunistic or vicarious calibrations. Techniques such as simultaneous nadir overpasses (SNOs) (Zou et al. 2006) often establish thresholds in space and time that approximate simultaneity. This technique can establish objective representativity guidelines which can optimize the quantity of matchup data, or enable constellations that feature frequent SNOs for operational calibration.

3.5 Data Availability Statement.

Data analyzed in this study were derived from existing public, openly available datasets cited in the reference section. See (National Data Buoy Center 1971; Global Modeling And Assimilation Office 2015).

Chapter 4 An Instrument Error Correlation Model for GNSS-R

This chapter presents a novel first-principles model for instrument error correlation for GNSS-R. This chapter is substantially derived from a work published in *Remote Sensing* under the title “An Instrument Error Correlation Model for Global Navigation Satellite System Reflectometry” (Powell, Ruf, McKague, et al. 2024). The research is presented as published, with deletions of introductory material and theory (covered in Chapter 1), integration of the published appendices into the main body, and minor updates to formatting and notation for consistency.

4.1 Abstract

All sensing systems have some inherent error. Often, these errors are systematic, and observations taken within a similar region of space and time can have correlated error structure. However, the data from these systems are frequently assumed to have completely independent and uncorrelated error. This work introduces a correlated error model for GNSS reflectometry (GNSS-R) using observations from NASA’s Cyclone Global Navigation Satellite System (CYGNSS). We validate our model against near-simultaneous observations between two CYGNSS satellites and double-difference our results with modeled observables to extract the correlated error structure due to the observing system itself. Our results are useful to catalog for future GNSS-R missions and can be applied to construct an error covariance matrix for weather data assimilation.

4.2 Introduction (Abridged)

The assumption that observation errors are independent and uncorrelated is not valid for any observatory. In practice, observation error correlations are minimized by thinning the dataset or ‘super-obbed’ samples, which improves overall model skill by mitigating the effects of correlated error, which may be interpreted by the model as a real signal (Bauer et al. 2011; Gao et al. 2019; Hoffman 2018).

The premise of not fully exploiting observation data is unappealing, considering that the process of acquiring these data usually requires the considerable capital expense of building and flying remote sensing satellites. The highest utility of dense observations in space and time occurs when the application (i.e., the model) resolution matches the observation resolution. For many imagers and sounders, there may be a reasonable case that thinning is an appropriate way to scale dense observation data to match model gridding.

For CYGNSS, however, thinning the data is an especially unpalatable solution, as the peak utility of CYGNSS data is during the relatively uncommon occurrence that a specular point passes through the eyewall of a tropical cyclone. Thinning may miss this observation or misrepresent the structure of the storm in the model. By a similar token, ‘super-obbing’ is not a practicable remedy because tracks are one dimensional and collocated observations are potentially hours apart. Blending data from a wide temporal window risks misrepresenting the dynamics of the tropical cyclone and negates the fast-revisit utility of the observatory.

Estimating the correlated error structure for observations can improve model skills in NWP (Stewart, Dance, and Nichols 2013; Rainwater, Bishop, and Campbell 2015; Stewart, Dance, and Nichols 2008; Simonin et al. 2019; Daley 1992). A number of methods have been demonstrated

to estimate the observation error correlation matrix (Waller, Dance, and Nichols 2016; Dee and Da Silva 1999), most notably Desroziers' diagnostic (Desroziers et al. 2005).

Our work does not estimate the full observation error correlation matrix that is typical of these prior works. Instead, we provide a first-principle, bottom-up, tunable engineering model for how the CYGNSS instrumentation itself can cause observations to contain correlated errors. This correlation model is validated and tuned against empirical observation data during a period when two CYGNSS assets were in a specific orbital geometry where observations nearly coincided in space and time. We further discuss the challenges and opportunities that result from this correlated error model, as well as highlight the potential for applicability to future assimilation investigations.

4.3 Materials and Methods (Abridged)

Each of the terms in Equation (1.2) can potentially be a source of correlated error, but for practical reasons, this study only evaluates those sources with the largest error magnitude and, therefore, the largest potential utility for future data assimilation users. We omitted consideration of the smallest error sources that, when combined as a root sum of squares, contribute to less than 1% of the total error magnitude. The five largest sources of error are explored, as shown in Table 4.1, representing a 1-sigma error magnitude for a reference wind speed of 10 m/s. The following sections evaluate each of these terms in depth.

Table 4-1. The magnitudes of 1-sigma errors for each term in Equation (1.2). The shaded rows indicate that the absolute magnitudes of these terms are negligible compared to the dominant error terms and are neglected in the construction of the error model in this work.

Error Term	Error Magnitude [dB]
$E(G^R)$	0.43 (Ruf et al. 2022)
$E(P^g)$	0.23 (Powell, Ruf, and Russel 2022)
$E(G^Z)$	0.20 (Wang, Ruf, Gleason, et al. 2021)
$E(P^Z)$	0.18 (Wang, Ruf, Gleason, et al. 2021)
$E(\zeta^E)$	0.15 (Wang, Ruf, Gleason, et al. 2021)
$E(\bar{A})$	0.05 (Gleason 2020)
$E(L^{atm})$	0.04 (Gleason 2020)
$E(R^{tot})$	<0.01 (Gleason 2020)
$E(L^Z)$	<0.01 (assumed $\sim E(R^{tot})$)

4.3.1 The Bottom-up Correlated Error Model

The CYGNSS Level 1 correlated error model is constructed to represent the major error sources based on the physical and operational characteristics of the CYGNSS observatories. These errors are correlated over both space and time, and depend on several variables relating to the operation of the CYGNSS constellation.

The construction of the error model is intended to be realistic enough to represent measurable and plausible correlated errors, yet flexible enough to allow for tuning and parameterization. The general modeled structure of correlated error takes the form:

$$R_{mod}(i, j) = \frac{1}{\mathcal{N}} \sum_n K_n(i, j) \quad (4.1)$$

where $R_{mod}(i, j)$ is the normalized correlated error between any two samples i and j with a domain $-1 \leq R_{mod} \leq 1$. There are n component terms that are added together, and each of the n th term corresponds to a unique source of error. $K_n(i, j)$ represents the error covariance between any two samples i and j for a specific component n . \mathcal{N} is a normalization constant to ensure that the diagonal terms of the R_{mod} matrix is 1.

Each of the K_n terms can be further generalized as follows:

$$K_n(i, j) = (E(n))^2 \cdot Ecorr_n(i, j) \quad (4.2)$$

where $E(n)$ represents the magnitude of each error component n as calculated in (Gleason 2018; 2020; Wang, Ruf, Gleason, et al. 2021) and can be thought of as the variance of the error, and $Ecorr_n(i, j)$ is a derived function with a domain $-1 \leq Ecorr \leq 1$ that represents the correlation in error between samples i and j for each component n .

Each $Ecorr$ function depends on different arguments depending on the nature of the error source. The individual errors are explored in depth in the following sections. Further, several tuning parameters have been added to assist with validation. The initial values of the tuning parameters are 1 and essentially leave the model output unmodified. However, this process assumes that the relative magnitude of the error components could be incorrectly specified, and by varying the tuning parameters, the model can be shaped to match a validation dataset.

4.3.2 Model Assumptions

Several assumptions are made regarding the overall model and each individual $Ecorr$ term. The first is that the decorrelation scales of interest are on the order of seconds to minutes. This restricts the problem to errors that decorrelate on timescales that would be of relevance to numerical weather prediction users. CYGNSS observations may have a correlated error structure that evolves during longer timescales, say, daily or seasonally. However, this structure would probably be better resolved using other calibration and processing techniques and would not add significant value to a correlated error matrix. As a result, all errors between CYGNSS samples i and j are assumed to be 100% uncorrelated if there is more than 10 min of separation in observation time. This corresponds to roughly 4500 km of distance, where observations could reasonably be treated as wholly uncorrelated.

Further, all error terms $Ecorr_n$ are assumed to be uncorrelated with each other. This assumption simplifies the implementation and calculation of the model but is less supported theoretically. There are a number of plausible rationales for why many of these error sources are correlated. However, this model hopes to correct for these empirically with the implementation of tuning parameters.

Finally, there is a general assumption that all the error sources in the model exhibit wide-sense stationarity during the timescales of interest. Therefore, the error magnitude values $E(n)$ are treated as constants. In practice, the magnitude of errors and their correlation time and space scale sizes may well vary (e.g., due to seasonal dependence). A fully operational implementation of the methodology developed here would take these dependencies into account, e.g., by an appropriate parameterization of the properties of the errors. Here, we only consider stationary error properties in order to illustrate how they are determined and how they would be combined to be used by a DA scheme.

4.4 Covariance Matrix Construction

The construction of R_{mod} makes a few important assumptions that are worth examining in depth. First, we will describe the general mathematical framework. In general, if X and Y are random vectors of length N , then the covariance matrices of X and Y are constructed:

$$K_{XX} = \langle X - \langle X \rangle \rangle \langle X - \langle X \rangle \rangle^T \quad (4.3a)$$

$$K_{YY} = \langle Y - \langle Y \rangle \rangle \langle Y - \langle Y \rangle \rangle^T \quad (4.3b)$$

where K_{XX} and K_{YY} are the N -by- N covariance matrices of X and Y , respectively, the bracket operator denotes the expectation operation, and the superscript T denotes the transpose vector.

Notably, the covariance matrix is symmetric and positive semi-definite. The covariance matrices can also be constructed in the following manner:

$$K_{XX} = S_X R_{XX} S_X \quad (4.4)$$

where S_X is a diagonal N -by- N matrix with the standard deviations of X_i :

$$S_X = \begin{bmatrix} \sigma_{X1} & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \sigma_{XN} \end{bmatrix} \quad (4.5)$$

and R_{XX} is the correlation matrix with correlation coefficients $\rho_{i,j}$:

$$R_{XX} = \begin{bmatrix} 1 & \rho_{X1,X2} & \cdots & \rho_{X1,XN} \\ \rho_{X2,X1} & 1 & & \\ \vdots & & \ddots & \\ \rho_{XN,X1} & & & 1 \end{bmatrix} \quad (4.6)$$

As standard, $-1 \leq \rho_{i,j} \leq 1$, and the diagonals all must equal 1. The *cross-covariance* matrices can also be defined in a similar fashion:

$$K_{XY} = \langle X - \langle X \rangle \rangle \langle Y - \langle Y \rangle \rangle^T \quad (4.7a)$$

$$K_{YX} = \langle Y - \langle Y \rangle \rangle \langle X - \langle X \rangle \rangle^T \quad (4.7b)$$

where K_{XY} and K_{YX} are the cross-covariance matrices and are not generally identical. The error magnitudes in Table 4-1 are calculated using a standard error propagation technique as highlighted in Gleason et al. (2019):

$$E_n = \left| \frac{\partial F}{\partial n} \right| \Delta n \quad (4.8)$$

E_n is the error magnitude of term n , Δn is the estimated 1-sigma dynamic range of n and F is an arbitrary function. We can propagate these errors as the sums of covariance matrices with the derivation below, inspired by Chapter 9 in Taylor (Taylor 1997).

For an arbitrary function F with random vector arguments X and Y , $F(X, Y)$, where $X = (X_1, X_2, \dots, X_N)$ and $Y = (Y_1, Y_2, \dots, Y_N)$, we can approximate F by its first order Taylor Series expansion about the mean value, assuming that the errors are generally small compared to the arguments:

$$F_i = F(X_i, Y_i) \quad (4.9)$$

$$F_i \approx F(\langle X \rangle, \langle Y \rangle) + \frac{\partial F}{\partial X}(X_i - \langle X \rangle) + \frac{\partial F}{\partial Y}(Y_i - \langle Y \rangle) \quad (4.10)$$

Noting that $\langle F \rangle = F(\langle X \rangle, \langle Y \rangle)$, and using $F \approx F(\langle X \rangle, \langle Y \rangle) + \frac{\partial F}{\partial X}(X - \langle X \rangle) + \frac{\partial F}{\partial Y}(Y - \langle Y \rangle)$, then,

$$K_{FF} = \langle F - \langle F \rangle \rangle \langle F - \langle F \rangle \rangle^T \quad (4.11)$$

and expanding using Equations (4.7), (4.9), and (4.10),

$$\begin{aligned} K_{FF} &= \left\langle \frac{\partial F}{\partial X}(X - \langle X \rangle) + \frac{\partial F}{\partial Y}(Y - \langle Y \rangle) \right\rangle \left\langle \frac{\partial F}{\partial X}(X - \langle X \rangle) + \frac{\partial F}{\partial Y}(Y - \langle Y \rangle) \right\rangle^T \\ &= \left[\left\langle \frac{\partial F}{\partial X}(X - \langle X \rangle) \right\rangle + \left\langle \frac{\partial F}{\partial Y}(Y - \langle Y \rangle) \right\rangle \right] \left[\left\langle \frac{\partial F}{\partial X}(X - \langle X \rangle) \right\rangle + \left\langle \frac{\partial F}{\partial Y}(Y - \langle Y \rangle) \right\rangle \right]^T \\ &= \left[\left\langle \frac{\partial F}{\partial X}(X - \langle X \rangle) \right\rangle \right] \left[\left\langle \frac{\partial F}{\partial X}(X - \langle X \rangle) \right\rangle \right]^T + \left[\left\langle \frac{\partial F}{\partial Y}(Y - \langle Y \rangle) \right\rangle \right] \left[\left\langle \frac{\partial F}{\partial Y}(Y - \langle Y \rangle) \right\rangle \right]^T \\ &\quad + \left[\left\langle \frac{\partial F}{\partial X}(X - \langle X \rangle) \right\rangle \right] \left[\left\langle \frac{\partial F}{\partial Y}(Y - \langle Y \rangle) \right\rangle \right]^T + \left[\left\langle \frac{\partial F}{\partial Y}(Y - \langle Y \rangle) \right\rangle \right] \left[\left\langle \frac{\partial F}{\partial X}(X - \langle X \rangle) \right\rangle \right]^T \\ &= \left(\frac{\partial F}{\partial X} \right)^2 K_{XX} + \left(\frac{\partial F}{\partial Y} \right)^2 K_{YY} + \left(\frac{\partial F}{\partial X} \right) \left(\frac{\partial F}{\partial Y} \right) K_{XY} + \left(\frac{\partial F}{\partial X} \right) \left(\frac{\partial F}{\partial Y} \right) K_{YX} \end{aligned} \quad (4.12)$$

Substituting $E_n^2 = \left(\frac{\partial F}{\partial n} \right)^2 \sigma_n^2$ from Equation (4.8), and noting the decomposition in Equation (4.4),

$$K_{FF} = E_x^2 R_{XX} + E_y^2 R_{YY} + E_x E_y R_{XY} + E_x E_y R_{YX} \quad (4.13)$$

The construction of R_{mod} ignores the cross-correlation between component terms, i.e., we assert

$R_{XY} = R_{YX} = 0$. This is for two main reasons:

1. The terms K_n , as described in the following sections, are not constructed from random variables but rather through analytic specification to emulate the expected correlated behavior. We generally have insufficient knowledge to measure or estimate the cross-correlation *between* error components. Instead, this model simply estimates

the cross-correlation of error *within* individual components, which are then added independently.

2. Any residual cross-correlation between error components can be tuned per our tuning parameters.

This formulation assumes that each component term of the error E_n comes from a wide-sense stationary distribution where the error magnitudes are treated as constants. This means that for each component of K_{FF} , and noting Equation (4.4),

$$\left(\frac{\partial F}{\partial X}\right)^2 K_{XX} = \left(\frac{\partial F}{\partial X}\right)^2 S_X R_{XX} S_X = \left(\frac{\partial F}{\partial X}\right)^2 \sigma_X^2 R_{XX} = E_X^2 R_{XX} \quad (4.14)$$

For our error correlation model R_{mod} , we simply construct R_{XX} analytically for each term, for which we adopt the nomenclature $Ecrr$ to emphasize that it is an error correlation function. Therefore, from Equation (4.1),

$$K_n(i, j) = E_n^2 \cdot Ecrr_n(i, j) \quad (4.15)$$

For this model, F is drawn from Equation (1.2) (reproduced here):

$$\sigma^o = \frac{P^g 4\pi L^{atm} L^Z G^Z}{P^Z \zeta^E G^R L^R \bar{A}}$$

where we use the parameters with the largest error magnitude, $\sigma^o = F(P^g, P^Z, G^Z, G^R, \zeta^E)$. The

relevant partial derivatives $\frac{\partial F}{\partial N}$ from Equation (4.12) become the following:

$$\left|\frac{\partial F}{\partial P^g}\right| = \frac{4\pi L^{atm} L^Z G^Z}{P^Z \zeta^E G^R L^R \bar{A}} \quad (4.16a)$$

$$\left|\frac{\partial F}{\partial P^Z}\right| = \frac{P^g 4\pi L^{atm} L^Z G^Z}{(P^Z)^2 \zeta^E G^R L^R \bar{A}} \quad (4.16b)$$

$$\left|\frac{\partial F}{\partial G^Z}\right| = \frac{P^g 4\pi L^{atm} L^Z}{P^Z \zeta^E G^R L^R \bar{A}} \quad (4.16c)$$

$$\left| \frac{\partial F}{\partial G^R} \right| = \frac{P^g 4\pi L^{atm} L^Z G^Z}{P^Z \zeta^E (G^R)^2 L^R \bar{A}} \quad (4.16d)$$

$$\left| \frac{\partial F}{\partial \zeta^E} \right| = \frac{P^g 4\pi L^{atm} L^Z G^Z}{P^Z (\zeta^E)^2 G^R L^R \bar{A}} \quad (4.16e)$$

To estimate E_n , as shown in Table 4-1, these partial derivatives are evaluated at the 1-sigma value for a reference 10 m/s wind speed using Equation (4.8). The total error model becomes the following:

$$R_{mod}(i, j) = \frac{1}{\mathcal{N}} [K_{p^g}(i, j) + K_P^Z(i, j) + K_{G^Z}(i, j) + K_{G^R}(i, j) + K_{\zeta^E}(i, j)] \quad (4.17)$$

where K_{p^g} is the correlated error component from the calibrated nadir receiver power and is discussed described in Section 4.4.1, K_P^Z is the correlated error component from the zenith receiver and is described in Section 4.4.2, K_{G^Z} and K_{G^R} are the correlated error components from the nadir and zenith CYGNSS antenna patterns, respectively, and are discussed in Section 4.4.3, and K_{ζ^E} is the correlated error component due to the zenith–specular ratio used for dynamic EIRP estimation and is discussed in Section 4.4.4. \mathcal{N} is a normalization constant that forces R_{mod} to behave like a correlation such that $-1 \leq R_{mod} \leq 1$ and can be thought of as an estimate for the rolled-up variance. Because this model has tunable parameters, it is not necessarily representative of the true variance of σ^o but rather of the modeled variance from our bottom-up model construction.

4.4.1 Calibrated Nadir Power P^g Error Model

The term P^g represents calibrated received power from GPS signals reflected from the Earth’s surface. P^g is calibrated both from pre-launch characterizations as well as on an on-orbit blackbody at a known temperature, which, as of 2022, takes a reading every 10 min to re-

compute gain that may have changed due to the dynamic thermal environment in orbit (Powell, Ruf, and Russel 2022).

For CYGNSS, P^g is computed by the Level 1A algorithm in Equation (1.4), and instrument gain is measured in orbit by performing readings from an onboard blackbody at a known temperature and calibrated via Equation (2.1). Equations (1.4) and (2.1) can be combined:

$$P^g = \frac{(C - C_N)}{C_B} (P_B + P_r) \quad (4.18)$$

The magnitude of the errors for the components calculating P^g is displayed in Table 4-2 and has a similar calculation as before.

Table 4-2. The magnitudes of 1-sigma errors for each term in Equation B3. The shaded rows indicate that the absolute magnitude is negligible compared to the dominant error terms and is neglected in the construction of the error model in this work.

Error Term	Error Magnitude [dB]
$E(C_N)$	0.14 (Gleason 2018)
$E(P_r)$	0.14 (Gleason 2018)
$E(C)$	0.10 (Gleason 2018)
$E(C_B)$	0.07 (Powell, Ruf, and Russel 2022)
$E(P_B)$	~0.04 (Gleason 2018)

C_B , P_B , and P_r all vary with temperature, and the instrument gain G will fluctuate as the satellite enters different thermal conditions in orbit. The most significant errors will occur just as the satellite crosses the terminator. At that point, CYGNSS will go from a nearly steady-state thermal environment, such as approximately half an orbit of illumination or eclipse, and then quickly enter the opposite state. The fraction of orbit spent illuminated is determined by the orbit beta angle, which varies on scales of weeks to months.

The dominant error term in the Level 1A algorithm is C_N , which also varies with temperature. The calibration sequence is designed to correct for this, and we assume that the errors vary slowly with the timescale of interest, which is defined to be on orders of seconds to

minutes. Therefore, all errors from C_N are assumed to be 100% correlated in time within a given track of CYGNSS observations, i.e., when a series of samples adjacent in space and time share a GPS transmitter and a CYGNSS receiver. C_N is also very sensitive to radio-frequency interference (RFI), which will present as non-physical signals above the specular point in a delay–Doppler map. We do not aim to model the complex phenomenologies of RFI in this work and assume there are no correlated error structures from RFI.

Errors in P_r occur because of a variety of reasons. The low noise amplifiers were all characterized on the ground prior to launch to establish the relationship of the noise figure with respect to temperature. The values of this relationship were stored in a look-up table (LUT) for processing science data. However, as the amplifiers age, the noise floor characteristics may have evolved, producing errors in this mapping. Further, the thermal environment of the thermocouple may not be exactly the same as experienced by the amplifier itself. For the purposes of this model, we assume all errors due to incomplete or erroneous knowledge of the true receiver noise power are 100% correlated with each other for a given track, as we assume that the errors evolve slowly compared to the timescales of interest.

Therefore, the error correlation terms for C_N and P_r for any arbitrary CYGNSS samples x_i and x_j are the following:

$$E_{corr_{C_N}}(x_i, x_j) = E_{corr_{P_r}}(x_i, x_j) = \begin{cases} 1, & \text{conditions met} \\ 0, & \text{else} \end{cases} \quad (4.19)$$

The conditions that must be met are the following: x_i and x_j must share a CYGNSS receiver and a GPS transmitter and be observed within 10 min of each other.

C is the measured parameter, the raw counts of power from a science observation near the region of the specular point. The analog-to-digital processing chain is the primary source of errors, such as quantization errors and non-common-mode interference. We assume that these

error terms are 100% uncorrelated with each other; that is, for every sample, it can be treated as white noise. The error correlation term for C is then the following:

$$E_{corr_C}(x_i, x_j) = \begin{cases} 1, & x_i = x_j \\ 0, & x_i \neq x_j \end{cases} \quad (4.20)$$

The error term for counts measured during a blackbody sample C_B is a source of analytically defined correlated error. Every 10 min (earlier in the mission, every 1 min), the receiver is switched from the nadir science antenna to look at the onboard blackbody source for a period of 4–6 s. Science observation processing linearly interpolates the counts between the nearest blackbody looks. When errors are made in estimating C_B , those errors are correlated linearly with all adjacent samples due to this interpolation. Correlation due to linear interpolation has an analytical form. Assuming a blackbody look happens at timesteps 0 and n , then, the correlation between any two samples x_i and x_j at arbitrary timesteps i and j where $0 \leq i < j \leq n$ is the following:

$$E_{corr_{C_B}}(x_i, x_j) = \frac{1}{\sigma_i \sigma_j} \left(\frac{(n-i)(n-j)}{n^2} + \frac{ij}{n^2} \right) \quad (4.21)$$

where $\sigma_i = \sqrt{\frac{(n-i)^2 + i^2}{n^2}}$ and $\sigma_j = \sqrt{\frac{(n-j)^2 + j^2}{n^2}}$. The actual values of the sampled blackbodies do not matter, as the correlated error is simply a function of how far the samples are from the blackbody looks in time.

Errors in P_B are due to misestimations of the blackbody's true noise power, which may be because the thermocouple is measuring incorrectly. We assume that the errors of this nature not only are slowly varying compared to the timescales of interest but, because of the marginal absolute magnitude, factor a negligible and unmeasurable amount in the overall correlated error structure. As such, the model ignores this term.

The rolled-up correlated error model from the sources in P^g between any arbitrary samples i and j can be expressed as follows:

$$K_{P^g}(i, j) = \begin{bmatrix} \alpha \cdot E(C)^2 \cdot Ecorr_C(i, j) + \\ \beta \cdot E(C_N)^2 \cdot Ecorr_{C_N}(i, j) + \\ \beta \cdot E(P_r)^2 \cdot Ecorr_{P_r}(i, j) + \\ E(C_B)^2 \cdot Ecorr_{C_B}(i, j) \end{bmatrix} \quad (4.22)$$

Because this model estimates the correlated error in each term that calculates P^g , this contribution to the overall correlated error is not multiplied by the rolled-up error magnitude $E(P^g)$. As such, we do not normalize this construction, as it will be normalized when combined with the other constituent terms in R_{mod} . R_{P^g} contains two tuning parameters: α is used to size uncorrelated white-noise error, and β is used to size the magnitude of totally correlated errors.

4.4.2 Zenith Power P^Z Error Model

The zenith receiver on CYGNSS works much the same way as the nadir science receiver. However, the zenith receiver does not have an onboard calibration system, and data are processed from the pre-flight characterizations of the electronics. The counts of the receiver are converted to watts via a quadratic regression. Because the satellite is subject to the same thermal dynamics as the nadir receiver, one can expect that the zenith power estimate contains errors due to thermally driven gain variations. We model the $Ecorr_{P^Z}$ similarly to the nadir receiver but with some important distinctions. Because the zenith receiver operates without an onboard blackbody to calibrate against, a number of simplifying assumptions are made. Errors are broken up into just two components: correlated $Ecorr_{P^Z_1}$ and uncorrelated $Ecorr_{P^Z_2}$. We further assume that correlated error due to the absence of an onboard blackbody are outside the timescale of

interest. The correlated error term $Ecorr_{P_1^Z}$ is assumed to be totally correlated, provided that the matching conditions are met:

$$Ecorr_{P_1^Z}(x_i, x_j) = \begin{cases} 1, & \text{conditions met} \\ 0, & \text{else} \end{cases} \quad (4.23)$$

The conditions for $Ecorr_{P_1^Z}$ are that samples x_i and x_j are made with the same CYGNSS observatory and within 10 min of each other. In addition, an uncorrelated component is allowed:

$$Ecorr_{P_2^Z}(x_i, x_j) = \begin{cases} 1, & x_i = x_j \\ 0, & x_i \neq x_j \end{cases} \quad (4.24)$$

The rolled-up correlated error model from the sources in P^Z between any arbitrary samples i and j can be expressed as follows:

$$K_{P^Z}(i, j) = \left[\beta \cdot E(P_1^Z)^2 \cdot Ecorr_{P_1^Z}(i, j) \right] + \left[\alpha \cdot E(P_2^Z)^2 \cdot Ecorr_{P_2^Z}(i, j) \right] \quad (4.25)$$

with the same tuning parameters α and β , as in Appendix A. Note that we assume $E(P_1^Z) = 0.18$ dB, as suggested in (Wang, Ruf, Gleason, et al. 2021). $E(P_2^Z)$ is estimated directly from a 24 h of CYGNSS data as approximately 1% of the magnitude of the signal P^Z ; therefore, this model assumes $E(P_2^Z) = 0.04$ dB.

4.4.3 Error Model for Antenna Gain Patterns G^R and G^Z

The CYGNSS observatory has three antennas: one zenith antenna that is used for direct GPS-to-CYGNSS signal tracking, as well as two nadir science antennas that are used to capture the scattered signal from Earth's surface. Each of the eight spacecraft had all three antennas characterized pre-launch, and values were stored in a look-up table for science processing. Errors in the antenna gain pattern can arise for a variety of reasons. First, the measurement equipment on the ground is essentially a receiver but in controlled conditions. This means that while systematic and correlated errors are likely well-constrained, uncorrelated speckle-type

error can still occur. To produce realistic antenna gain patterns, the results of the ground characterization were smoothed with various filters and techniques.

Another source of error is the fact that CYGNSS antennas were not characterized while integrated with the spacecraft. This was a cost-saving measure decided by the mission management team. However, the electromagnetic properties of the antenna couple in some fashion with the spacecraft bus, and that will inevitably change the gain patterns.

Initial analysis of CYGNSS data shortly after launch showed significant retrieval performance dependence on the observation azimuthal angle with respect to the CYGNSS body frame, which was later hypothesized to originate in errors in the CYGNSS antenna patterns. To compensate for this deviation from measured patterns, the CYGNSS antenna patterns have been updated at several instances over the mission life via empirical calibration. The nadir antenna patterns are updated by comparing a climatology of CYGNSS measurements of σ^o (> 2 years) with model-generate σ_{mod}^o and plotting a scaling factor in the antenna reference coordinate system. σ_{mod}^o is generated by using modeled reanalysis winds to generate mean-squared slope with the L-band spectrum extension model, as described in (Wang et al. 2019). However, during the generation of these updated patterns, a number of smoothing filters are applied.

This prompts a discussion of a conjecture used extensively for this section:

Conjecture 4.1. *Uncorrelated errors can become correlated by post-processing with averaging and filters.*

This insight drives much of this section's analysis. Smoothing and filtering will necessarily impose a correlation in error between previously uncorrelated errors. For white noise, that implies that the choice of filter will add color and structure to the noise.

In particular, a handy lemma allows us to demonstrate that for white noise, the information required to capture correlated error structure is the filtering kernel itself. We will explore this behavior for a one-dimensional case, but it is generalizable to higher dimensions in our application, as the two-dimensional filters used for antenna smoothing are separable by construction.

Lemma 4.2. *For a filtered signal $F(t) = K(t) * D(t)$, where $K(t)$ is a filtering kernel and $D(t)$ is an arbitrary data signal, the autocorrelation $F \star F$ is the convolution of the autocorrelated kernel $K \star K$ and the autocorrelated signal $D \star D$.*

Proof of Lemma 4.2. Assume convolution and cross-correlation have the standard definitions for two real-valued timeseries $K(t)$ and $D(t)$, that is,

$$K(t) * D(t)[n] = \sum_{t=-\infty}^{\infty} K(t)D(n-t) = \sum_{t=-\infty}^{\infty} K(n-t)D(t) \quad (4.26)$$

and

$$K(t) \star D(t)[n] = \sum_{t=-\infty}^{\infty} K(t)D(n+t) = \sum_{t=-\infty}^{\infty} K(t-n)D(t) \quad (4.27)$$

where $*$ is the convolution operator, \star is the cross-correlation operator, and n is the lag argument. Observe that convolution operations are commutative and, further, that the cross-correlation can be written as a convolution by exploiting its symmetry:

$$K(t) \star D(t)[n] = K(-t) * D(t)[n] \quad (4.28)$$

Therefore, to evaluate the correlated error imposed by kernel $K(t)$,

$$\begin{aligned} F(t) \star F(t) &= \{K(t) * D(t)\} \star \{K(t) * D(t)\} \\ &= K(-t) * D(-t) * K(t) * D(t) \\ &= \{K(-t) * K(t)\} * \{D(-t) * D(t)\} \end{aligned}$$

$$= \{K(t) \star K(t)\} * \{D(t) \star D(t)\}$$

If the arbitrary signal $D(t)$ happens to be white noise, it is completely uncorrelated, and its autocorrelation collapses to a Dirac delta function centered at $t = 0$. Therefore, the entire structure of the correlated error is from the filter itself:

$$F(t) \star F(t) = K(t) \star K(t) \tag{4.29}$$

□

For the purposes of this work’s error model, we have no knowledge of the potential correlated structure in the actual errors in the gain pattern. The nadir antenna patterns are updated after applying a 6-degree boxcar averaging filter in both the azimuthal and elevation in the spacecraft coordinate frame and then an additional 10-degree two-dimensional smoothing window. We assume that zenith antenna patterns use a similar post-processing technique during their generation.

These filtering kernels act like low-pass filters. All correlated structure on scales ~ 5 degrees and smaller and uncorrelated error will be strongly influenced by the filtering process, and the correlated structure can be estimated from the Filter Lemma. This model assumes that there is no residual larger-scale structure in correlated error in the antenna gain patterns.

The generated filter kernel, which applies to both the nadir and zenith antenna patterns, can be shown in Figure 4.1. The correlated error is a function of how close any two observations are with respect to the relevant antenna gain pattern coordinates.

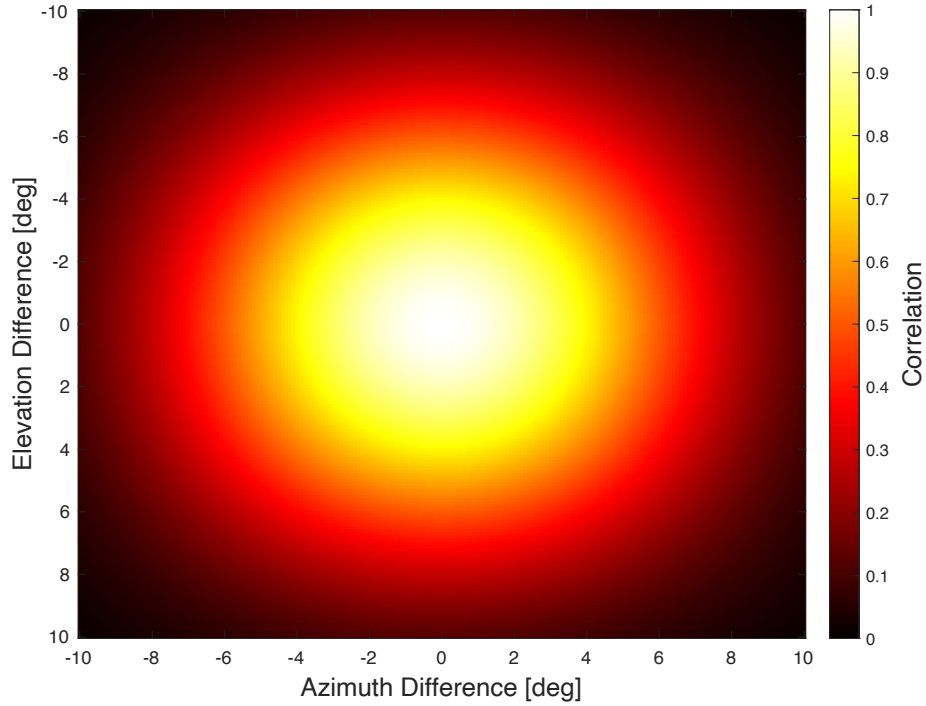


Figure 4.1 - The filtering kernel used to smooth nadir and zenith antenna gain patterns. This kernel imposes correlated error structure onto the antenna gain patterns. The coordinate system should be read as distance in the relevant antenna reference frame. Therefore, if two observations are nearby in the antenna pattern, they will have strongly correlated errors. However, if two observations are far apart in the pattern, the correlated structure decays.

For any arbitrary samples i and j , we compute the gain pattern coordinates (θ_i, ϕ_i) and (θ_j, ϕ_j) in the relevant antenna reference frame. To retrieve how correlated the error is, we compute the distance between the two observations in the reference frame:

$$\Delta\theta = \theta_i - \theta_j$$

$$\Delta\phi = \phi_i - \phi_j$$

The error correlation function is computed via a LUT of the filter kernel K :

$$E_{corr_{G^Z}} = K(\Delta\theta_Z, \Delta\phi_Z) \quad (4.30a)$$

$$E_{corr_{G^R}} = K(\Delta\theta_R, \Delta\phi_R) \quad (4.30b)$$

This error correlation holds if the samples i and j share the same antenna and are on the same spacecraft. If they are on separate antennas or spacecrafts, the correlated error is zero. The rolled-up correlated errors for the gain patterns can be expressed as follows:

$$K_{G^Z}(i, j) = \gamma \cdot E(G^Z)^2 \cdot \frac{Ecorr_{G^Z}(i, j)}{\delta \cdot (1 + \Delta\phi^2 + \Delta\theta^2)} \quad (4.31a)$$

$$K_{G^R}(i, j) = \gamma \cdot E(G^R)^2 \cdot \frac{Ecorr_{G^R}(i, j)}{\delta \cdot (1 + \Delta\phi^2 + \Delta\theta^2)} \quad (4.31b)$$

where two new tuning parameters have been introduced. γ is used to tune the overall magnitude of the correlated error from these components, and δ is used to scale the decorrelation roll-off rate as the samples spread in antenna coordinates.

4.4.4 Zenith-Specular Ratio ζ Error Model

Errors in the zenith–specular ratio ζ are defined as a function of specular incidence angle θ_{inc} , which is a function of the geometry of a given GPS transmitter, a CYGNSS receiver at any given sample time.

ζ is used to estimate GPS EIRP and is derived via the following, as described in (Wang, Ruf, Gleason, et al. 2021b):

$$\zeta \equiv \frac{EIRP_z}{EIRP_s} = \frac{EIRP(t, \theta_z, \phi_z)}{EIRP(t, \theta_s, \phi_s)} = \frac{P^T(t)G^T(\theta_z, \phi_z)}{P^T(t)G^T(\theta_s, \phi_s)} = \frac{G^T(\theta_z, \phi_z)}{G^T(\theta_s, \phi_s)} \quad (4.32)$$

where the angles are defined in the GPS reference frame. For specular geometries, the azimuthal angles in the zenith direction are nearly identical to the specular direction, so $\phi_z = \phi_s = \phi$. In addition, the elevation angles in the GPS antenna reference frame θ_z and θ_s can be estimated from the angle of incidence of specular reflection from Earth θ_{inc} :

$$\theta_z \cong \theta_z(\theta_{inc}) \quad (4.33a)$$

$$\theta_s \cong \theta_s(\theta_{inc}) \quad (4.33b)$$

As a result, ζ can be expressed as a function of the specular incidence angle and azimuthal angle in the GPS antenna reference frame. While GPS antenna patterns are known to exhibit azimuthal dependence, this variation is less significant than the elevation angle, and CYGNSS uses the azimuthal average for its EIRP estimate:

$$\zeta(\theta_{inc}) \equiv \frac{1}{2\pi} \int_0^{2\pi} \frac{G_T(\theta_z(\theta_{inc}), \phi)}{G_T(\theta_s(\theta_{inc}), \phi)} d\phi \quad (4.34)$$

The estimated correlated error in ζ , however, comes with two steps of this processing. First is the mapping of Earth scattering incidence angle θ_{inc} to GPS antenna elevation angles θ_z and θ_s in Equations (4.33a) and (4.33b). This particular mapping is coarse, as even the high-fidelity-derived GPS antenna maps are plotted to 0.5-degree increments. Because the dynamic range of θ_z only extends to about 15 degrees, that only leaves ~ 30 data points to map the full dynamic range of scattering incidence angles.

The second aspect has to do with the way in which ζ is processed and generated and invokes the same logic as in the Filter Lemma. For every GPS satellite, a ζ LUT is generated as a function of observation incidence angle θ_{inc} . To minimize discontinuities, a fourth-order power series is fit. We argue that this smoothing is the predominant source of correlated error structure. An example of this is demonstrated in Figure 4.2.




Figure 4.2 – This figure illustrates a calculated zenith–specular ratio ζ as a function of observation incidence angle θ_{inc} at a fixed GPS antenna azimuth for GPS PRN 2. The blue trace is interpolated from raw observations over a two-year period at each of the elevation gridpoints in the GPS antenna pattern for a single azimuthal cut of PRN 2. The red trace is a generated smoothed zenith–specular ratio ζ that would be similar to the ones used in the operational LUTs using a 4th-order power series fit. Note that at large incidence angles, i.e., grazing observations, there is a great deal of uncertainty in ζ because there are few valid observations in those regions. In practice, only data at incidence angles < 60 degrees constrain error in ζ .

While linear interpolation itself imparts some degree of error structure, we believe it is the most representative way to express ‘raw’ data in a continuous series for the purposes of exploring correlated error due to the power series smoothing. For each GPS PRN, we calculate the difference between these estimates:

$$\Delta\zeta_{\phi}(\theta_{inc}) = \left(\frac{G_T(\theta_z(\theta_{inc}), \phi)}{G_T(\theta_s(\theta_{inc}), \phi)} \right)_{interp} - \left(\frac{G_T(\theta_z(\theta_{inc}), \phi)}{G_T(\theta_s(\theta_{inc}), \phi)} \right)_{smooth}$$

Then, the error correlation is simply the following:

$$E_{corr_{\zeta}}(\theta_i, \theta_j) = corr(\Delta\zeta_{\phi}(\theta_i), \Delta\zeta_{\phi}(\theta_j)) \quad (4.35)$$

where θ_i is the incidence angle of the observation at sample i , and θ_j is the incidence angle of the observation at sample j . In practice, the correlation is computed by using each azimuthal cut as an instance and building a LUT of correlation as a function of incidence angles for samples i and j . An example of this LUT for GPS PRN 2 is shown in Figure 4.3. The rolled-up correlated error for ζ can then be expressed as follows:

$$R_\zeta(i, j) = \left(\beta \cdot E(\zeta)^2 \cdot Ecorr_\zeta(i, j) \right)^{\frac{1}{2}} \quad (4.36)$$

with the same tuning parameter β as introduced in Section 4.4.1.




Figure 4.3 – This figure depicts an error correlation matrix derived from Equation (E4) for GPS PRN 2, which is used to produce $Ecorr_\zeta$. The matrix is a function of the observation incidence angle for sample points i and j . Note that the mapping from coordinates in GPS elevation angle θ_z and θ_s to Earth scattering incidence angle θ_{inc} is coarse and produces the checkerboard-like pattern near the diagonal. A single error in the measurement of G_T in the GPS antenna pattern is will highly correlate within a range of incidence angles in θ_{inc} , as mapped. At high incidence angles, errors are strongly correlated as the power series fit is likely to be wrong in the same direction.

4.5 Verification Techniques

One of the primary challenges in constructing the CYGNSS Level 1 correlated error model R_{mod} is identifying a plausible validation scenario. In practice, it can be difficult to disentangle the various sources of correlated observation error structure, e.g., those caused by the inherent behavior and calibration of the instrument, by the geophysical retrieval and inversion process, and by the representativity errors imposed when observations are gridded and ingested into models. Further, the choice of ground truth may impose an additional source of correlated error, such as when using reanalysis data or another observation source.

In an effort to disentangle these correlated errors and isolate only those due to instrumental sources, this work validates the correlated error structure R_{mod} by matching up near-simultaneous collocated observations made by two different observatories at nearly identical geometries, generating modeled NBRCS σ_{mod}^o from reanalysis data, and single- and double-differencing the results.

4.5.1 Curating Matchup Observations

As a constellation of eight small satellites in the same orbital plane, CYGNSS is generally unable to collect collocated, near-simultaneous observations from different satellites. If the CYGNSS observatories were equally distributed across the orbital plane, each asset would follow the next at approximately a 10 min lag. However, in that time, both of the uniquely defining features of the CYGNSS observation changes: (1) the surface of the Earth rotates underneath the observatories, changing the ground track, and (2) the GPS satellites that serve as the source of radar signal advance in their orbit. Therefore, in 10 min, it can be quite challenging to develop one-to-one matchup conditions suitable for investigating the correlated error structure at timescales of seconds to minutes.

CYGNSS has no onboard thrusters, and orbit phasing is controlled solely through differential drag. At several junctures during the CYGNSS mission, one of the observatories advanced within the plane to be nearly overlapping with another yet at a slightly different altitude. Generally, the greater the altitude difference between the observatories, the faster the relative precession within the orbit plane.

An exhaustive review of each observatory’s ephemerides for the life of the mission identified a few matchup opportunities. Matchups near the beginning of the mission were preferred, as the orbit planes of the satellites tend to drift apart over the lifetime of the mission. After further filtering by the operational status of each observatory to ensure that they were in similar attitude configurations and operating in similar science modes, a 24 h period starting 0Z on 11 SEP 2019 was chosen, where FM1 and FM5 were in a nearly identical orbital phase for a sustained period. A representation of the ground sampling during this period is shown in Figure 4.4.

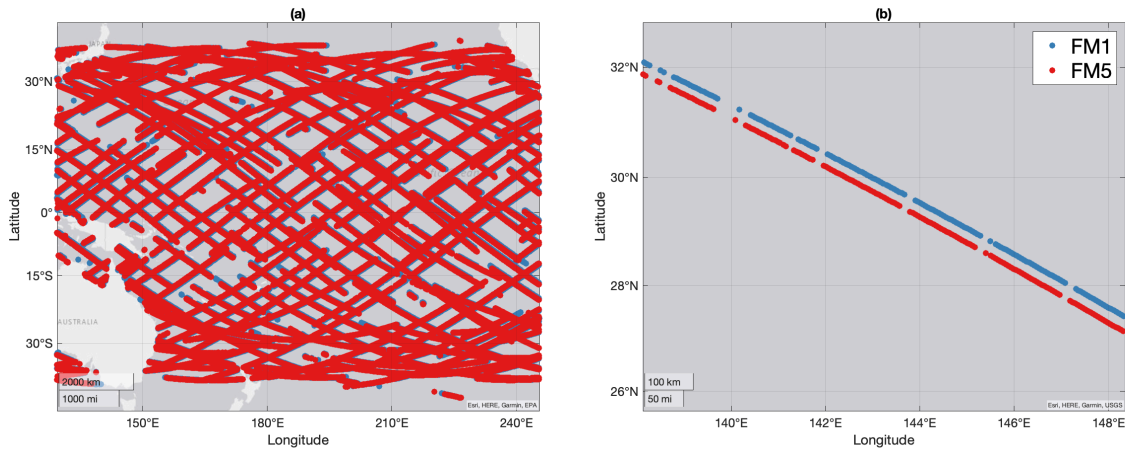


Figure 4.4 – Comparison of FM1 and FM5 observations for 11 SEP 2019. (a) Rendering of the ground samples captured during the 24 h period in the near-overlap condition; (b) A single track of samples for both FM1 and FM5 that has been matched sample-for-sample to facilitate one-to-one comparisons of the observations. The samples from each observatory are no more than 0.5 deg apart in distance and were acquired approximately 3 s apart in time.

Each ‘track’ of observations (when a series of observations are made in close succession sharing a CYGNSS receiver and GPS transmitter) was matched sample-for-sample between FM1 and FM5 by first minimizing the distance between observations and then quality controlling for several factors:

- Matched tracks must both contain more than 300 samples;
- Individual sample matchups are valid if samples are within 0.5 degrees (great circle distance);
- Individual samples are screened to ensure no quality control flags apply; and
- The matched track is only valid if 60% of the data remains after all other matchup criteria apply.

The resulting matchup conditions produced 103 matched tracks within the 24 h period of near-simultaneous observations.

4.5.2 Generating Model NBRCS

A forward model for the CYGNSS observatory (generating σ_{mod}^0 from wind data) can be challenging. Operationally, CYGNSS uses a GMF LUT that maps between the two quantities. However, for the purposes of this analysis, we prefer to use a physically representative forward operator that represents the physical dynamics of the roughening of the ocean surface due to locally driven winds.

For each sample matched up, σ_{mod}^0 is generated by computing a spectrally corrected ocean surface MSS from an ERA5-forced (Hersbach et al. 2020) WaveWatch III wave model as described in (Wang et al. 2019). Because the resolution of ERA5 is much coarser than CYGNSS Level 1 observations in both space and time, the σ_{mod}^0 is matched up with σ^0 via a tri-linear interpolation across three dimensions (latitude, longitude, and time).

Therefore, for every track of data, there are observations from two different CYGNSS observatories and two modeled NBRCS from reanalysis data.

4.5.3 Estimating Total Correlated Error

Single- and double-differencing are common techniques to calibrate remote sensing instruments, especially radiometers (Berg et al. 2021; Kroodma, McKague, and Ruf 2012) and radars (Zec et al. 2017). Both are useful for quantifying the correlated error structures in CYGNSS observations.

For each track of data, there are two single-differenced datasets,

$$SD_{obs} = \sigma_1^o - \sigma_5^o \quad (4.37a)$$

$$SD_{mod} = \sigma_{mod,1}^o - \sigma_{mod,5}^o \quad (4.37b)$$

where SD_{obs} is the single-differenced data from the two matched-up observations from FM1 and FM5, and SD_{mod} is the single-differenced data of the model-generated NBRCS for the specific coordinates of FM1 and FM5. The double-difference is computed by differencing these two quantities:

$$DD = SD_{obs} - SD_{mod} \quad (4.38)$$

Both the SDs and DD are useful for our analysis. The SD_{obs} term represents the difference in σ^o measured by two different CYGNSS satellites. Despite the strict matchup conditions, these assets are still measuring different fields of view at different times. These differences may result in some residual correlated structure. In addition, any correlated structure in SD_{obs} may not reveal structure imposed from fundamental properties of the earth system. Therefore, SD_{mod} allows us to identify the correlated structure of any systematic differences in observation target. An example of the utility of these differencing techniques is shown in Figure 4.5.




Figure 4.5 – A display of model (blue) and observed (green) single-differenced matchup data for a single CYGNSS track, with the double-differenced data (red) overlaid. Note that both the model and observed single differences are similarly high early in the track, suggesting that the differences in the observations may be because the samples are observing fundamentally different targets, captured in SD_{mod} . The double-difference accounts for these types of errors, and for the entire track, the double-difference is quite stable. The sparsity of data is caused by quality control parameters, such as when the CYGNSS observatory is performing its onboard calibration procedure, which occurs once every minute. Because the two matched observatories do not have synchronized calibration clocks, the matched dataset typically flags out two of these cadences for every minute of sampling.

Since we are primarily interested in correlated error, and not absolute error, our primary investigative tool is the autocorrelation function. The correlated error for each track can be computed by autocorrelating the DD using the standard formulation:

$$\rho(\tau) = \frac{1}{(N - \tau)} \frac{1}{\sigma_{x_i}} \frac{1}{\sigma_{x_{i+\tau}}} \sum_{i=1}^{N-\tau} (X_i - \bar{X})(X_{i+\tau} - \bar{X}) \quad (4.39)$$

where τ is the lag, N is the total number of lags observed, X_i is the value of timeseries X at time i , \bar{X} is the mean of timeseries X , and the standard deviation is formulated as normal, $\sigma_i =$

$$\sqrt{\frac{\sum_{i=1}^{N-\tau} (X_i - \bar{X})^2}{N - \tau}}.$$

Because the error correlation behavior can vary from track-to-track, we also construct a bulk autocorrelation which approximates the autocorrelation behavior for all tracks sampled. For this we consolidate M tracks, the bulk autocorrelation is:

$$\mathcal{R}(\tau) = \frac{1}{M(N-\tau)} \frac{1}{\sigma_{y_i}} \frac{1}{\sigma_{y_{i+\tau}}} \sum_{i=1}^{N-\tau} (Y_i - \bar{Y})(Y_{i+\tau} - \bar{Y}) \quad (4.40)$$

where $Y_i = \sum_{j=1}^M X_{i,j}$, j is the index for track number, and the standard deviation and mean are calculated as before, but with the consolidated track series Y .

4.5.4 Model Tuning

The constructed correlated error model $R_{mod}(i, j)$ is designed to estimate the correlated error between arbitrary samples i and j . Validating the correlated error with empirical matches is challenging since the empirical error correlation can only be estimated in a broader, statistical sense. To create an appropriate comparison, we introduce the *modeled* error autocorrelation $\tilde{\rho}_{mod}$, which is a function of lag τ :

$$\tilde{\rho}_{mod}(\tau) = \frac{1}{N-\tau} \sum_{i=1}^{N-\tau} R_{mod}(i, i+\tau) \quad (4.41)$$

where R_{mod} is calculated using Equation (4.17), τ is the lag, and N is the total number of lags observed. The quantity $\tilde{\rho}_{mod}$ is reasonably comparable to the autocorrelation $\rho(\tau)$ of observed error for a single track of samples. A similar analog can be made for the *bulk modeled* error autocorrelation $\tilde{\mathcal{R}}_{mod}(\tau)$, which estimates the autocorrelation behavior of the error model across M tracks:

$$\tilde{\mathcal{R}}_{mod}(\tau) = \frac{1}{M} \sum_{j=1}^M \tilde{\rho}_{mod,j}(\tau) \quad (4.42)$$

where $\tilde{\rho}_{mod,j}(\tau)$ is the modeled autocorrelation for track j at lag τ .

A parametrized version of R_{mod} is constructed using the following form:

$$R_{mod}(\alpha, \beta, \gamma, \delta)|_{i,j} = \frac{1}{\mathcal{N}} [R_{Pg}(\alpha, \beta) + R_{Pz}(\alpha, \beta) + R_{GR}(\gamma, \delta) + R_{GZ}(\gamma, \delta) + R_{\zeta}(\beta)] \quad (4.43)$$

Where the following is true:

- α represents the relative magnitude of the white noise component of the error, which decorrelates at $\tau = 1$;
- β represents the relative magnitude long-decay pedestal or any residual correlated errors at the edge of our timescales of interest;
- γ represents the relative magnitude of the correlated error caused by the terms R_{GR} and R_{GZ} , which exhibit smooth decay as samples spread apart when projected through the nadir and zenith antenna coordinates, respectively [see Appendix D for an in-depth discussion]; and
- δ represents the relative decorrelation roll-off in terms R_{GR} and R_{GZ} .

With an appropriate benchmark, the tuning parameters α , β , γ , and δ are iterated such that $\tilde{\mathcal{R}}_{mod}$ matches a target signal. The tuning parameters are applied such that they can be modified to change specific characteristics of the modeled behavior. The target for matching the modeled autocorrelation $\tilde{\mathcal{R}}_{mod}$ is the bulk autocorrelation of the double-differenced data DD:

$$\mathcal{R}[DD](\tau) \cong \tilde{\mathcal{R}}_{mod}(\tau) \quad (4.44)$$

Because this is an under-determined system, there is no unique solution to optimizing the tuning parameters. Further, these parameters all relate to one another, and changing one will impact the others. Instead, R_{mod} is tuned heuristically such that the modeled behavior matches the empirical data at key points: to match the white noise component at lag $\tau = 1$; to match the rolloff at lags $\tau = 5$ and $\tau = 30$; and to match the endpoint behavior at lag $\tau = 100$.

4.6 Results

4.6.1 Bulk Behavior

The constructed model R_{mod} is first tuned to match the overall behavior of $\mathcal{R}[DD](\tau)$, as described in Section 4.5.4. With the appropriate tuning parameters, the bulk-modeled autocorrelation $\tilde{\mathcal{R}}_{mod}$ closely resembles $\mathcal{R}[DD](\tau)$ in most important aspects, including the relative magnitude of the white noise component, the relative rate of decorrelation roll-off, and endpoint behavior at large lags. The final tuned $\tilde{\mathcal{R}}_{mod}$ is shown in Figure 4.6 with $\mathcal{R}[SD_{obs}]$, $\mathcal{R}[SD_{mod}]$, and $\mathcal{R}[DD]$.

The behavior in Figure 4.6 is worth discussing in detail. The single-differenced model generated σ_{mod}^o decays at a much slower rate than the single-differenced observations, and double-differenced data suggests that the errors of slight mismatch in observation location and time are generating errors on a fundamentally different scale than the instrument errors. Further, the fact that the double-differenced decorrelation is almost identical to the observation single-differencing suggests that single-differencing near-simultaneous observations are a reasonable approximation for bulk error correlation.




Figure 4.6 – Comparison of the bulk error autocorrelation behavior for all matched tracks. The single difference of modeled σ_{mod}^o (blue), the single difference of observed σ^o (green), the double-difference DD (purple), the untuned bulk model error correlation $\tilde{\mathcal{R}}_{mod}$ (dotted red), and the final tuned bulk model error correlation $\tilde{\mathcal{R}}_{mod}$ (orange) are shown. The shading indicates the estimated 1-sigma standard deviation of the population decorrelation behavior for each individual trace. The lags are computed in seconds, which correspond to single samples at 1 Hz sampling for the CYGNSS observatory. Note that DD generally follows the SD_{obs} trace, suggesting that single-differencing for this particular use case may be a reasonable representation of bulk error correlation behavior.

Evaluating the double-difference trace in Figure 4.6 also reveals important qualities of the CYGNSS error correlation structure. First, the uncorrelated error component accounts for roughly 5% of the total system error. Therefore, assumptions that samples may be treated as independent with uncorrelated error are empirically refuted. Further, the error decorrelation roll-off is swift: about 50% of the error decorrelates within 7 s of observation. Using a flat-plane approximation of Earth, this corresponds to a distance of approximately 50 km, which is generally the scale of the gridding of modern global weather models but much coarser than the

resolution of state-of-the-art regional models. Finally, the endpoint behavior at large lag times suggests a small, positive correlated error ($\sim 2\%$) at larger timescales. The statistical properties of the lag correlation become less stable at larger lags, but the existence of a long-timescale correlation is not surprising, considering all observations from a single receiver share a common electronics and processing chain.

The fact that our tuned model R_{mod} can approximate the bulk error correlation structure $\tilde{\mathcal{R}}_{mod}$ with similar features as $\mathcal{R}[DD]$ validates the assumption that the overall instrument-correlated error can be modeled from the fundamental components of the instrument observable: in our case, the radar-range equation (Equation (1.1)). Further, that R_{mod} can be generated between arbitrary samples suggests that an observation error correlation matrix \mathbf{R} can be generated dynamically from first principles given appropriate knowledge about the instrument and retrieval.

The value of the chosen tuning parameters is also worth investigating. The untuned model is defined by having the parameters $\alpha = \beta = \gamma = \delta = 1$. Figure 4.6 demonstrates that the modification of tuning parameters can change the overall model behavior significantly to match observed behavior. Figure 4.6 further suggests that the untuned R_{mod} generally overestimates both the relative magnitudes of uncorrelated error (i.e., white noise, tuned by α) and totally correlated error (i.e., endpoint correlation, tuned by β). The chosen parameters for tuning are shown in Table 4-3.

Table 4-3 –The chosen magnitudes of model tuning parameters for R_{mod} .

Tuning Parameter	Magnitude	Function
α	0.005	Relative magnitude of uncorrelated error
β	0.01	Relative magnitude of endpoint correlated error
γ	1	Relative magnitude of nearby roll-off
δ	1	Steepness of roll-off component

The tuning parameter γ is used to vary the relative magnitude of the near-sample correlated error roll-off. In our tuned model, this parameter remains at 1. The tuning parameter δ is used to enhance or reduce the steepness of the roll-off in correlation caused by two observations moving farther apart in the nadir and zenith coordinate systems. If $\delta > 1$, it enhances the steepness of decay. The fact that our optimized tuning maintains $\delta = 1$ suggests that the filtering theory discussed in Lemma 4.2 is a reasonable model of decay.

4.6.2 Single-Track Comparisons

The tuned model can be compared to single-track autocorrelations $\rho(\tau)$ of matched-sample double-differences using the modeled error autocorrelation $\tilde{\rho}_{mod}$ as described in Equation (4.41). We compare three exemplar tracks in Figure 4.7.

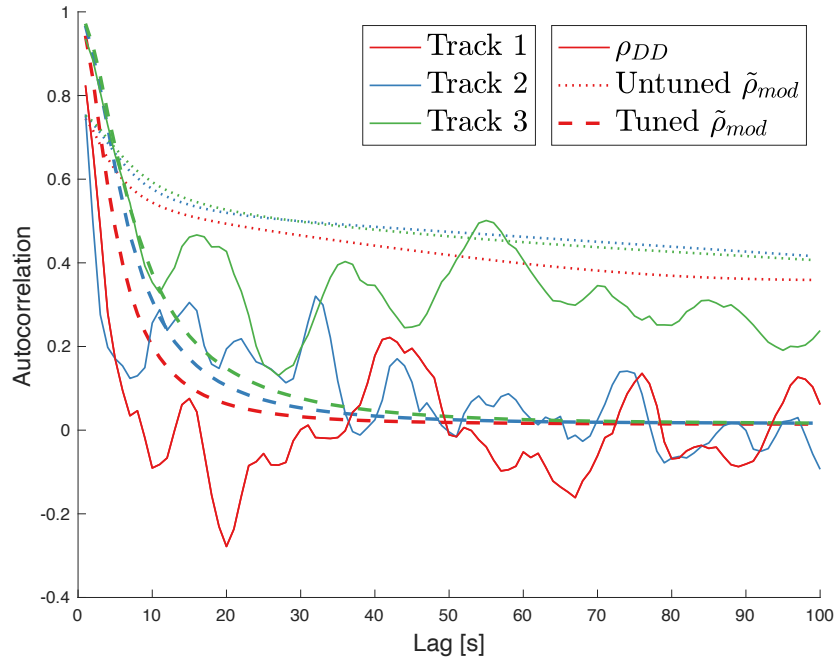


Figure 4.7 – Comparison of the single-track error autocorrelation behavior for three selected tracks. The autocorrelation of the NBRCS double-difference (solid lines), untuned modeled error correlation $\tilde{\rho}_{mod}$, and tuned modeled error autocorrelation $\tilde{\rho}_{mod}$ (dashed lines) for each of the three tracks are shown. Each track is painted a different color. This figure demonstrates the wide variability of single-track autocorrelations of double-differences, arising from the variability of the observable, the limited amount of data in a single track, and the challenges in quality controlling sufficient observation data. It is possible that the correlated error does vary this much from track to track. In contrast, both the tuned and untuned modeled error autocorrelation $\tilde{\rho}_{mod}$ are much more stable from track to track.

The single-track autocorrelation behavior shown in Figure 4.7 is also revealing of the limitations of this work. The single-track error autocorrelation for double-differenced data is highly variable. This may be due to both artifacts in the data (processing, quality control, insufficient data), as well as the real behavior of the observation. It is worth articulating that the autocorrelation of data is generally less stable with smaller datasets and at longer lags. With our quality control parameters, we flag out significant quantities of data, which decreases the stability of autocorrelations from track to track. We further note the difficulty in exploring the dynamics of this behavior in a statistical sense: our aperture of observation where two CYGNSS assets are in a near-overlap condition is quite rare, and as the mission progresses, the orbit planes of the satellites have drifted, making further matchup scenarios less representative. These

behaviors may evolve differently day-to-day or as a function of observed wind speed, but the paucity of data in this configuration makes it challenging to establish sufficient baselines to test for significance.

4.6.3 Dynamic Correlated Error Estimation and Impact of Tuning

The model R_{mod} can produce plausible dynamics, suggesting the broader importance of a realistic dynamic correlated error model. Figure 4.8 plots two nearby tracks captured nearly simultaneously by the same receiver (but different GPS transmitters).

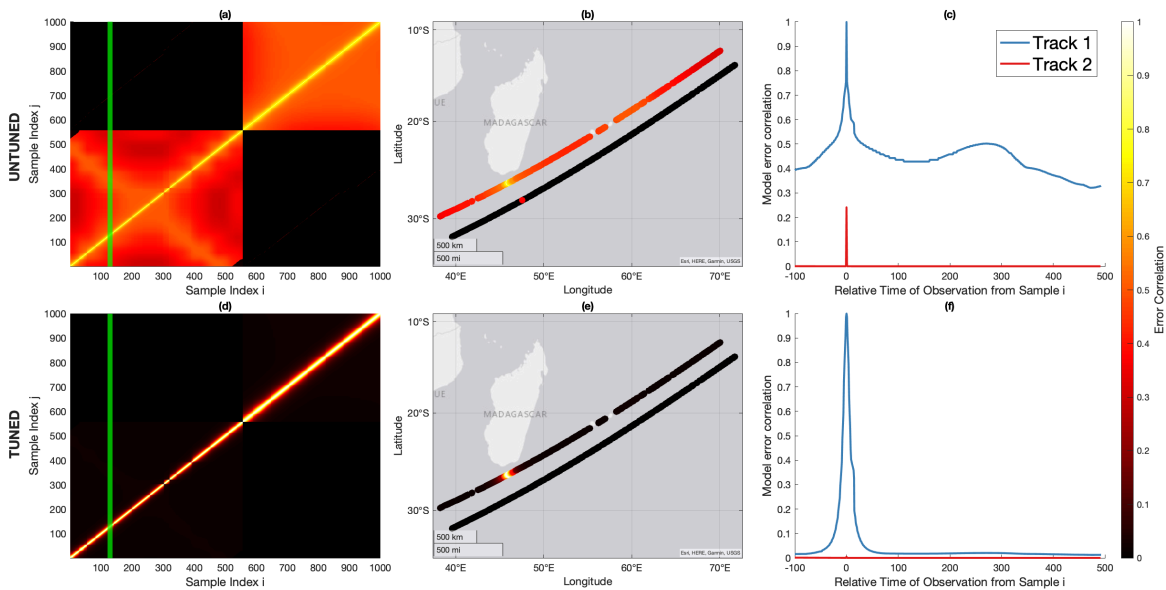


Figure 4.8 – Generated untuned and tuned R_{mod} for two nearby tracks captured by the same CYGNSS receiver, each with three different representations. (a) Untuned R_{mod} as represented by sample index. The two tracks are concatenated in a vector, and the modeled error correlation is calculated for each combination of indices. The green vertical line represents an exemplar index ($i = 129$) along the track, where the correlated error is estimated for every other observation in the neighborhood. (b) Untuned R_{mod} for the same exemplar index as represented by the physical location of sample acquisition. (c) Untuned R_{mod} for the same exemplar index represented by time of acquisition. The traces for the two different tracks are plotted in different colors. The error correlation along the same track is in blue, while the error correlation for the adjacent track is in red. (d) Tuned R_{mod} as represented by sample index for the same two tracks as in (a). (e) Tuned R_{mod} for the same tracks represented in the physical location. (f) Tuned R_{mod} for the same exemplar index represented by the time of acquisition. Note that for both the untuned and tuned cases, the correlated error for the adjacent track is non-zero but generally very small compared to the error correlation along the track.

Because the correlated error R_{mod} can be generated for any arbitrary observation using the bottom-up model, the dynamics of how instrument errors decorrelate can be explored. To illustrate, we choose an arbitrary exemplar index ($i = 129$) to demonstrate that this calculation can be performed for any sample within a track. If the observations were treated as completely independent without any correlated error, the matrix in Figure 5a would contain non-zero elements exclusively along the main diagonal. However, we observe several structural elements. The ‘pixelation’ pattern is largely a result of the R_{ζ} term and originates from the coarse mapping of the estimated GPS antenna gain pattern as a function of scattering incidence θ_{inc} . This phenomenon is explored in depth in Section 4.4.3. The smooth decorrelation roll-off is from the R_{GR} and R_{GZ} components of the model and represent the direct and reflected signals moving about the nadir and zenith antenna patterns in CYGNSS receivers as the observation is collected along the track. We note that the model allows for cross-track correlation where the two different tracks share a CYGNSS receiver but have a different GPS transmitter. In both the untuned and tuned cases, these tracks appear to have uncorrelated cross-track error. There are residual correlations between the tracks when the scattering geometry is such that the two tracks share similar incidence angles or are near similar antenna coordinates, in addition to the shared receiver noise at lag $\tau = 0$.

We can also determine R_{mod} for our matchup tracks where CYGNSS observes similar track geometries with the same GPS transmitter but with two different receivers. The correlated structure for one of the tracks in the previous example (but matched by different receivers) is shown in Figure 4.9.

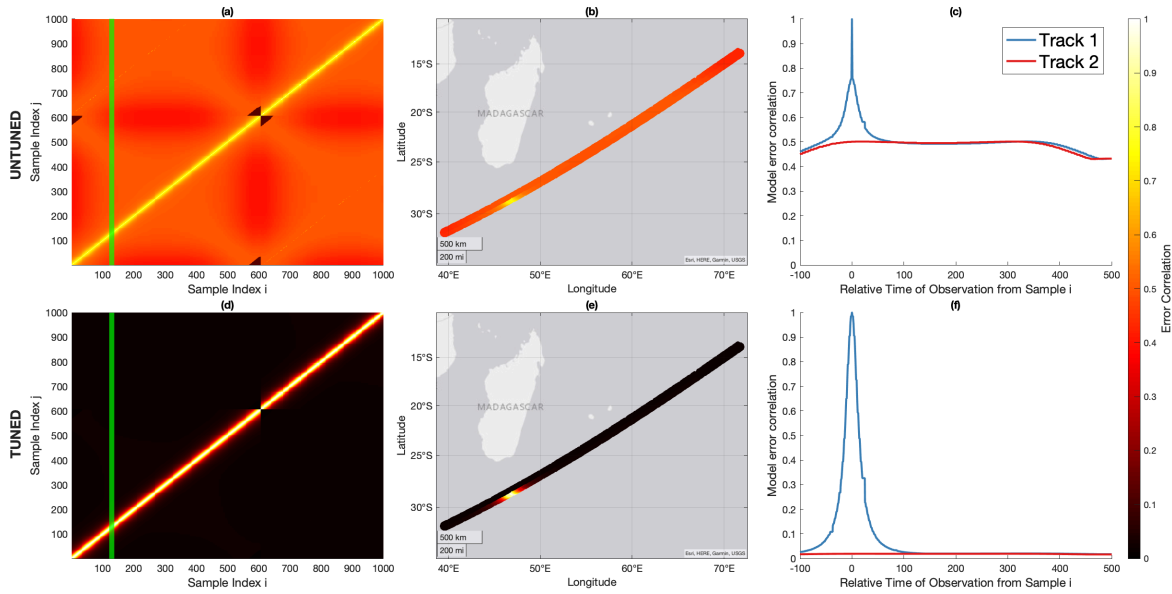


Figure 4.9 – Generated untuned and tuned R_{mod} for two nearby tracks captured by two different CYGNSS receivers during the overlap period, each with three different representations. **(a)** Untuned R_{mod} as represented by sample index. The two tracks are concatenated in a vector, and the modeled error correlation is calculated for each combination of indices. The green vertical line represents an exemplar index ($i = 129$) along the track where the correlated error is estimated for every other observation in the neighborhood. The correlated error model allows for correlated errors across receivers. **(b)** Untuned R_{mod} for the same exemplar index as represented by the physical location of sample acquisition. **(c)** Untuned R_{mod} for the same exemplar index represented by the time of acquisition. The traces for the two different tracks are plotted in different colors. The error correlation along the same track is in blue, while the error correlation for the adjacent track is in red. **(d)** Tuned R_{mod} as represented by the sample index for the same two tracks as in **(a)**. **(e)** Tuned R_{mod} for the same tracks represented by the physical location. **(f)** Tuned R_{mod} for the same exemplar index represented by the time of acquisition. Note that the tuning has virtually eliminated cross-track error correlation.

The model R_{mod} allows for correlated error between two different receivers that share a GPS transmitter. We note that observations sharing the same GPS transmitter may contain correlated error due to the correlated misestimation of GPS transmit power. This may be an incomplete articulation of the full cross-receiver error correlation. For example, CYGNSS assets in similar orbital regimes may experience similar thermal and environmental conditions that cause correlated errors during our timescales of interest. We assume that correlated error due to misestimation of GPS transmit power is nearly constant for the timescales of interest.

As demonstrated in Figure 4.9, the cross-track correlation virtually disappears after tuning the model R_{mod} . Taken with the data presented in Figure 4.8, there appears to be

negligible cross-track error correlation, either when two tracks share a CYGNSS receiver or when two tracks share a GPS transmitter. This is a novel result but challenging to validate in practice, as the double-differencing may eliminate any shared error correlation between two tracks.

4.7 Discussion

R_{mod} was designed to maintain several key features. First, R_{mod} can be generated for arbitrary samples in any dataset given the appropriate inputs. Second, R_{mod} is designed to maintain traceability from component error sources with reasonable physical assumptions. Third, R_{mod} is designed to be tuned to allow for calibration and validation as more data becomes available. Finally, R_{mod} is intended to be implementable in a dynamic error model without significant computational expense.

4.7.1 Limitations

The correlated error model R_{mod} exhibits complex dynamics that are generally plausible but difficult to validate outside of the large-scale statistical estimation discussed in Section 4.6.1. The fact that the bulk statistical behavior between the estimated autocorrelation model averaged over all tracks is consistent with the observed double-difference autocorrelation (cf. Figure 4.6) is reassuring, but the disparity between the individual track model estimated autocorrelations and the single-track double-difference estimations (c.f. Figure 4.7) suggests that either the model is insufficient at capturing the dynamics track-to-track or that the correlated error may not be as stationary as assumed.

With only about a day of data during the overlap period, there remains the possibility that the validation dataset is unrepresentative of the overall statistical behavior of the observation.

R_{mod} is designed only to account for errors at short spatial and temporal scales relevant for data assimilation purposes; therefore, this model may not necessarily account for changes of the underlying statistical distributions at larger scales due to seasonality or orbital precession.

The error model assumes that errors are generally isotropic and that $R_{mod}(i, j) = R_{mod}(j, i)$. This assumption has important practical utility for simplifying the implementation of R_{mod} and the conditioning of the required matrix inversion. Anisotropies may exist in any of the components of the error model but are assumed to be small compared to the modeled behavior. Further, this model generally assumes no spatial dependence. Spatial dependence of the correlated error is primarily driven by the errors in the retrieval and not the instrumentation, which orbits Earth every 95 min. The ‘spatial’ dependence of error from CYGNSS transiting its orbit has largely to do with the dynamics of the thermal loading of the receiver as lighting conditions evolve in orbit. This is accounted for in the assumptions of the component models R_{pg} and R_{pz} as discussed in Sections 4.4.2 and 4.4.3, respectively. The correlated error resulting from the evolution of the observing geometries as both GPS and CYGNSS propagate in their orbits is captured in the correlated structure of the antenna gain maps and the zenith–specular ratio explored in Sections 4.4.4 and 4.4.5.

4.7.2 Impact of Tuning Parameters

The application of tuning parameters is a design decision to capture the fact that initial estimates of the bottom-up error magnitudes may be incomplete and to facilitate rapid adjustments as new calibration and validation information becomes available. The tuning parameters attempt to weigh the value of individual component errors in comparison to each other based on aggregate data from one day of observations during the overlap period.

The specific values of the tuning parameters recovered suggest that the correlated error model R_{mod} overestimates both the uncorrelated and highly correlated components of instrument error. This is useful information for future studies because, as an opportunistic measurement, GNSS-R has limited insight into the error structures in the source signals from GPS.

For observations within a single track, these parameters provide evidence that nearby samples in space and time can add significant new information to a forecast. The tuning parameter $\alpha = 0.005$ indicates that the uncorrelated component of the error is significantly overestimated, suggesting that the constituent error magnitudes in (Gleason 2018) may be overestimated. Further, the fact that the tuning parameters γ and $\delta = 1$ suggest that the overall structure of the error correlation decay is consistent with the theory posited in Appendix D (that the primary source of error correlation is from the application of smoothing filters in the production of antenna gain patterns) and that the application of the Filtering Lemma (Lemma D2) is reasonable.

For observations between tracks, the tuned error model R_{mod} indicates that errors are nearly uncorrelated between tracks (cf. Figures 4.8 and 4.9). This is driven primarily by the tuning of the parameter $\beta = 0.01$. This has significant practical utility to future assimilation strategies for CYGNSS, suggesting that two nearby tracks are essentially independent measurements with independent instrument error (correlated error may still result from correlated errors in the retrieval or representation).

Finally, it should be noted that the overall effect of the tuning somewhat obviates the need for a dynamic correlated error model, as the overall behavior of the tuned model is quite stable from track to track (cf. Figure 4.7). This suggests a simplified instrument-correlated error

model could be derived from the bulk behavior of $\tilde{\rho}_{mod}$, reducing many of the required input data.

4.8 Conclusions

This work produces a first-principles estimate of correlated instrument error with results that approximate observed statistical behavior. We believe that this model presents a significant advancement in the estimation of the spatial and temporal correlation structure of instrument error for remote sensing systems and, in particular, for the unique considerations of GNSS-R measurements by CYGNSS.

In essence, this work answers a theoretical exercise for enumerating the plausible engineering reasons why two data points from the same observing constellation can share correlated sources of error. We evaluate the correlated error structure for CYGNSS by examining the potential plausible sources of correlated structure from individual components of the instrumentation, combining these sources from first principles as a tuned engineering model, and evaluating the efficacy of the model via a robust validation during a period when two satellites with nearly identical observing geometry captured near-simultaneous and near-collocated samples.

The instrument-originating sources of correlated error is likely to be a small component of the overall correlated structures of observation error. For instance, a likely significant source of correlated error structure is the Geophysical Model Function retrieval that converts observed NBRCS to surface wind speed. These errors can be multifaceted, both encompassing representation error, as the ground truth for training this retrieval is reanalysis data, which may not capture the spatial or temporal dynamics of wind speed (Ruf and Balasubramaniam 2019). A

companion work will explore how the retrieval that maps from NBRCS to wind speed produces correlated error structures in space and time.

Further, the fact that the single-differenced observations generally drive the double-differenced autocorrelation behavior suggests that Powell et al.'s metric of simultaneity and collocation (Powell, Ruf, Gleason, et al. 2024) generally applies for CYGNSS when satellites are in a near overlap condition. This interesting result suggests that double-differencing may not be required for statistical estimations of GNSS-R errors given a sufficiently large dataset of observations that nearly overlap. This implication may relax further validation requirements of R_{mod} and may enable near-real-time calibration strategies when samples are sufficiently close without the need to generate reanalysis-driven forward model NBRCS, which introduces significant latencies.

Finally, while the correlated error model R_{mod} was designed to compute the estimated correlated error from arbitrary CYGNSS samples, this may be unnecessary in practice as the tuning makes the model quite stable between samples and tracks. Further, given sufficient assumptions about the stationarity of the instrument-correlated error and the overall magnitude of the instrument-originating sources of correlated error, this information could be conveyed as a static look-up table.

Data Availability Statement: Publicly available datasets were analyzed in this study. CYGNSS Level 1 data can be found at NASA's Physical Oceanography Distributed Active Archive Center (PO.DAAC) (CYGNSS 2021, 1). ERA5 reanalysis can be found at (Hersbach et al. 2023).

Chapter 5 Error Covariance in Observation Space and Geophysical Partitions

This chapter propagates the error covariance modeled in Chapter 4 into observation space and suggests a method for partitioning CYGNSS representation-retrieval error into geophysical components. This chapter is substantially derived from a manuscript submitted to AMS's *Journal of Atmospheric and Oceanic Technology* under the title "A Method for Estimating and Partitioning the Covariance of Windspeed Observation Error from GNSS Reflectometry" (Powell and Ruf 2024). The introductory material is discussed in Chapter 1. Copyright in this work may be transferred without notice.

5.1 Abstract

This work introduces a mechanism for identifying and partitioning the sources of correlated error in windspeed observations made by spaceborne GNSS-reflectometry (GNSS-R). Observations by the CYGNSS constellation of GNSS-R satellites are matched with reanalysis windspeeds to determine the retrieval error. A correction for some of the error is constructed using ancillary environmental parameters. As a result of the corrections, the windspeed retrieval is improved by 11%. Further, we examine the correlated structure of both the windspeed error and the correction to provide physically plausible explanations for the correlated errors that are present in GNSS-R observations.

5.2 Introduction (abridged)

This work expands upon the (Powell, Ruf, McKague, et al. 2024) instrument correlation model (discussed in Chapter 4) to include the correlated error in CYGNSS windspeed retrieval, and applies a “top-down” estimation to assist in partitioning sources of correlated error between the instrument and the combined retrieval and representation errors.

5.3 Error Model Construction (abridged)

5.3.1 Definitions

In this section, we introduce a simple model for how observation errors occur. We derive our language from Janjić (Janjić et al. 2018), but apply slightly modified definitions for both brevity and clarity. We assert that *observation error* is simply the difference between an observed and true state for a state vector \mathbf{y} :

$$\mathbf{e}_{obs} = \mathbf{y}_{obs} - \mathbf{y}_{true} \quad (5.1)$$

where \mathbf{e}_{obs} is the observation error in state space, \mathbf{y}_{obs} is the observation vector in state space, and \mathbf{y}_{true} is the true value of the state space parameter. With this formulation, an \mathbf{R} matrix can be constructed for samples i and j by taking the covariance of \mathbf{e}_{obs} :

$$R(i, j) = cov(\mathbf{e}_{obs}(i), \mathbf{e}_{obs}(j)) \quad (5.2)$$

We further assert that the sources of observation error consist of *instrument error*, *retrieval error*, and *representation error*, which are assumed to be independent and can be summed to produce the total error:

$$\mathbf{e}_{obs} = \mathbf{e}_{inst} + \mathbf{e}_{ret} + \mathbf{e}_{rep} \quad (5.3)$$

Instrument and retrieval errors concern deviations from truth due to the design and operation of the observing system. Instrument errors are inherent in all physical systems because the

measurement is never perfectly accurate. This occurs from imperfections in manufacturing, limited knowledge of the operating environment, and deviations from operational assumptions.

Observatories rarely measure the target parameter directly, often these systems take a series of measurements in *measurement space*, which represents the physical quantity an instrument measures. These physical quantities – power received, time of flight, or Doppler shift – are representative of the target parameter, but require a transformation to *observation space* to be understood in terms of geophysical parameters such as temperature, pressure, and windspeed. Retrieval errors occur in any situation during which a measurement proxy is used for a target observation and when the mapping between measurement space and observation space is not one-to-one or unique.

Representation error occurs when the instrument and its reference measurement do not correspond to identical versions of the geophysical parameter. For example, they may represent different spatial or temporal averages of the parameter as discussed in Janjić, or they may be made at different times or locations as explored in (Powell, Ruf, Gleason, et al. 2024). Representation error can vary depending on the choice of “ground truth” reference.

We can define these error types with some mathematical formalism to illustrate the relationships between them. Instrument error occurs because the observing system is not measuring the true observation space parameter:

$$\mathbf{e}_{inst} = H(\mathbf{x}_{obs}) - H(\mathbf{x}_{true}) \quad (5.4)$$

where \mathbf{x}_{obs} is the observation vector in measurement space, \mathbf{x}_{true} is the true observable in measurement space, and H is the observation operator that transforms observations from measurement space to geophysical parameter space. Using a similar logic, we define retrieval error as:

$$\mathbf{e}_{ret} = H(\mathbf{x}_{true}) - H_{perf}(\mathbf{x}_{true}) \cong H(\mathbf{x}_{obs}) - H_{perf}(\mathbf{x}_{obs}) \quad (5.5)$$

where H_{perf} is a perfect observation operator, or in other words, will perfectly map the true observable \mathbf{x}_{true} to geophysical parameter space \mathbf{y}_{true} with zero ambiguity or imposed error. For simplicity, we assume that the retrieval error for observed measurements is equivalent to the retrieval error for true measurement space observations. This formulation necessarily suggests that the mapping of H is imperfect. Additionally, the construction of (5.5) implicitly suggests that instrument and retrieval errors are independent. In practice, this assumption may not always hold, and imposes conditions on the sensitivity of the retrieval to instrumentation error. For our illustrative exercise, we consider these errors as independent and partitionable.

Because representation error is fundamentally comparative, it is challenging to generalize. For instance, a satellite observation compared against a buoy will have different representation errors than when compared to another type of satellite observation or a model output. Further, in our case, representation errors and retrieval errors may be explicitly related, as the training data for the empirical CYGNSS retrieval are ERA5 reanalysis matchups (C. S. Ruf and Balasubramaniam 2019). Because of this, we will treat the retrieval and representation error as a combined entity, \mathbf{e}_{rr} :

$$\mathbf{e}_{rr} = \mathbf{e}_{ret} + \mathbf{e}_{rep} \quad (5.6)$$

Observation error now becomes:

$$\mathbf{e}_{obs} = \mathbf{e}_{inst} + \mathbf{e}_{rr} \quad (5.7)$$

In this construction, we assert that \mathbf{e}_{inst} and \mathbf{e}_{rr} are separable and independent components of observation error.

5.3.2 Observation Error Covariance Model

This work considers how the observation error varies in space and time – in particular how its correlation varies with separation. We make two general assumptions: first, that the representation-retrieval error \mathbf{e}_{rr} is isotropic in space and can be parameterized by spatial distance between two samples i and j ; and second, that the instrument error \mathbf{e}_{inst} is isotropic in time and can be parameterized by temporal distance between two samples i and j .

These assumptions warrant discussion. The retrieval component of \mathbf{e}_{rr} signifies all of the reasons CYGNSS may be sensitive to variations in NBRCS *other* than locally-driven winds. For instance, the CYGNSS windspeed retrieval may be incorrect when the satellite samples areas of ocean surfaces covered by oil slicks, biomass, or some other contaminant; or when a significant long-wave swell modulates the overall roughness condition; or any host of other possible reasons why the mapping between surface roughness and windspeed is not as expected. The representation component of \mathbf{e}_{rr} signifies when CYGNSS captures true windspeed dynamics and variability that is not properly resolved by ERA5. In each of these cases, anisotropies may apply locally for various reasons. However, we posit that with a sufficiently large sample size, each constituent source of the combined representation-retrieval error has a characteristic length scale that does not vary significantly in the aggregate. We further assert that correlated errors from \mathbf{e}_{rr} occur at timescales longer than are relevant for operational data assimilation systems (i.e., seconds to minutes), and can be corrected via other means such as diurnal or seasonal trends.

The instrument error \mathbf{e}_{inst} was modeled by Powell et al. (Powell, Ruf, McKague, et al. 2024) and is primarily a function of time, as the errors relate due to variability in the environment of the CYGNSS receiver, the evolution of the observation geometry, and the condition of the GPS transmitter. CYGNSS collects streaks of observations known as “tracks”,

which occur when adjacent samples in time share a common CYGNSS receiver and a GPS transmitter; said tracks tend to exhibit related error characteristics. Because of the nature of how CYGNSS and GPS orbits evolve, the specular points along each track of CYGNSS observations can vary in spacing. At a 1 Hz sampling rate (for observations 2019 and prior), samples are between 2 km and 7 km apart. However, the distance between samples is nearly constant within a track. Therefore, a trackwise mapping between time and space can be calculated:

$$\tau(i, j) = t_1 - t_2 \rightarrow \Delta(i, j) = c_{track} \cdot \tau \quad (5.8)$$

where τ is the distance in time between samples i and j at observation times t_1 and t_2 , respectively; Δ is the distance in space on a great circle projection; and c_{track} is the nearly constant observation velocity for a given track in km/s. With this mapping, \mathbf{e}_{inst} and \mathbf{e}_{rr} can share a common spatial separation index, and the covarying errors can be expressed as a function of distance between samples.

The combined error covariance takes the form derived in Appendix A of Powell et al. 2024 (Powell, Ruf, McKague, et al. 2024):

$$R_{mod}(i, j) = R_{inst}(i, j) + R_{rr}(i, j) \quad (5.9)$$

where R_{mod} is the modeled observation error covariance between samples i and j , R_{inst} is the calculated instrument error covariance between samples i and j , R_{rr} is the calculated representation-retrieval error covariance between samples i and j . Further,

$$R_{inst}(i, j) = E_{inst}^2 K_{inst}(i, j) \quad (5.10a)$$

$$R_{rr}(i, j) = E_{rr}^2 K_{rr}(i, j) \quad (5.10b)$$

where E_{inst} is the error magnitude of the instrumentation error in state space, K_{inst} is the instrument error correlation between samples i and j , E_{rr} is the error magnitude of the combined representation-retrieval error, and K_{rr} is the representation-retrieval error correlation between

samples i and j . The correlation functions K_n have the range of $[-1,1]$, and $K_n = 1$ if $i = j$. This construction necessarily implies that the error magnitudes E_{inst} and E_{rr} are stationary with respect to location and that the instrument and representation-retrieval errors are independent.

5.3.3 Correlated Instrument Errors in Observation Space

Powell et al. 2024 provides a model for correlated instrument error, but the errors are calculated in measurement space (i.e., in units of NBRCS). To propagate these errors to observation space (units of m/s), the instrument error covariance is run through the GMF. The instrument error correlation function K_{inst} is the tuned $\tilde{\mathcal{R}}_{mod}$ from Powell et al. 2024, which is nearly stationary across tracks and observing conditions. The instrument error magnitude takes the form:

$$E_{inst}(i) = \left[\left| \frac{\partial(GMF^{-1})}{\partial \sigma^o} \right| std(\sigma^o) \right]_{\sigma^o(i)} \quad (5.11)$$

where GMF^{-1} represents the inverse of GMF, i.e. the mapping from windspeed to NBRCS. The error magnitude is therefore sensitive to windspeed, as the GMF inversion is strongly sensitive to windspeed. This behavior is illustrated for the fully developed seas (FDS) GMF in Figure 5.1.

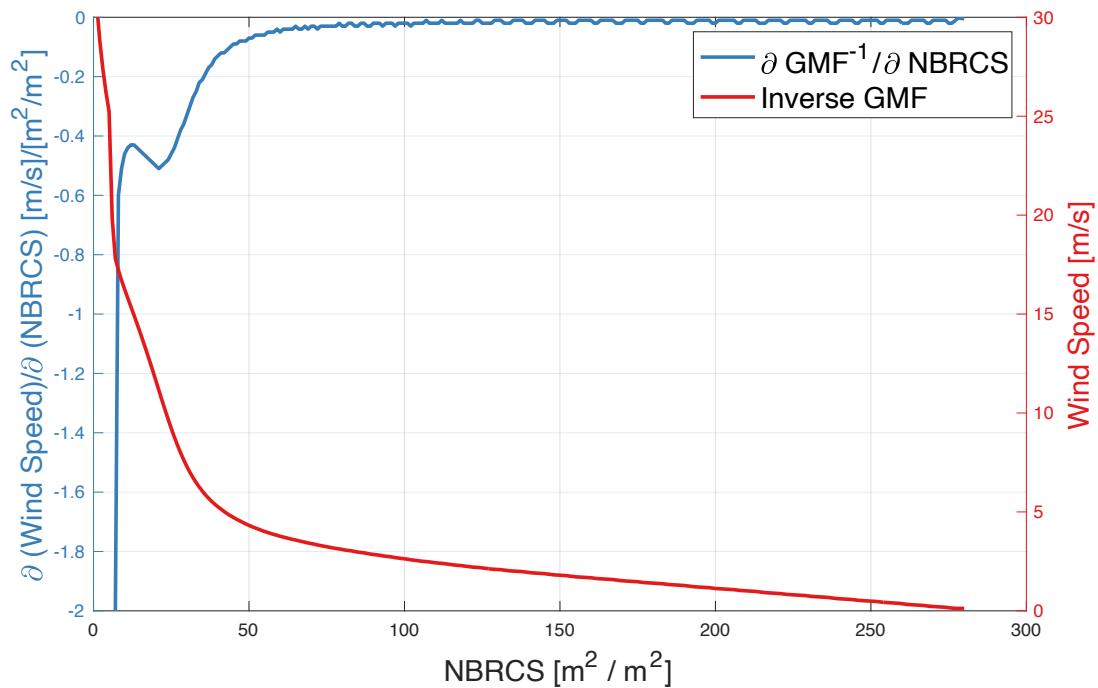


Figure 5.1 – The inverse GMF (red) and its partial derivative with respect to NBRCS (blue) for a fixed incidence angle $\theta = 30$ deg. Note that as NBRCS approaches 0, the retrieved windspeed increases substantially, as does its sensitivity to errors in NBRCS. At higher NBRCS, the sensitivity to errors approaches 0. The FDS GMF is valid for windspeeds < 25 m/s.

The instrument error model $R_{inst}(i, j)$ can be calculated as a function of windspeed and incidence angle, as shown in Figure 5.2. The error covariance is computed in terms of temporal distance τ between two observations i and j within the same track, and then mapped to great circle distance Δ using (5.8), such that R_{inst} takes the form $R_{inst}(i, \Delta)$, where the magnitude of the error depends on the windspeed and incidence angle at sample i , and the decorrelation rolloff is mapped to great circle distance Δ .




Figure 5.2 – Contribution to error covariance by propagating instrument error sources into observation space using the GMF. The colors correspond to different windspeeds: 5 m/s (red), 7 m/s (blue), and 15 m/s (green). The line styles correspond to observation incidence angle: 10 degrees (solid), 30 degrees (dashed), 45 degrees (dotted).

5.3.4 Evaluation of Representation-Retrieval Errors

Direct evaluation of $R_{rr}(i, j)$ is challenging, as both the magnitude and decorrelation rolloff of the representation-retrieval error is not well known. However, we can estimate $R_{rr}(i, j)$ by direct calculation of $R(i, j)$ as in (5.2) according to

$$R_{rr}(i, j) = R(i, j) - R_{inst}(i, j) \quad (5.12)$$

where R_{rr} is the representation-retrieval covariance, For calculation of \mathbf{e}_{obs} , \mathbf{y}_{obs} is the CYGNSS retrieved fully-developed seas windspeed observation at a given point in time and space, and \mathbf{y}_{true} is the ERA5 reanalysis 10 m neutral winds matched up via a nearest-neighbor

algorithm in space and time. It is well-documented that using reanalysis winds as ground truth is suboptimal, as the reanalysis tends to underestimate high wind speeds and smear out natural variability (Wu et al. 2024), but since the CYGNSS GMF is trained on ERA5, it is *a priori* defined as truth.

$R(i, j)$ is calculated by taking the autocorrelation along individual tracks such that

$$R(i, j) = R(\Delta) = \sigma_{\epsilon, obs}^2 \cdot \rho_{\epsilon, obs}(\Delta) \quad (5.13)$$

where $\rho_{\epsilon, obs}(\Delta)$ is the single-track observation error autocorrelation at spatial lag Δ and $\sigma_{\epsilon, obs}$ is the standard deviation of the observation error. This construction necessarily assumes that the error is stationary and isotropic. Because we assume stationarity, we incorporate a large number of tracks and calculate a *bulk* observation correlation for M tracks:

$$\mathcal{R}_{obs}(\Delta) = \frac{1}{M(N - \tau)} \frac{1}{\sigma_{z_i} \sigma_{z_{i+\Delta}}} \sum_{i=1}^{N-\tau} (Z_i - \bar{Z})(Z_{i+\Delta} - \bar{Z}) \quad (5.14)$$

where $Z_i = \sum_{k=1}^M e_{obs_{i,k}}$, σ_{z_i} is the standard deviation for at an index i , and k is the index for track number. By substituting $\mathcal{R}(\Delta)$ for $\rho(\Delta)$, R_{rr} can be estimated as:

$$R_{rr}(\Delta) \cong \sigma_{\epsilon, obs}^2 \cdot \mathcal{R}_{obs}(\Delta) - R_{inst}(\Delta) \quad (5.15)$$

where all values are a function of spatial great circle distance and are isotropic in space.

5.3.5 Data Selection

A month of CYGNSS data was collected (1-31 JAN 2019) using version 3.1 of the Level 2 science data product (CYGNSS 2021). These observations were filtered to only include tracks with sufficiently long track lengths ($N > 250$ samples) to enable investigation of the spatial and temporal observation error relationships using trackwise autocorrelation. Additionally, CYGNSS observations were filtered to include only windspeeds between 5 and 25 m/s, which is the domain where the CYGNSS fully-developed seas retrieval is most representative. These data

were matched with ERA5 hourly reanalysis winds and wave model outputs using a nearest-neighbor algorithm (Hersbach et al. 2020). The ERA5 atmospheric product has grid spacing of 31 km, while the wave model parameters operate on a 0.36 deg grid. Both the atmospheric and wave parameters were pre-processed and interpolated for regular latitude-longitude gridding by the Copernicus Climate Data Store upon data access. With the filters applied, a total of 62497 valid tracks were matched, with global observations +/- 40 deg latitude, across all 8 CYGNSS satellites.

5.4 Partition of Representation-Retrieval Error

In this section, we introduce a methodology for partitioning the representation-retrieval error in a manner that suggests a plausible physical explanation. The processing consists of four main steps:

- i. Identify model variables that correlate with windspeeds and windspeed errors;
- ii. Parameterize windspeed error as a function of a selected model variable;
- iii. Correct the retrieved windspeed with the parameterized functions; and
- iv. Repeat this process using corrected windspeed.

By applying successive corrections using this strategy, we can approximate individual components of error. To mitigate the influence of cross-correlation between model variables, we use the corrected windspeed as the starting point for the successive correction. The explainable error covariance is then:

$$R_{explained,k}(\Delta) \cong R[\epsilon_k](\Delta) \quad (5.16)$$

where $R_{explained,k}$ is the explainable error covariance for a given correction k as a function of spatial lag Δ , ϵ_k is the vector of windspeed correction in observation space for a given correction

cycle k , and each $R[\mathbf{x}] = \sigma_x^2 \mathcal{R}_x$, where σ_x^2 is the variance of term x and \mathcal{R}_x is the bulk autocorrelation of x as in (5.14).

5.4.1 Selection of Correctable Variables

Because ERA5 contains a fully-coupled wind-wave model, there are a number of model variables that may modulate MSS, local windspeed output, or some combination of the two. In general, it is counterproductive to correct with variables that are closely related to the target observation, which in this case is local windspeed. This is because we end up “correcting” with another representation of the ground truth, which can decrease the independence of the observed windspeed. We have compiled a list of variables and correlated the matchups with CYGNSS FDS observed windspeed, ERA5 reanalysis windspeed, and the windspeed error e_{obs} . Table 5-1 displays an abbreviated list of 10 ERA5 model variables sorted by the absolute value of windspeed error correlation. This list was curated from an initial analysis of 20 variables to illustrate the considerations for selecting variables for this process.

Table 5-1 – Abbreviated list of ERA5 model variables matched up with the CYGNSS windspeed observations. These variables are sorted by the absolute magnitude of correlation with windspeed error e_{obs} . Also displayed is the correlation of the model matchup with CYGNSS observations and ERA5 model windspeed. Variables in *italics* were passed over in analysis because they are too strongly correlated with ERA5 windspeed. Variables in **bold** are used in this analysis because they are strongly correlated with windspeed error but are sufficiently decorrelated from ERA5 windspeed.

Model Variable (ERA5 variable Parameter ID)	Correlation with Windspeed Error e_{obs}	Correlation with CYGNSS Windspeed	Correlation with ERA5 Windspeed
<i>Coefficient of Drag w/ Waves (140233)</i>	-0.5310	0.5725	0.9331
Mean Wave Period (140232)	0.4609	0.1841	-0.1942
<i>Charnock (148)</i>	-0.4463	0.3324	0.6498
Wave Spectral Skewness (140207)	-0.3033	0.3897	0.5896
Sig. Wave Height (140229)	0.1583	0.7767	0.5782
Wave Spectral Peakedness (140254)	0.1521	0.1569	0.0228
Air Seat Temp Diff [(167) – (34)]	-0.1115	-0.2295	-0.1198
2m Temp (167)	-0.1020	-0.3768	-0.2602
Sea Surface Temp (34)	-0.0707	-0.3359	-0.2478
Surface Pressure (134)	0.0364	-0.1478	-0.1619

The variables listed in Table 5-1 are illustrative of a primary heuristic for determining optimal model variables using this methodology: we are looking primarily for variables that correlate *strongly* with windspeed error e_{obs} , but *not strongly* with ERA5 windspeed output. For the purposes of this analysis, we pass over any variable that has absolute value correlation coefficient with ERA5 windspeed greater than 0.6. This threshold is primarily designed to avoid using model variables that have nearly one-to-one mappings with CYGNSS ground truth. This would have the effect of correcting errors with ground truth and provide limited insight to the plausible physical basis of error sources. However, this method considers variables with a low

level of partial correlation with ground truth, provided it is sufficiently decorrelated that there may exist a plausible physical rationale for the error.

The first variable in Table 5-1, Coefficient of Drag with Wind Waves, is nearly completely correlated with ERA5 windspeed. Upon inspection, the ERA5 wave model calculates this value as a linear function of the neutral windspeed over the ocean, so this high correlation is to be expected. The third item, Charnock Parameter, accounts for the aerodynamic roughness of the ocean surface, and depends directly on the sea state via wind-induced stress. Because these variables are primarily proxies for ground truth, we do not consider them in our analysis.

The top two remaining variables that correlate most strongly with windspeed error are the Mean Wave Period (MWP) and Wave Spectral Skewness (WSS). Both describe general statistical properties of the ocean surface and are calculated from the 2-dimensional wave spectrum produced by the wave model. Note that another parameter is Significant Wave Height (SWH), which is a known error source in the CYGNSS retrieval (Pascual, Clarizia, and Ruf 2021). The version of the CYGNSS windspeed product used for this analysis includes a correction for SWH, but there remains some residual correlation with windspeed error.

5.4.2 Generating Windspeed Correction

We will start with MWP, which is the average time between two wave crests, inclusive of both wind waves and swell. This variable is strongly correlated with windspeed error but only moderately correlated with ERA5 windspeed. However, because there is some residual dependence on ERA5 windspeed, we develop a correction using a two-step process:

1. Correct for windspeed dependence to generate model variable windspeed anomalies; and
2. Parameterize windspeed errors for variable windspeed anomalies with respect to the model.

Both steps can be visualized by depicting the two variables in a two-dimensional histogram, as shown in Figure 5.3. Because the MWP appears to be roughly linearly dependent on ERA5 windspeed, a linear regression is computed through the region of highest density ($N > 1e4$ counts) to minimize the impact of outliers, as shown in Fig. 5.3(a). This regression provides an expected MWP as a function of windspeed $\lambda_{exp}(\mathbf{y})$. The MWP anomaly is now:

$$\lambda_{anom}(\mathbf{y}) = \lambda_{exp}(\mathbf{y}) - \lambda_{obs} \quad (5.17)$$

where λ_{anom} is the MWP anomaly, λ_{exp} is the parameterized expected MWP from the empirical linear regression, and λ_{obs} is the observed MWP for the matchup data. The MWP anomaly is then matched with the observation error \mathbf{e}_{obs} as shown in Fig. 5.3(b). A similar regression analysis can be performed to create an estimator for windspeed error as a function of MWP anomaly $\mathbf{e}_{obs,\lambda}(\lambda_{anom})$, which is shown in red.

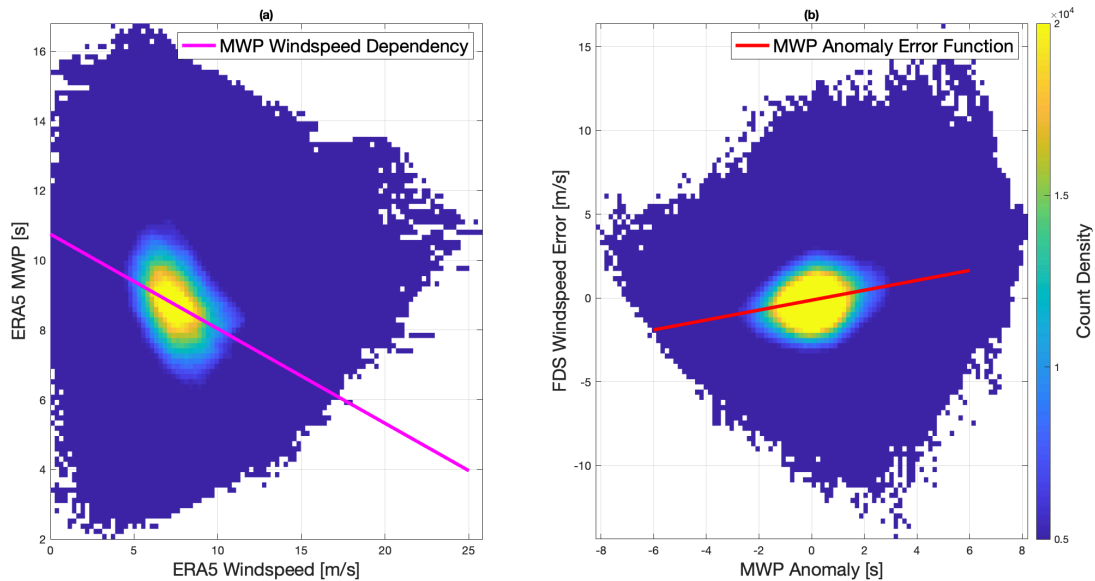


Figure 5.3 – (a) A 2-dimensional histogram of MWP versus ERA5 windspeed. The color shading of the bins represent the frequency of occurrence to illustrate the shape of the overall distribution of windspeed dependence on MWP. A linear regression is computed (magenta) through the region of highest count density ($N > 1e4$) to minimize the effects of outliers. This regression represents the windspeed dependence of MWP. The MWP anomaly is calculated as distance from the regression for a given windspeed. (b) A 2-dimensional histogram with windspeed error versus MWP anomaly (note the change in axes). A linear regression is computed (red) to estimate the impact of MWP anomaly on WS error.

The MWP-corrected windspeed can be now calculated as:

$$\mathbf{y}_\lambda = \mathbf{y} - \mathbf{e}_\lambda(\lambda_{anom}(\mathbf{y})) \quad (5.18)$$

where \mathbf{y}_λ is the MWP-corrected windspeed, \mathbf{y} is a windspeed vector, and $\mathbf{e}_\lambda(\lambda_{anom}(\mathbf{y}))$ is the MWP windspeed error, which is a function of MWP anomaly, which itself is a function of windspeed. This corrected windspeed is now used as the new basis for subsequent corrections.

The values in Table 5-1 can be updated to illustrate how the MWP correction impacts the correlations of other parameters. These new values are shown in Table 5-2.

Table 5-2 – List of ERA5 model variables matched up with the MWP-corrected windspeed. These variables are sorted by the absolute magnitude of correlation with MWP-corrected windspeed error $\mathbf{e}_{obs,\lambda}$. Also displayed is the correlation of the model matchup with MWP-corrected CYGNSS windspeed and MWP-corrected ERA5 windspeed. Variables in *italics* were passed over in analysis in the first step for being too strongly correlated with uncorrected ERA5 windspeed. Variables in **bold** are used in this analysis. Note that the MWP correction increases the WSS correlation with ERA5 windspeed, because MWP and WSS are related and the correction itself introduces additional windspeed dependence. For the purposes of this analysis, we assert that this does not violate the original rationale for excluding other parameters, because we know that WSS is not a robust proxy for ERA5 winds.

Model Variable (ERA5 variable name)	Correlation with MWP Corrected Windspeed Error	Correlation with MWP Corrected CYGNSS Windspeed	Correlation with MWP Corrected ERA5 Windspeed
<i>Coefficient of Drag w/ Waves (140233)</i>	-0.5943	0.6322	0.9173
<i>Charnock (148)</i>	-0.4877	0.3793	0.6449
Wave Spectral Skewness (140207)	-0.2433	0.5134	0.6527
Mean Wave Period (140232)	0.2091	-0.0592	-0.3886
Air Sea Temp Diff [(167) – (34)]	-0.1015	-0.2332	-0.112
Surface Pressure (134)	0.0605	-0.1463	-0.1499
Sig Wave Hght (140229)	-0.0583	0.656	0.4314
Wave Spectral Peakedness (140254)	0.0556	0.0748	-0.049
2m Temp (167)	-0.0315	-0.3453	-0.2134
Sea Surface Temp (34)	0.0058	-0.2985	-0.1977

A similar process occurs for the successive correction for WSS. WSS is a statistical property of the ocean surface and is a measure of deviation from Gaussianity. Operationally, this variable is used to forecast “freak” waves, and is calculated from third-order cumulants of the sea surface elevation in the wave model (ECMWF 2018). While one can expect WSS to have some physical coupling with surface windspeed, it can also be driven by other wave model parameters that impact the ocean surface elevation statistics. We calculate the WSS correction from the MSS-corrected ERA5 windspeed matchup as illustrated in Figure 5.4. The only differences from the prior iteration are that we started with the MWP-corrected windspeed and applied a quadratic fit (as opposed to linear) for the WSS anomaly error function based on the shape of the density plot. The parameters of both corrections are displayed in Table 5.3.




Figure 5.4 – (a) A 2-dimensional histogram with WSS versus MWP-corrected ERA5 windspeed. The color shading of the bins represents the frequency of occurrence to illustrate the shape of the overall distribution of (MWP-corrected) windspeed dependence on WSS. A linear regression is computed (magenta) through the region of highest count density ($N > 2e4$) to minimize the effects of outliers. This regression represents the windspeed dependence of WSS. The WSS anomaly is calculated as distance from the regression for a given windspeed. (b) A 2-dimensional histogram with MWP-corrected windspeed error versus WSS anomaly (note the change in axes). A quadratic regression is computed (red) to estimate the impact of WSS anomaly on windspeed error.

The WSS-corrected windspeed can be calculated in a similar fashion as (5.18):

$$\mathbf{y}_{\lambda,w} = \mathbf{y}_\lambda - \mathbf{e}_w(w_{anom}(\mathbf{y})) \quad (5.19)$$

where $\mathbf{y}_{\lambda,w}$ is the MWP- and WSS-corrected windspeed, \mathbf{y}_λ is the MWP-corrected windspeed, \mathbf{y} is a windspeed vector, and $\mathbf{e}_w(w_{anom}(\mathbf{y}))$ is the WSS windspeed error, which is a function of WSS anomaly, which itself is a function of windspeed.

Table 5-3 – List of regression parameters for the corrections applied in this section. Note that the regression coefficients c_i are listed in descending order, where c_i take the form $c_n x^n + c_{n-1} x^{n-1} + \dots$.

Abscissa	Ordinate	Type of Regression	Regression Coefficients
ERA5 Windspeed [m/s]	MWP [s]	Linear	[-0.2175,10.7503]
MWP Anomaly [s]	Windspeed Error [m/s]	Linear	[0.2957,-0.1286]
MWP-Corrected ERA5 Windspeed [m/s]	WSS [unitless]	Linear	[6.7731e-4, 0.0075]
WSS Anomaly [unitless]	Windspeed Error [m/s]	Quadratic	[-1.7584e3, 19.0555,0 .2402]

5.5 Results and Discussion

If the parametrized models for windspeed error as functions of MWP and WSS are applied as corrections to improve the windspeed retrieval (for the quality-controlled tracks at 5-25 m/s), the root-mean-square error decreases from 1.75 m/s with no corrections to 1.58 m/s with the MWP correction and further to 1.54 m/s with both the MWP and WSS corrections. This represents an overall improvement of 11% in the windspeed retrieval.

In addition, we can estimate the covariance structure for each error component individually and compare with the overall observation error covariance. The individual components of error covariance can be calculated via (17) and (19):

$$R_{\lambda}(\Delta) = \sigma_{\epsilon_{\lambda}}^2 \cdot \mathcal{R}[\epsilon_{\lambda}](\Delta) \quad (5.20a)$$

$$R_w(\Delta) = \sigma_{\epsilon_w}^2 \cdot \mathcal{R}[\epsilon_w](\Delta) \quad (5.20b)$$

where R_{λ} is the explainable error covariance due to MWP, R_w is the explainable error covariance due to WSS, $\sigma_{\epsilon_{\lambda}}$ is the standard deviation of the MWP correction in observation space, σ_{ϵ_w} is the standard deviation of the WSS correction in observation space, $\mathcal{R}[\epsilon_{\lambda}]$ is the bulk autocorrelation of the MWP correction in observation space as a function of spatial lag Δ , and $\mathcal{R}[\epsilon_w]$ is the bulk autocorrelation of the WSS correction in observation space as a function of spatial lag Δ .

The overall observation error covariance is calculated as suggested in (5.2) and (5.16):

$$R_o(\Delta) = R[\mathbf{y}_o - \mathbf{y}_{ERA}](\Delta) \quad (5.21a)$$

$$R_{corr}(\Delta) = R[\mathbf{y}_{\lambda,w} - \mathbf{y}_{ERA}](\Delta) \quad (5.21b)$$

where R_o is the initial observation error covariance as a function of spatial lag Δ , R_{corr} is the observation error covariance after the MWP and WSS corrections as a function of spatial lag Δ , \mathbf{y}_o is the uncorrected CYGNSS observed windspeed, $\mathbf{y}_{\lambda,w}$ is the MWP and WSS corrected CYGNSS observed windspeed, and \mathbf{y}_{ERA} is the ERA5 matched windspeed.

Figure 5.5 depicts the spatial error covariance for instrument error covariance R_{inst} as defined in (5.10a) for an exemplar track at the global average windspeed of 7 m/s and at 30 deg incidence, partitioned error covariances R_{λ} and R_w , and the overall observation error covariances R_o and R_{corr} .

Several key features can be identified from Fig. 5.5. First, the instrument-derived error covariance is generally a small component of the observation error covariance and decays with a

spatial scale size on the order of ~ 100 km. Second, the error covariance from MWP decays more slowly, while the WSS partition of observation error covariance is minimal at spatial lags greater than 500 km. These differences in characteristic scale sizes and rolloffs provide useful insights for future investigations of ocean surface roughness. Third, both the corrected and uncorrected topline windspeed error covariance decay to zero at ~ 1000 km.

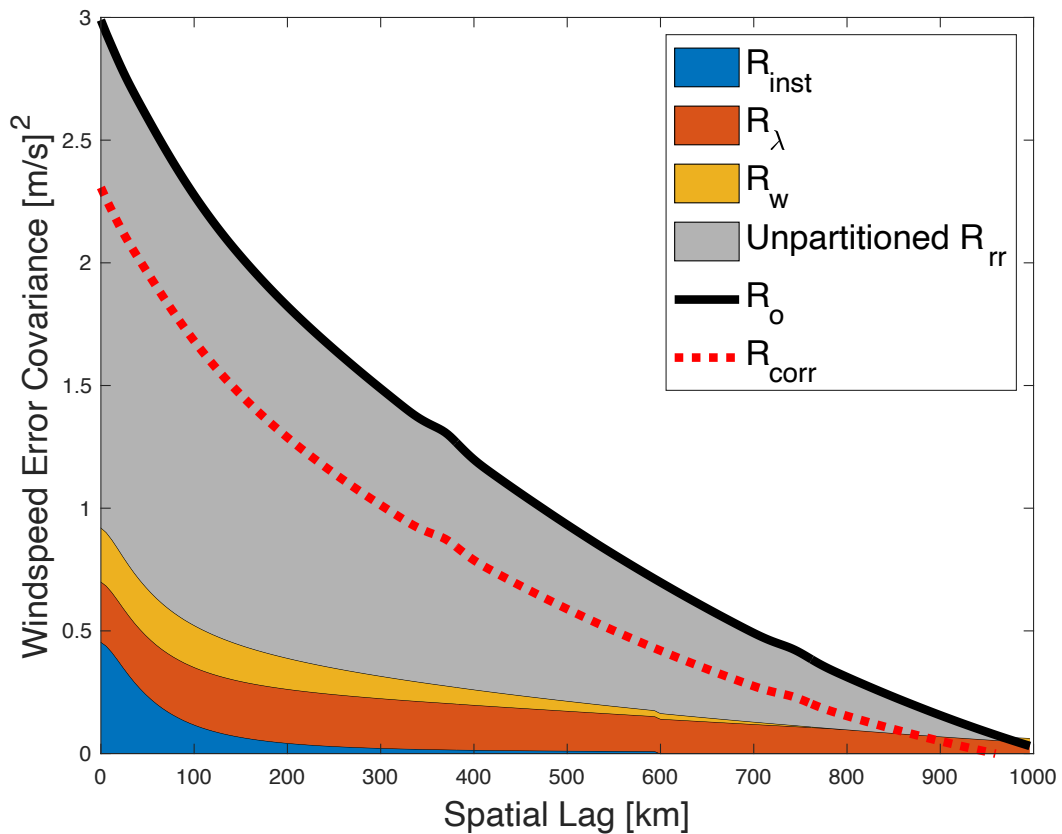


Figure 5.5 – A stacked area chart illustrating the error covariance partitions from this work. The blue region represents the observation error covariance from R_{inst} for a reference track at global average windspeed. The orange region represents the observation error covariance due to MWP. The yellow region represents the observation error covariance due to WSS. The three colored regions are defined in the model as independent and are plotted by stacking the contributions upon one another. The grey region is defined as undiagnosed representation-retrieval error in this model, and is constructed as the difference between R_o and the sum of the three defined error covariance contributions. R_o is traced in black and the corrected error covariance R_{corr} is dotted in red, demonstrating the across-the-board improvement in observation error covariance due to the corrections.

5.5.1 Validation

A key test of validation is determining whether the corresponding improvement in the corrected observation error covariance is due to the specific partitions identified. To calculate the impact of the corrections on overall error covariance structure, the cross correlation between R_o and R_{corr} needs to be accounted for. We asserted earlier that the error partitions are independent, but R_o and R_{corr} are expected to be highly correlated. The difference in error covariances becomes:

$$R_{o-corr}(\Delta) = R[(\mathbf{y}_o - \mathbf{y}_{ERA}) - (\mathbf{y}_{\lambda,w} - \mathbf{y}_{ERA})](\Delta) \quad (5.22a)$$

$$= R[\mathbf{y}_o - \mathbf{y}_{\lambda,w}](\Delta) \quad (5.22b)$$

$$= \sigma_{y_o - y_{\lambda,w}}^2 \cdot \mathcal{R}[\mathbf{y}_o - \mathbf{y}_{\lambda,w}](\Delta) \quad (5.22c)$$

$$= \sigma_{y_o - y_{\lambda,w}}^2 \cdot \mathcal{R}[\boldsymbol{\epsilon}_\lambda + \boldsymbol{\epsilon}_w](\Delta) \quad (5.22d)$$

where R_{o-corr} is the difference in covariance between the initial and corrected windspeed errors, and the variance of the difference in windspeeds $\sigma_{y_o - y_{\lambda,w}}^2$ is computed by

$$\sigma_{y_o - y_{\lambda,w}}^2 = \sigma_{y_o}^2 + \sigma_{y_{\lambda,w}}^2 - 2\rho_{o-c}(\sigma_{y_o} \sigma_{y_{\lambda,w}}) \quad (5.23)$$

In (5.23), σ_{y_o} is the standard deviation of uncorrected windspeed, $\sigma_{y_{\lambda,w}}$ is the standard deviation of the corrected windspeed, and ρ_{o-c} is the correlation coefficient between the uncorrected and corrected windspeeds. The error covariance difference $R_{o-corr}(\Delta)$ is plotted together with R_λ and R_w in Figure 5.6.

Figure 5.6 suggests that the improvement in windspeed error due to the corrections, and the resulting impacts to the shape of the observation error covariance, nearly perfectly corresponds with the sum of the individual error covariance partitions. Further, this plot also suggests that the process of successive corrections mitigates enough residual correlation between error terms that the correction terms can be considered independent between one another.

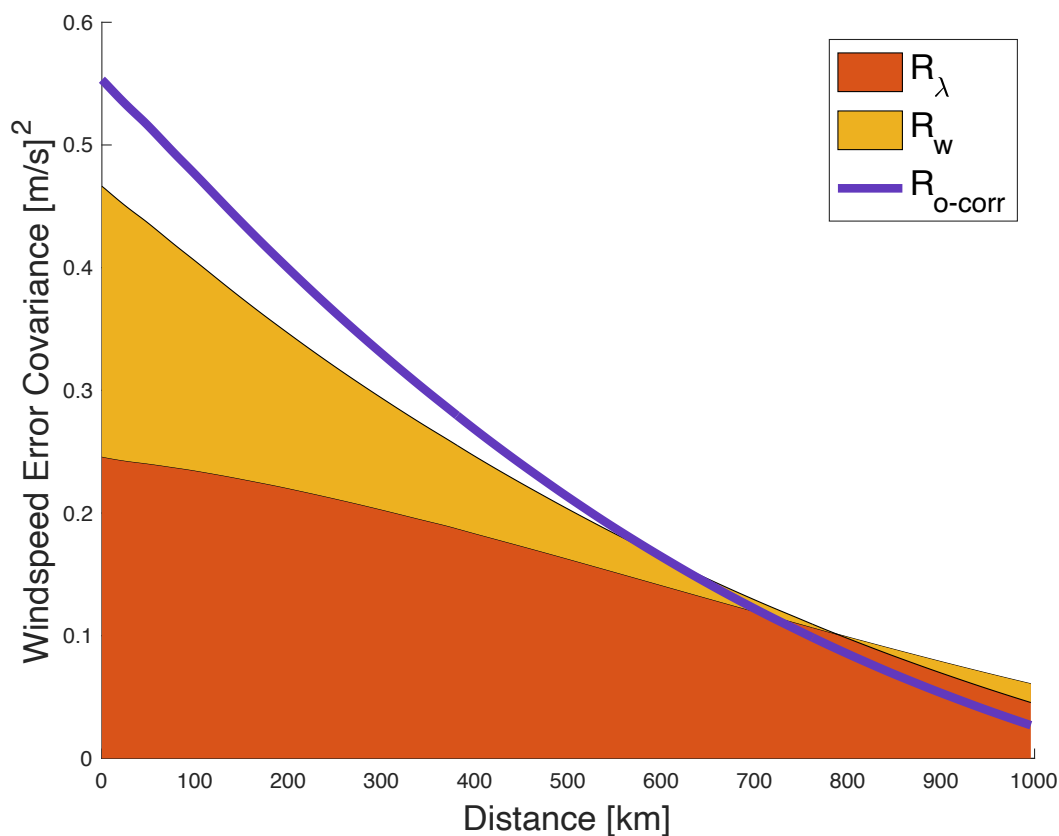


Figure 5.6 – A stacked area chart illustrating the validation test for the error partitions. The orange region represents the observation error covariance due to MWP. The yellow region represents the observation error covariance due to WSS. The error covariance due to the correction R_{o-corr} is plotted in purple. The overall shape and magnitude between the combined MWP and WSS error covariance and the calculated R_{o-corr} are reasonably consistent, which suggests that the retrieval improvement is the result of the correction of these specific error sources.

5.5.2 Limitations

As seen in Figure 5.5, a significant amount of representation-retrieval error remains undiagnosed. While this work can begin to partition certain components of e_{rr} , it is unlikely that this method will be able to completely partition observation error until the reference ground truth is matched to the observation appropriately in scale and representation. Matching disparate datasets for commensurate representation is a much larger topic that is outside the scope of this

work, but this limitation suggests an opportunity to begin quantifying the error structures of representation error.

This analysis also makes several assumptions about the isotropy and distribution of observation error. These assumptions are enabled by sufficiently large datasets to characterize broad statistical behavior of observation errors, but these relationships may change based on location, season, or another confounding variables.

Data Availability Statement.

All data used in this work are publicly available. The CYGNSS v3.1 science data record is available from the NASA Physical Oceanography Distributed Active Archive Center (PO.DAAC) (CYGNSS 2021). ERA5 Reanalysis is publicly available through the Copernicus Climate Data Store (Copernicus Climate Change Service (C3S) Climate Data Store (CDS) 2023)

Chapter 6 Conclusions

6.1 Contributions

The overall objective of this dissertation is to improve the return on investment of space-based Earth remote sensing systems, with particular focus on applications to GNSS-R technology. GNSS-R itself represents a groundbreaking new method to exploit signals-of-opportunity to extract information about the Earth system. As an opportunistic measurement, GNSS-R enables global measurements of various Earth surface parameters – windspeed, soil moisture, inundation, and others – in a form factor and price point unachievable with other measurements. But as a relatively newer technique with unique sampling characteristics, there are structural barriers to fully realizing the value GNSS-R offers to many observation and prediction applications.

There are four primary contributions in this dissertation, each with the aim of enhancing CYGNSS's value, and in many cases, the enhancements are generalizable to all current and future observing systems.

Chapter 2 presents an analysis of CYGNSS's thermally-dependent gain calibration errors stemming from a suboptimal pre-launch calibration sequencing estimate. This work improves the science operations duty cycle from 90% to 98%, vastly increasing the amount of science-quality data collected over the life of the mission. If the value of the data is equivalent to CYGNSS's approximately \$150M lifecycle cost, simply recovering this data produced ~\$12M in value return.

Chapter 3 presents a technique for calculating the representativity error for when observations are not exactly collocated or simultaneous. This will become more important as new satellite architectures come online and observations are disaggregated from single platforms. For many phenomena, matchup conditions of space and time were determined arbitrarily, or by simply the limitation of the available data on-hand. However, with new miniaturized sensors and proliferated constellations, space system architects can design use the technique presented in this chapter to design constellations to meet specific representation requirements.

Chapter 4 develops and validates a novel first-principle based instrument error correlation model for CYGNSS. This model has four main innovations: (1) it provides physically explainable reasons why observation error is correlated in space in time, which enables optimal weighting and assimilation in numerical weather prediction; (2) the model construction is generalizable, enabling the calculation of dynamic “bottom-up” instrument error covariance for other observing systems and techniques; (3) the validation suggests that errors between tracks of observations are reasonably independent; and (4) the development of a the “Filter Lemma” provides a theoretical basis for how previously uncorrelated errors can become correlated.

Finally, Chapter 5 propagates the instrument error model into observation space for the CYGNSS windspeed product, and develops a method for partitioning retrieval-representation error covariance for GNSS-R. This work has three immediate applications: (1) improvement of the CYGNSS FDS windspeed product with the inclusion of two ancillary wave parameters (wave spectral skewness and mean wave period), which reduced RMSE by 11% against ERA5 ground truth; (2) exploration of confounding geophysical causes and modulators of sea surface roughness, enabling future physical oceanography investigations; and (3) construction of

observation error covariance matrices for optimal data assimilation in numerical weather prediction.

6.2 Future Work

This dissertation also suggests a large body of natural follow-on investigations in three major thrusts:

1. Data assimilation: evaluating how weather model skill enhances when observation error covariance is specified.
2. Physical and biogeochemical oceanography: exploring the physical behavior of how ocean surface roughness is modulated by sources other than locally driven winds, such as surface contaminants, long-range swell, and stability condition.
3. Representation Error Theory: development of mathematical methods for characterizing, untangling, and reducing representation error.

Each of these activities adds value to both CYGNSS but also a host of other current and future observing systems. Data assimilation is a ripe target for future investigations, as small investments in extracting more value from observations can significantly improve the utility and usability of the data collected. Since collecting observations is almost always the most significant cost driver in the value chain, enhancements in data usability and utility represent some of the cheapest opportunities to improve model skill.

At the same time, a robust method to measure the correlated error structure in GNSS-R allows for deeper investigations of error terms or nuisance parameters in the retrieval, which may recover novel science findings in oceanography. For instance, the fundamental scales and

structures of surface biomass such as algae and sargassum may be an error term in the ocean windspeed product.

Finally, observation error covariance was examined closely for instrumentation error and likely sources of retrieval error, but there remains a significant opportunity to enhance the body of work regarding representation error covariance. Chapter 2 provides a first step in pricing the representativity error in terms of isotropic statistical estimates from high-quality estimates of variability of geophysical parameters, but the larger *representation* problem remains an opportunity for significant future theoretical development.

6.3 Lessons learned for future GNSS-R missions

This dissertation provides a few opportunities for lessons learned for future GNSS-R investigations, and some generalizable to other observing systems. The premise of Chapter 2 originates from a prelaunch oversight – the clock that governs the receiver load switching between the onboard blackbody calibration target and the science antenna were not synchronized, leading to a much-degraded science collection duty cycle. However, the analysis also presents another lesson learned for future missions, which is that inexpensive temperature sensors located at various places externally and internally can be extremely useful for characterizing how the temperature cycling impacts science measurements.

In fact, several other innovations by colleagues at the University of Michigan have been in response to invalid pre-launch assumptions. Through no fault of the CYGNSS team, there was a design assumption that GPS transmitters are relatively stable in time and across block types. Unfortunately, the recent GPS generations feature a new dynamic signal power that changes the effective isotropic radiated power throughout its orbit. The “flex power” events are not pre-

announced and required the development of a real-time dynamic power measurement technique to account for these fluctuations (Wang, Ruf, Gleason, et al. 2021b).

Additionally, CYGNSS was architected to capture measurements of local windspeeds near tropical cyclones to enhance cyclone initialization and forecasting in weather models. However, because the GNSS-R technique has less signal-to-noise at higher windspeeds, the retrieval degrades in tropical cyclone conditions. Simple enhancements in nadir gain can significantly improve the SNR and also the retrieval (Balasubramaniam and Ruf 2023).

Finally, the work presented in Chapter 4 suggests that careful documentation of prelaunch and operational processing can be useful guides for identifying sources of correlated error. In particular, the smoothing of CYGNSS's measured antenna gain patterns are a likely source of correlated error structure, and the specific details of these types of science data processing usually resides in unpublished technical memos and science team notes. There is residual value for all current and future missions to pay particular attention to operations that may correlate errors, such as smoothing or filtering, and documenting these operations in a systematic and accessible fashion.

References

- Arnold, Charles P, and Clifford H Dey. 1986. "Observing-Systems Simulation Experiments: Past, Present, and Future" 67 (6). [https://doi.org/10.1175/1520-0477\(1986\)067<0687:OSSEPP>2.0.CO;2](https://doi.org/10.1175/1520-0477(1986)067<0687:OSSEPP>2.0.CO;2).
- Balasubramaniam, Rajeswari, and Christopher S. Ruf. 2023. "Development and Application of a GNSS-R Error Model for Hurricane Winds." *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 1–12. <https://doi.org/10.1109/JSTARS.2023.3344371>.
- Bannister, R. N. 2008a. "A Review of Forecast Error Covariance Statistics in Atmospheric Variational Data Assimilation. I: Characteristics and Measurements of Forecast Error Covariances." *Quarterly Journal of the Royal Meteorological Society* 134 (637): 1951–70. <https://doi.org/10.1002/qj.339>.
- . 2008b. "A Review of Forecast Error Covariance Statistics in Atmospheric Variational Data Assimilation. II: Modelling the Forecast Error Covariance Statistics." *Quarterly Journal of the Royal Meteorological Society* 134 (637): 1971–96. <https://doi.org/10.1002/qj.340>.
- Bauer, Peter, Roberto Buizza, Carla Cardinali, and Jean Noël Thépaut. 2011. "Impact of Singular-Vector-Based Satellite Data Thinning on NWP." *Quarterly Journal of the Royal Meteorological Society* 137 (655): 286–302. <https://doi.org/10.1002/qj.733>.
- Berg, Wesley, Shannon T. Brown, Boon H. Lim, Steven C. Reising, Yuriy Goncharenko, Christian D. Kummerow, Todd C. Gaier, and Sharmila Padmanabhan. 2021. "Calibration and Validation of the TEMPEST-D CubeSat Radiometer." *IEEE Transactions on Geoscience and Remote Sensing* 59 (6): 4904–14. <https://doi.org/10.1109/TGRS.2020.3018999>.
- Bormann, Niels, and Peter Bauer. 2010. "Estimates of Spatial and Interchannel Observation-Error Characteristics for Current Sounder Radiances for Numerical Weather Prediction. I: Methods and Application to ATOVS Data." *Quarterly Journal of the Royal Meteorological Society* 136 (649): 1036–50. <https://doi.org/10.1002/qj.616>.
- Bormann, Niels, Massimo Bonavita, Rossana Dragani, Reima Eresmaa, Marco Matricardi, and Anthony McNally. 2016. "Enhancing the Impact of IASI Observations through an Updated Observation-error Covariance Matrix." *Quarterly Journal of the Royal Meteorological Society* 142 (697): 1767–80. <https://doi.org/10.1002/qj.2774>.
- Bretherton, Francis P., Russ E. Davis, and C.B. Fandry. 1976. "A Technique for Objective Analysis and Design of Oceanographic Experiments Applied to MODE-73." *Deep Sea Research and Oceanographic Abstracts* 23 (7): 559–82. [https://doi.org/10.1016/0011-7471\(76\)90001-2](https://doi.org/10.1016/0011-7471(76)90001-2).

- Bussy-Virat, Charles D., Aaron J. Ridley, Abhay Masher, Kyle Nave, and Marissa Intelisano. 2019. "Assessment of the Differential Drag Maneuver Operations on the CYGNSS Constellation." *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 12 (1): 7–15. <https://doi.org/10.1109/JSTARS.2018.2878158>.
- Campbell, William F., Elizabeth A. Satterfield, Benjamin Ruston, and Nancy L. Baker. 2017. "Accounting for Correlated Observation Error in a Dual-Formulation 4D Variational Data Assimilation System." *Monthly Weather Review* 145 (3): 1019–32. <https://doi.org/10.1175/MWR-D-16-0240.1>.
- Carreno-Luengo, H., S. T. Lowe, C. Zuffada, S. Esterhuizen, and S. Oveisgharan. 2017. "GNSS-R from the SMAP and CyGNSS Missions: Application to Polarimetric Scatterometry and Ocean Altimetry." In *International Geoscience and Remote Sensing Symposium (IGARSS)*. <https://doi.org/10.1109/IGARSS.2017.8128130>.
- Carreno-Luengo, Hugo, and Christopher S. Ruf. 2022. "Retrieving Freeze/Thaw Surface State From CYGNSS Measurements." *IEEE Transactions on Geoscience and Remote Sensing* 60: 1–13. <https://doi.org/10.1109/TGRS.2021.3120932>.
- Chew, C. C., and E. E. Small. 2018. "Soil Moisture Sensing Using Spaceborne GNSS Reflections: Comparison of CYGNSS Reflectivity to SMAP Soil Moisture." *Geophysical Research Letters* 45 (9): 4049–57. <https://doi.org/10.1029/2018GL077905>.
- Chew, Clara, John T. Reager, and Eric Small. 2018. "CYGNSS Data Map Flood Inundation during the 2017 Atlantic Hurricane Season." *Scientific Reports* 8 (1). <https://doi.org/10.1038/s41598-018-27673-x>.
- Christophersen, Hui W., Brittany A. Dahl, Jason P. Dunion, Robert F. Rogers, Frank D. Marks, Robert Atlas, and William J. Blackwell. 2021. "Impact of TROPICS Radiances on Tropical Cyclone Prediction in an OSSE." *Monthly Weather Review*, May. <https://doi.org/10.1175/MWR-D-20-0339.1>.
- Chu, Peter C., Wang Guihua, and Yuchun Chen. 2002. "Japan Sea Thermohaline Structure and Circulation. Part III: Autocorrelation Functions." *Journal of Physical Oceanography* 32 (12): 3596–3615. [https://doi.org/10.1175/1520-0485\(2002\)032<3596:JSTSAC>2.0.CO;2](https://doi.org/10.1175/1520-0485(2002)032<3596:JSTSAC>2.0.CO;2).
- Clarizia, Maria Paola, and Christopher S. Ruf. 2016a. "On the Spatial Resolution of GNSS Reflectometry." *IEEE Geoscience and Remote Sensing Letters* 13 (8): 1064–68. <https://doi.org/10.1109/LGRS.2016.2565380>.
- . 2016b. "Wind Speed Retrieval Algorithm for the Cyclone Global Navigation Satellite System (CYGNSS) Mission." *IEEE Transactions on Geoscience and Remote Sensing* 54 (8): 4419–32. <https://doi.org/10.1109/TGRS.2016.2541343>.
- Colosi, John, and Tim P. Barnett. 1990. "The Characteristic Spatial and Temporal Scales for SLP, SST, and Air Temperature in the Southern Hemisphere." *Journal of Applied Meteorology* 29 (8): 694–703. [https://doi.org/10.1175/1520-0450\(1990\)029<0694:TCSATS>2.0.CO;2](https://doi.org/10.1175/1520-0450(1990)029<0694:TCSATS>2.0.CO;2).
- Copernicus Climate Change Service (C3S) Climate Data Store (CDS). 2023. "Copernicus Climate Change Service (2023): ERA5 Hourly Data on Single Levels from 1940 to Present." <https://doi.org/10.24381/cds.adbb2d47>.
- Cordoba, M., S. L. Dance, G. A. Kelly, N. K. Nichols, and J. A. Waller. 2017. "Diagnosing Atmospheric Motion Vector Observation Errors for an Operational High-resolution Data Assimilation System." *Quarterly Journal of the Royal Meteorological Society* 143 (702): 333–41. <https://doi.org/10.1002/qj.2925>.

- CYGNSS. 2020a. “CYGNSS Level 1 Science Data Record Version 3.0.” PO.DAAC. <https://doi.org/10.5067/CYGNSS-L1X30>.
- . 2020b. “CYGNSS Level 2 Science Data Record Version 3.0.” PO.DAAC. <https://doi.org/10.5067/CYGNSS-L2X30>.
- . 2021. “CYGNSS Level 1 Science Data Record Version 3.1.” NASA Physical Oceanography Distributed Active Archive Center. <https://doi.org/10.5067/CYGNSS-L1X31>.
- Daley, Roger. 1992. “The Effect of Serially Correlated Observation and Model Error on Atmospheric Data Assimilation.” *Monthly Weather Review* 120 (1): 164–77. [https://doi.org/10.1175/1520-0493\(1992\)120<0164:TEOSCO>2.0.CO;2](https://doi.org/10.1175/1520-0493(1992)120<0164:TEOSCO>2.0.CO;2).
- Dee, Dick P., and Arlindo M. Da Silva. 1999. “Maximum-Likelihood Estimation of Forecast and Observation Error Covariance Parameters. Part I: Methodology.” *Monthly Weather Review* 127 (8): 1822–34. [https://doi.org/10.1175/1520-0493\(1999\)127<1822:MLEOFA>2.0.CO;2](https://doi.org/10.1175/1520-0493(1999)127<1822:MLEOFA>2.0.CO;2).
- Delcroix, Thierry, Michael J. McPhaden, Alain Dessier, and Yves Gouriou. 2005. “Time and Space Scales for Sea Surface Salinity in the Tropical Oceans.” *Deep Sea Research Part I: Oceanographic Research Papers* 52 (5): 787–813. <https://doi.org/10.1016/j.dsr.2004.11.012>.
- Desroziers, G., L. Berre, B. Chapnik, and P. Poli. 2005. “Diagnosis of Observation, Background and Analysis-error Statistics in Observation Space.” *Quarterly Journal of the Royal Meteorological Society* 131 (613): 3385–96. <https://doi.org/10.1256/qj.05.108>.
- Donlon, Craig J., Robert Cullen, Luisella Giulicchi, Pierrick Vuilleumier, C. Richard Francis, Mieke Kuschnerus, William Simpson, et al. 2021. “The Copernicus Sentinel-6 Mission: Enhanced Continuity of Satellite Sea Level Measurements from Space.” *Remote Sensing of Environment* 258 (June): 112395. <https://doi.org/10.1016/j.rse.2021.112395>.
- ECMWF. 2018. “IFS Documentation Part VII: ECMWF Wave Model.” Cy45r1.
- Eden, Carsten. 2007. “Eddy Length Scales in the North Atlantic Ocean.” *Journal of Geophysical Research* 112 (C6): C06004. <https://doi.org/10.1029/2006JC003901>.
- Evans, Madeline C., and Christopher S. Ruf. 2022. “Toward the Detection and Imaging of Ocean Microplastics With a Spaceborne Radar.” *IEEE Transactions on Geoscience and Remote Sensing* 60: 1–9. <https://doi.org/10.1109/TGRS.2021.3081691>.
- Fowler, A. M., S. L. Dance, and J. A. Waller. 2018. “On the Interaction of Observation and Prior Error Correlations in Data Assimilation.” *Quarterly Journal of the Royal Meteorological Society* 144 (710): 48–62. <https://doi.org/10.1002/qj.3183>.
- Gandin, L.S. 1965. *Objective analysis of meteorological fields (Ob'ektivnyi analiz meteorologicheskikh polei)*. Jerusalem: Israel Program for Scientific Translations.
- Gao, Yudong, Hui Xiao, Dehai Jiang, Qilin Wan, Pak Wai Chan, Kai Kwong Hon, and Guo Deng. 2019. “Impacts of Thinning Aircraft Observations on Data Assimilation and Its Prediction during Typhoon Nida (2016).” *Atmosphere* 10 (12): 754. <https://doi.org/10.3390/ATMOS10120754>.
- Gelaro, Ronald, Will McCarty, Max J. Suárez, Ricardo Todling, Andrea Molod, Lawrence Takacs, Cynthia A. Randles, et al. 2017. “The Modern-Era Retrospective Analysis for Research and Applications, Version 2 (MERRA-2).” *Journal of Climate* 30 (14): 5419–54. <https://doi.org/10.1175/JCLI-D-16-0758.1>.
- Gerlein-Safdi, Cynthia, A. Anthony Bloom, Genevieve Plant, Eric A. Kort, and Christopher S. Ruf. 2021. “Improving Representation of Tropical Wetland Methane Emissions With

- CYGNSS Inundation Maps.” *Global Biogeochemical Cycles* 35 (12): e2020GB006890. <https://doi.org/10.1029/2020GB006890>.
- Gille, Sarah T. 2005. “Statistical Characterization of Zonal and Meridional Ocean Wind Stress.” *Journal of Atmospheric and Oceanic Technology* 22 (9): 1353–72. <https://doi.org/10.1175/JTECH1789.1>.
- Gille, Sarah T., and Kathryn A. Kelly. 1996. “Scales of Spatial and Temporal Variability in the Southern Ocean.” *Journal of Geophysical Research: Oceans* 101 (C4): 8759–73. <https://doi.org/10.1029/96JC00203>.
- Gleason, Scott. 2014. “Cyclone Global Navigation Satellite System (CYGNSS) Algorithm Theoretical Basis Document Level 1A DDM Calibration.”
- . 2018. “Level 1A DDM Calibration Algorithm Theoretical Basis Document (Rev. 2).” Algorithm Theoretical Basis Document 148–0136. University of Michigan. https://cygnss.engin.umich.edu/wp-content/uploads/sites/534/2021/07/148-0136_ATBD_L1A_DDMCalibration_Rev2_Aug2018_release.pdf.
- . 2020. “Level 1B DDM Calibration Algorithm Theoretical Basis Document (Rev. 3).” Algorithm Theoretical Basis Document 148–0137. University of Michigan. https://cygnss.engin.umich.edu/wp-content/uploads/sites/534/2021/07/148-0137_ATBD-L1B-DDM-Calibration_R3_release.pdf.
- Gleason, Scott, Christopher S. Ruf, Maria Paola Clarizia, and Andrew J. O’Brien. 2016. “Calibration and Unwrapping of the Normalized Scattering Cross Section for the Cyclone Global Navigation Satellite System.” *IEEE Transactions on Geoscience and Remote Sensing* 54 (5): 2495–2509. <https://doi.org/10.1109/TGRS.2015.2502245>.
- Gleason, Scott, Christopher S. Ruf, Andrew J. O’Brien, and Darren S. McKague. 2019. “The CYGNSS Level 1 Calibration Algorithm and Error Analysis Based on On-Orbit Measurements.” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 12 (1): 37–49. <https://doi.org/10.1109/JSTARS.2018.2832981>.
- Global Modeling and Assimilation Office. 2015. “MERRA-2 inst1_2d_asm_Nx: 2d,3-Hourly,Instantaneous,Single-Level,Assimilation,Single-Level Diagnostics V5.12.4.” NASA Goddard Earth Sciences Data and Information Services Center. <https://doi.org/10.5067/3Z173KIE2TPD>.
- Grey, John Petrus, Ian R. Mann, Michael D. Fleischauer, and Duncan G. Elliott. 2020. “Analytic Model for Low Earth Orbit Satellite Solar Power.” *IEEE Transactions on Aerospace and Electronic Systems* 56 (5): 3349–59. <https://doi.org/10.1109/TAES.2020.3009510>.
- Harrington, J.D. 2012. “NASA Selects Low Cost, High Science Earth Venture Space System.” NASA. <https://www.nasa.gov/news-release/nasa-selects-low-cost-high-science-earth-venture-space-system/>.
- Hersbach, H, B Bell, P Berrisford, G Biavati, A Horányi, J Muñoz Sabater, J Nicolas, et al. 2023. “ERA5 Hourly Data on Single Levels from 1940 to Present.” Copernicus Climate Change Service (C3S) Climate Data Store (CDS). <https://doi.org/10.24381/cds.adbb2d47>.
- Hersbach, Hans, Bill Bell, Paul Berrisford, Shoji Hirahara, András Horányi, Joaquín Muñoz-Sabater, Julien Nicolas, et al. 2020. “The ERA5 Global Reanalysis.” *Quarterly Journal of the Royal Meteorological Society* 146 (730): 1999–2049. <https://doi.org/10.1002/qj.3803>.
- Hoffman, Ross N. 2018. “The Effect of Thinning and Superobservations in a Simple One-Dimensional Data Analysis with Mischaracterized Error.” *Monthly Weather Review* 146 (4): 1181–95. <https://doi.org/10.1175/MWR-D-17-0363.1>.

- Hoffman, Ross N., and Robert Atlas. 2016. "Future Observing System Simulation Experiments." *Bulletin of the American Meteorological Society* 97 (9): 1601–16. <https://doi.org/10.1175/BAMS-D-15-00200.1>.
- Janjić, T., N. Bormann, M. Bocquet, J. A. Carton, S. E. Cohn, S. L. Dance, S. N. Losa, et al. 2018. "On the Representation Error in Data Assimilation." *Quarterly Journal of the Royal Meteorological Society* 144 (713): 1257–78. <https://doi.org/10.1002/qj.3130>.
- Kim, Edward, Cheng-Hsuan J. Lyu, Kent Anderson, R. Vincent Leslie, and William J. Blackwell. 2014. "S-NPP ATMS Instrument Prelaunch and on-Orbit Performance Evaluation." *Journal of Geophysical Research: Atmospheres* 119 (9): 5653–70. <https://doi.org/10.1002/2013JD020483>.
- Kohler, Craig, Lex LeBlanc, and James Elliott. 2015. "SCOOP - NDBC's New Ocean Observing System." In *OCEANS 2015 - MTS/IEEE Washington*, 1–5. Washington, DC: IEEE. <https://doi.org/10.23919/OCEANS.2015.7401834>.
- Kroodsma, Rachael A., Darren S. McKague, and Christopher S. Ruf. 2012. "Inter-Calibration of Microwave Radiometers Using the Vicarious Cold Calibration Double Difference Method." *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 5 (3): 1006–13. <https://doi.org/10.1109/JSTARS.2012.2195773>.
- Kuragano, Tsurane, and Masafumi Kamachi. 2000. "Global Statistical Space-Time Scales of Oceanic Variability Estimated from the TOPEX/POSEIDON Altimeter Data." *Journal of Geophysical Research: Oceans* 105 (C1): 955–74. <https://doi.org/10.1029/1999JC900247>.
- Lazo, Jeffrey K., Megan Lawson, Peter H. Larsen, and Donald M. Waldman. 2011. "U.S. Economic Sensitivity to Weather Variability." *Bulletin of the American Meteorological Society* 92 (6): 709–20. <https://doi.org/10.1175/2011BAMS2928.1>.
- Li, Weiqiang, Estel Cardellach, Fran Fabra, Serni Ribo, and Antonio Rius. 2020. "Assessment of Spaceborne GNSS-R Ocean Altimetry Performance Using CYGNSS Mission Raw Data." *IEEE Transactions on Geoscience and Remote Sensing* 58 (1): 238–50. <https://doi.org/10.1109/TGRS.2019.2936108>.
- Li, Zhenglong, Jun Li, Timothy J. Schmit, Pei Wang, Agnes Lim, Jinlong Li, Fredrick W. Nagle, et al. 2019. "The Alternative of CubeSat-Based Advanced Infrared and Microwave Sounders for High Impact Weather Forecasting." *Atmospheric and Oceanic Science Letters* 12 (2): 80–90. <https://doi.org/10.1080/16742834.2019.1568816>.
- Li, Zhenglong, Jun Li, Pei Wang, Agnes Lim, Jinlong Li, Timothy J. Schmit, Robert Atlas, Sid-Ahmed Boukabara, and Ross N. Hoffman. 2018. "Value-Added Impact of Geostationary Hyperspectral Infrared Sounders on Local Severe Storm Forecasts—via a Quick Regional OSSE." *Advances in Atmospheric Sciences* 35 (10): 1217–30. <https://doi.org/10.1007/s00376-018-8036-3>.
- Lipiec, Eva. 2023. "National Oceanic and Atmospheric Administration (NOAA) FY2024 Budget Request and Appropriations." IF12406. In Focus. Washington, DC: Congressional Research Service. <https://crsreports.congress.gov/product/pdf/IF/IF12406>.
- Liu, Z. Q., and F. Rabier. 2002. "The Interaction between Model Resolution, Observation Resolution and Observation Density in Data Assimilation: A One-Dimensional Study." *Quarterly Journal of the Royal Meteorological Society* 128 (582): 1367–86. <https://doi.org/10.1256/003590002320373337>.
- Lorenc, A. C., S. P. Ballard, R. S. Bell, N. B. Ingleby, P. L. F. Andrews, D. M. Barker, J. R. Bray, et al. 2000. "The Met. Office Global Three-dimensional Variational Data

- Assimilation Scheme.” *Quarterly Journal of the Royal Meteorological Society* 126 (570): 2991–3012. <https://doi.org/10.1002/qj.49712657002>.
- Marcuccio, Salvo, Silvia Ullo, Marco Carminati, and Olfa Kanoun. 2019. “Smaller Satellites, Larger Constellations: Trends and Design Issues for Earth Observation Systems.” *IEEE Aerospace and Electronic Systems Magazine* 34 (10): 50–59. <https://doi.org/10.1109/MAES.2019.2928612>.
- Marquis, Willard A., and Daniel L. Reigh. 2015. “The GPS Block IIR and IIR-M Broadcast L-Band Antenna Panel: Its Pattern and Performance: GPS Block IIR and IIR-M L-Band Antenna Panel Pattern.” *Navigation* 62 (4): 329–47. <https://doi.org/10.1002/navi.123>.
- Mayers, David, and Christopher Ruf. 2019. “Tropical Cyclone Center Fix Using CYGNSS Winds.” *Journal of Applied Meteorology and Climatology* 58 (9): 1993–2003. <https://doi.org/10.1175/JAMC-D-19-0054.1>.
- McLean, Lorna M. 2010. “The Determination Of Ocean Correlation Scales Using Argo Float Data.” Southampton, UK: University of Southampton.
- Morgan, Daniel. 2024. “NASA Appropriations and Authorizations: A Fact Sheet.” R43419. Congressional Research Service.
- Morss, R. E., and K. A. Emanuel. 2002. “Influence of Added Observations on Analysis and Forecast Errors: Results from Idealized Systems.” *Quarterly Journal of the Royal Meteorological Society* 128 (579): 285–321. <https://doi.org/10.1256/00359000260498897>.
- Nag, Sreeja, Jacqueline LeMoigne, David W. Miller, and Olivier De Weck. 2015. “A Framework for Orbital Performance Evaluation in Distributed Space Missions for Earth Observation.” In *2015 IEEE Aerospace Conference*, 1–20. Big Sky, MT: IEEE. <https://doi.org/10.1109/AERO.2015.7119227>.
- National Academies of Sciences, Engineering, and Medicine. 2023. *Assessment of Commercial Space Platforms for Earth Science Instruments: Report Series-Committee on Earth Science and Applications from Space*. Washington, D.C.: National Academies Press. <https://doi.org/10.17226/27019>.
- National Data Buoy Center. 1971. “Meteorological and Oceanographic Data Collected from the National Data Buoy Center Coastal-Marine Automated Network (C-MAN) and Moored (Weather) Buoys.” NOAA National Centers for Environmental Information. <https://www.ncei.noaa.gov/archive/accession/NDBC-CMANWx>.
- Pascual, Daniel, Maria Paola Clarizia, and Christopher S. Ruf. 2021. “Improved CYGNSS Wind Speed Retrieval Using Significant Wave Height Correction.” *Remote Sensing* 13 (21): 4313. <https://doi.org/10.3390/rs13214313>.
- Powell, C. E., Christopher S. Ruf, Scott Gleason, and Scot C. R. Rafkin. 2024. “Sampled Together: Assessing the Value of Simultaneous Collocated Measurements for Optimal Satellite Configurations.” *Bulletin of the American Meteorological Society* 105 (1): E285–96. <https://doi.org/10.1175/BAMS-D-23-0198.1>.
- Powell, C. E., Christopher S. Ruf, Darren S. McKague, Tianlin Wang, and Anthony Russel. 2024. “An Instrument Error Correlation Model for Global Navigation Satellite System Reflectometry.” *Remote Sensing* 16 (5): 742. <https://doi.org/10.3390/rs16050742>.
- Powell, C. E., Christopher S. Ruf, and Anthony Russel. 2022. “An Improved Blackbody Calibration Cadence for CYGNSS.” *IEEE Transactions on Geoscience and Remote Sensing* 60: 1–7. <https://doi.org/10.1109/TGRS.2022.3165001>.

- Powell, C.E., and Christopher S. Ruf. 2024. "A Method for Estimating and Partitioning the Covariance of Windspeed Observation Error from GNSS Reflectometry." *Submitted to J. Atmos. Oceanic Tech, in Review*, March.
- Privé, N. C., Matthew McLinden, Bing Lin, Isaac Moradi, Meta Sienkiewicz, G. M. Heymsfield, and Will McCarty. 2023. "Impacts of Marine Surface Pressure Observations from a Spaceborne Differential Absorption Radar Investigated with an Observing System Simulation Experiment." *Journal of Atmospheric and Oceanic Technology* 40 (8): 897–918. <https://doi.org/10.1175/JTECH-D-22-0088.1>.
- Rainwater, Sabrina, Craig H. Bishop, and William F. Campbell. 2015. "The Benefits of Correlated Observation Errors for Small Scales." *Quarterly Journal of the Royal Meteorological Society* 141 (693): 3439–45. <https://doi.org/10.1002/qj.2582>.
- Reynolds, Richard W., and Thomas M. Smith. 1994. "Improved Global Sea Surface Temperature Analyses Using Optimum Interpolation." *Journal of Climate* 7 (6): 929–48. [https://doi.org/10.1175/1520-0442\(1994\)007<0929:IGSSTA>2.0.CO;2](https://doi.org/10.1175/1520-0442(1994)007<0929:IGSSTA>2.0.CO;2).
- Richardson, Lewis F. 1922. *Weather Prediction by Numerical Process*. 1st ed. Cambridge, UK.
- Romanou, Anastasia, William B. Rossow, and Shu-Hsien Chou. 2006. "Decorrelation Scales of High-Resolution Turbulent Fluxes at the Ocean Surface and a Method to Fill in Gaps in Satellite Data Products." *Journal of Climate* 19 (14): 3378–93. <https://doi.org/10.1175/JCLI3773.1>.
- Ruf, Christopher S., Robert Atlas, Paul S. Chang, Maria Paola Clarizia, James L. Garrison, Scott Gleason, Stephen J. Katzberg, et al. 2016. "New Ocean Winds Satellite Mission to Probe Hurricanes and Tropical Convection." *Bulletin of the American Meteorological Society* 97 (3): 385–95. <https://doi.org/10.1175/BAMS-D-14-00218.1>.
- Ruf, Christopher S., and Rajeswari Balasubramaniam. 2019. "Development of the CYGNSS Geophysical Model Function for Wind Speed." *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 12 (1): 66–77. <https://doi.org/10.1109/JSTARS.2018.2833075>.
- Ruf, Christopher, M.P. Clarizia, S. Gleason, Z. Jelenak, J. Murray, M. Morris, S. Musko, et al. 2016. *CYGNSS Handbook - Cyclone Global Navigation Satellite System*. NASA.
- Ruf, Christopher S., Scott Gleason, and Darren S. McKague. 2019. "Assessment of CYGNSS Wind Speed Retrieval Uncertainty." *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 12 (1): 87–97. <https://doi.org/10.1109/JSTARS.2018.2825948>.
- Ruf, Christopher, Darren McKague, Derek J. Posselt, Scott Gleason, Maria Paola Clarizia, Valery U Zavorotny, Tim Butler, et al. 2022. *CYGNSS Handbook*. 2nd ed. Michigan Publishing Services. <https://doi.org/10.3998/mpub.12741920>.
- Rutherford, Ian D. 1972. "Data Assimilation by Statistical Interpolation of Forecast Error Fields." *Journal of the Atmospheric Sciences* 29 (5): 809–15. [https://doi.org/10.1175/1520-0469\(1972\)029<0809:DABSIO>2.0.CO;2](https://doi.org/10.1175/1520-0469(1972)029<0809:DABSIO>2.0.CO;2).
- Schoeberl, M.R. 2002. "The Afternoon Constellation: A Formation of Earth Observing Systems for the Atmosphere and Hydrosphere." In *IEEE International Geoscience and Remote Sensing Symposium*, 1:354–56. Toronto, Ont., Canada: IEEE. <https://doi.org/10.1109/IGARSS.2002.1025038>.
- Schutgens, Nick A. J., Edward Gryspeerdt, Natalie Weigum, Svetlana Tsyro, Daisuke Goto, Michael Schulz, and Philip Stier. 2016. "Will a Perfect Model Agree with Perfect

- Observations? The Impact of Spatial Sampling.” *Atmospheric Chemistry and Physics* 16 (10): 6335–53. <https://doi.org/10.5194/acp-16-6335-2016>.
- Schutgens, Nick, Svetlana Tsyro, Edward Gryspeerd, Daisuke Goto, Natalie Weigum, Michael Schulz, and Philip Stier. 2017. “On the Spatio-Temporal Representativeness of Observations.” *Atmospheric Chemistry and Physics* 17 (16): 9761–80. <https://doi.org/10.5194/acp-17-9761-2017>.
- Semane, Nouredine, Richard Anthes, Jeremiah Sjoberg, Sean Healy, and Benjamin Ruston. 2022. “Comparison of Desroziers and Three-Cornered Hat Methods for Estimating COSMIC-2 Bending Angle Uncertainties.” *Journal of Atmospheric and Oceanic Technology* 39 (7): 929–39. <https://doi.org/10.1175/JTECH-D-21-0175.1>.
- Simonin, David, Joanne A. Waller, Susan P. Ballard, Sarah L. Dance, and Nancy K. Nichols. 2019. “A Pragmatic Strategy for Implementing Spatially Correlated Observation Errors in an Operational System: An Application to Doppler Radial Winds.” *Quarterly Journal of the Royal Meteorological Society* 145 (723): 2772–90. <https://doi.org/10.1002/qj.3592>.
- Sjoberg, Jeremiah P., Richard A. Anthes, and Therese Rieckh. 2021. “The Three-Cornered Hat Method for Estimating Error Variances of Three or More Atmospheric Datasets. Part I: Overview and Evaluation.” *Journal of Atmospheric and Oceanic Technology* 38 (3): 555–72. <https://doi.org/10.1175/JTECH-D-19-0217.1>.
- Smith, David, Samuel E. Hunt, Mireya Etxaluz, Dan Peters, Tim Nightingale, Jonathan Mittaz, Emma R. Woolliams, and Edward Polehampton. 2021. “Traceability of the Sentinel-3 SLSTR Level-1 Infrared Radiometric Processing.” *Remote Sensing* 13 (3): 374. <https://doi.org/10.3390/rs13030374>.
- St. Germain, Karen, Frank W. Gallagher, and W. Maier. 2018. “Opportunities Highlighted By the NSOSA Study.” In *IGARSS 2018 - 2018 IEEE International Geoscience and Remote Sensing Symposium*, 7387–90. Valencia: IEEE. <https://doi.org/10.1109/IGARSS.2018.8517710>.
- Steigenberger, Peter, Steffen Thölert, and Oliver Montenbruck. 2019. “Flex Power on GPS Block IIR-M and IIF.” *GPS Solutions* 23 (1): 1–12. <https://doi.org/10.1007/s10291-018-0797-8>.
- Stephens, Graeme L., Deborah G. Vane, Ronald J. Boain, Gerald G. Mace, Kenneth Sassen, Zhien Wang, Anthony J. Illingworth, et al. 2002. “The Cloudsat Mission and the A-Train: A New Dimension of Space-Based Observations of Clouds and Precipitation.” *Bulletin of the American Meteorological Society* 83 (12): 1771–90. <https://doi.org/10.1175/BAMS-83-12-1771>.
- Stewart, L. M., S. L. Dance, and N. K. Nichols. 2008. “Correlated Observation Errors in Data Assimilation.” *International Journal for Numerical Methods in Fluids* 56 (8): 1521–27. <https://doi.org/10.1002/flid.1636>.
- Stewart, L. M., S. L. Dance, N. K. Nichols, J. R. Eyre, and J. Cameron. 2014. “Estimating Interchannel Observation-error Correlations for IASI Radiance Data in the Met Office System.” *Quarterly Journal of the Royal Meteorological Society* 140 (681): 1236–44. <https://doi.org/10.1002/qj.2211>.
- Stewart, Laura M., Sarah L. Dance, and Nancy K. Nichols. 2013. “Data Assimilation with Correlated Observation Errors: Experiments with a 1-D Shallow Water Model.” *Tellus A: Dynamic Meteorology and Oceanography* 65 (1): 19546. <https://doi.org/10.3402/tellusa.v65i0.19546>.

- Tan, David G. H., Erik Andersson, Michael Fisher, and Lars Isaksen. 2007. "Observing-System Impact Assessment Using a Data Assimilation Ensemble Technique: Application to the ADM–Aeolus Wind Profiling Mission." *Quarterly Journal of the Royal Meteorological Society* 133 (623): 381–90. <https://doi.org/10.1002/qj.43>.
- Taylor, John R. 1997. *An Introduction to Error Analysis: The Study of Uncertainties in Physical Measurements*. 2nd ed. Sausalito, Calif: University Science Books.
- Ulaby, Fawwaz T., and David G. Long. 2014. *Microwave Radar and Radiometric Remote Sensing*. Ann Arbor, MI: University of Michigan Press.
- Waller, J. A., S. L. Dance, and N. K. Nichols. 2016. "Theoretical Insight into Diagnosing Observation Error Correlations Using Observation-minus-background and Observation-minus-analysis Statistics." *Quarterly Journal of the Royal Meteorological Society* 142 (694): 418–31. <https://doi.org/10.1002/qj.2661>.
- Waller, J. A., D. Simonin, S. L. Dance, N. K. Nichols, and S. P. Ballard. 2016. "Diagnosing Observation Error Correlations for Doppler Radar Radial Winds in the Met Office UKV Model Using Observation-Minus-Background and Observation-Minus-Analysis Statistics." *Monthly Weather Review* 144 (10): 3533–51. <https://doi.org/10.1175/MWR-D-15-0340.1>.
- Waller, Joanne, Susan Ballard, Sarah Dance, Graeme Kelly, Nancy Nichols, and David Simonin. 2016. "Diagnosing Horizontal and Inter-Channel Observation Error Correlations for SEVIRI Observations Using Observation-Minus-Background and Observation-Minus-Analysis Statistics." *Remote Sensing* 8 (7): 581. <https://doi.org/10.3390/rs8070581>.
- Wang, Tianlin, Christopher Ruf, Scott Gleason, Darren McKague, Andrew O'Brien, and Bruce Block. 2020. "Monitoring GPS Eirp for Cygnss Level 1 Calibration." In *IGARSS 2020 - 2020 IEEE International Geoscience and Remote Sensing Symposium*, 6293–96. Waikoloa, HI, USA: IEEE. <https://doi.org/10.1109/IGARSS39084.2020.9324491>.
- Wang, Tianlin, Christopher Ruf, Darren McKague, Anthony Russel, Andrew O'Brien, and Scott Gleason. 2021. "The Important Role of Antenna Pattern Characterization in the Absolute Calibration of GNSS-R Measurements." In *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*, 144–46. Brussels, Belgium: IEEE. <https://doi.org/10.1109/IGARSS47720.2021.9554506>.
- Wang, Tianlin, Christopher S. Ruf, Scott Gleason, Andrew J. O'Brien, Darren S. McKague, Bruce P. Block, and Anthony Russel. 2021. "Dynamic Calibration of GPS Effective Isotropic Radiated Power for GNSS-Reflectometry Earth Remote Sensing." *IEEE Transactions on Geoscience and Remote Sensing*, 1–12. <https://doi.org/10.1109/TGRS.2021.3070238>.
- Wang, Tianlin, Valery U Zavorotny, Joel Johnson, Yuchan Yi, and Christopher Ruf. 2019. "Integration of Cygnss Wind and Wave Observations with the Wavewatch III Numerical Model." In *IGARSS 2019 - 2019 IEEE International Geoscience and Remote Sensing Symposium*, 8350–53. IEEE. <https://doi.org/10.1109/IGARSS.2019.8900481>.
- Werner, Debra. 2023. "NOAA Kicks off NEON Weather Satellite Program." *SpaceNews*, May 9, 2023. <https://spacenews.com/noaa-kicks-off-neon-weather-satellite-program/>.
- White, Warren B. 1995. "Design of a Global Observing System for Gyre-Scale Upper Ocean Temperature Variability." *Progress in Oceanography* 36 (3): 169–217. [https://doi.org/10.1016/0079-6611\(95\)00017-8](https://doi.org/10.1016/0079-6611(95)00017-8).
- Williams, Edwin A., William A. Crossley, and Thomas J. Lang. 2001. "Average and Maximum Revisit Time Trade Studies for Satellite Constellations Using a Multiobjective Genetic

- Algorithm.” *The Journal of the Astronautical Sciences* 49 (3): 385–400.
<https://doi.org/10.1007/BF03546229>.
- World Meteorological Organization. 2015. *Valuing Weather and Climate: Economic Assessment of Meteorological and Hydrological Services*. Geneva, Switzerland: World Meteorological Organization.
- Wu, Longtao, Hui Su, Xubin Zeng, Derek J. Posselt, Sun Wong, Shuyi Chen, and Ad Stoffelen. 2024. “Uncertainty of Atmospheric Winds in Three Widely Used Global Reanalysis Datasets.” *Journal of Applied Meteorology and Climatology* 63 (2): 165–80.
<https://doi.org/10.1175/JAMC-D-22-0198.1>.
- Yang, John Xun, Yalei You, William Blackwell, Cheng Da, Eugenia Kalnay, Christopher Grassotti, Quanhua (Mark) Liu, et al. 2023. “SatERR: A Community Error Inventory for Satellite Microwave Observation Error Representation and Uncertainty Quantification.” *Bulletin of the American Meteorological Society*, April. <https://doi.org/10.1175/BAMS-D-22-0207.1>.
- Zavorotny, Valery U., Scott Gleason, Estel Cardellach, and Adriano Camps. 2014. “Tutorial on Remote Sensing Using GNSS Bistatic Radar of Opportunity.” *IEEE Geoscience and Remote Sensing Magazine* 2 (4): 8–45. <https://doi.org/10.1109/MGRS.2014.2374220>.
- Zec, Josko, W. Jones, Ruaa Alsabah, and Ali Al-Sabbagh. 2017. “RapidScat Cross-Calibration Using the Double Difference Technique.” *Remote Sensing* 9 (11): 1160.
<https://doi.org/10.3390/rs9111160>.
- Zhang, Shixuan, Zhaoxia Pu, Derek J. Posselt, and Robert Atlas. 2017. “Impact of CYGNSS Ocean Surface Wind Speeds on Numerical Simulations of a Hurricane in Observing System Simulation Experiments.” *Journal of Atmospheric and Oceanic Technology* 34 (2): 375–83. <https://doi.org/10.1175/JTECH-D-16-0144.1>.
- Zou, Cheng-Zhi, Mitchell D. Goldberg, Zhaohui Cheng, Norman C. Grody, Jerry T. Sullivan, Changyong Cao, and Dan Tarpley. 2006. “Recalibration of Microwave Sounding Unit for Climate Studies Using Simultaneous Nadir Overpasses.” *Journal of Geophysical Research* 111 (D19): D19114. <https://doi.org/10.1029/2005JD006798>.