EXPLORING ANTARCTIC LAND SURFACE TEMPERATURE EXTREMES

USING CONDENSED ANOMALY DATABASES

by

GLENN EDWIN GRANT

B.Sc., Southern Oregon University, 1983

M.A., University of Colorado, 2012

A dissertation submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Geography

2017

This dissertation entitled:
Exploring Antarctic Land Surface Temperature Extremes Using Condensed Anomaly Databases
written by Glenn Edwin Grant
has been approved for the Department of Geography

_____
Professor Mark Serreze (Geography)

_____
Mr. David Gallaher (National Snow and Ice Data Center)

Date: _____

The final copy of this dissertation has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Grant, Glenn Edwin (Ph.D., Geography)

Exploring Antarctic Land Surface Temperature Extremes Using Condensed Anomaly Databases

Dissertation directed by Professor Mark Serreze

Satellite observations have revolutionized the Earth Sciences and climate studies. However, data and imagery continue to accumulate at an accelerating rate, and efficient tools for data discovery, analysis, and quality checking lag behind. In particular, studies of long-term, continental-scale processes at high spatiotemporal resolutions are especially problematic. The traditional technique of downloading an entire dataset and using customized analysis code is often impractical or consumes too many resources.

The *Condensate Database Project* was envisioned as an alternative method for data exploration and quality checking. The project's premise was that much of the data in any satellite dataset is unneeded and can be eliminated, compacting massive datasets into more manageable sizes. Dataset sizes are further reduced by retaining only anomalous data of high interest. Hosting the resulting "condensed" datasets in high-speed databases enables immediate availability for queries and exploration.

Proof of the project's success relied on demonstrating that the anomaly database methods can enhance and accelerate scientific investigations. The hypothesis of this dissertation is that the condensed datasets are effective tools for exploring many scientific questions, spurring further investigations and revealing important information that might otherwise remain undetected.

This dissertation uses condensed databases containing 17 years of Antarctic land surface temperature anomalies as its primary data. The study demonstrates the utility of the condensate database methods by discovering new information. In particular, the process revealed critical quality problems in the source satellite data. The results are used as the starting point for four case studies, investigating Antarctic temperature extremes, cloud detection errors, and the teleconnections between Antarctic temperature anomalies and climate indices.

The results confirm the hypothesis that the condensate databases are a highly useful tool for Earth Science analyses. Moreover, the quality checking capabilities provide an important method for independent evaluation of dataset veracity.

Acknowledgements

Funding for this research was provided by the National Science Foundation through the Polar Cyberinfrastructure Program, Project 1251257, “Condensate Database for Efficient Anomaly Detection and Quality Assurance of Massive Cryospheric Data”, David W. Gallaher and Qin Lv principal investigators. This dissertation was also made possible with the kind help and patient encouragement of many NSIDC staff.

Table of Contents

List of Tables

List of Figures

1 Introduction

Antarctica is a land of extremes.

Extreme cold temperatures, high winds, climatic isolation, and the highest average elevation of any continent make Antarctica an inhospitable place. This dissertation uses new tools to investigate one aspect of those conditions, the extremes of Antarctic land surface temperatures. The questions this study examines are: What are the typical extreme temperatures around the continent? Are the magnitudes of the extremes changing, i.e. are they getting lower or higher over time? Are extreme temperatures becoming more or less common? And are there particular weather events or climatological factors that influence the occurrence of extreme temperature events? These are broad questions, and traditional research methods would require considerable time and extensive data processing to address even one of them. This study, however, uses a new set of tools for data exploration that accelerate the research process as well as identifying data quality issues that would otherwise remain hidden.

This dissertation's overarching hypothesis is that the methods developed herein provide a faster, more efficient process for data exploration, anomaly discovery, and quality control than conventional methods. The methods are evaluated using a series of scientific investigations into Antarctic temperature extremes. The results successfully demonstrate that these methods have multiple benefits, including finding new, unexpected events, replicating previous work, and exposing quality problems. Methodological limitations and unforeseen challenges are also discussed. Of the results presented, the most compelling discovery is that

the underlying dataset suffers from serious quality defects, and that the methods developed here fulfill a critical need for verification and validation of remotely sensed datasets.

1.1 Background

1.1.1 Motivation

Satellite-based Earth Observing Systems have changed our view of the world. The data collected is a great asset for science, commerce, and public safety. However, the volume of data that is received and recorded expands every year, and has reached the point where it is difficult for scientists to download, process, or even perform a cursory examination on the bulk of it. As a result, many studies obtain data from trusted data brokers, effectively relying on third parties to ensure that the data is of high quality. Typically, the data is then subsampled into smaller spatial regions or coarse temporal increments prior to analysis. This traditional process has proven fallible [e.g., Eisenman et al., 2014; Grant, 2012; Grant and Gallaher, 2015]. Simply exploring a large dataset, looking for events of interest, can be prohibitively time consuming, and ensuring the data's quality is even more problematic.

In response to the challenge of processing these massive datasets, the Condensate Database Project (NSF project number 1251257, *Condensate Database for Efficient Anomaly Detection and Quality Assurance of Massive Cryospheric Data*) sought more efficient ways of handling, exploring, and quality checking satellite datasets. Datasets approaching so-called "Big Data" scales were of particular interest. This project acknowledges that the problems of Big Data analysis could be overcome, in large part, simply by using sufficiently powerful hardware. Yet, with limited budgets, bandwidth, and time, and with an ever-expanding collection of

satellite data, a more elegant answer was required, one that was not reliant on a brute-force hardware solution.

The Condensate Database Project was built on the knowledge gained through a previous project, Data Rods (NSF project number 1546243) [Gallaher and Grant, 2012]. The Data Rods Project attempted to accelerate temporal data analysis by reorienting spatially-oriented images (flat files) into temporal columns ("data rods"), where each data rod contained all the information at one location through time. A key finding of the project was that I/O, the action of sending data to and from the computer processors, was the primary bottleneck when processing massive datasets. No amount of additional processing power would accelerate an analysis when the processors were sitting idle most of the time, waiting for the data to arrive. The problem, then, was to reduce the volume of data that must be shuffled between storage and processing. If the analytic process could be front-loaded, where datasets are preprocessed to reduce their volume, analytic efforts may be accelerated. This was the motivation for the Condensate Database Project: eliminate any unnecessary data, "condensing" it, while retaining only the most crucial information. But which data should be retained? In the context of climate and Earth Science research, indicators of change, or data anomalies, are often the most important information.[1] These anomalies are the contents of the condensed datasets. The result is a series of reusable, compact datasets, housed in databases that can be queried for data exploration purposes. As a side benefit of this process, various data quality problems

[1] Nomenclature note: In scientific literature, the term *anomaly* sometimes refers to the difference between a data value from its mean; thus, virtually all data points may be anomalous, even if the difference is small. As used in this dissertation, however, an anomaly is a true anomalous data point, one that is, in some respect, an outlier in the dataset, often an extreme value.

become immediately apparent. Section 1.2.1 describes the criteria for determining anomalous data, and Section 3.1 details the Condensate Project's architecture.

Since the condensation process necessarily excludes much of the original data, not all science questions can be posed against the condensed datasets. Nor are the condensed datasets considered archival quality; they are tools for data exploration, and users may need to revert to the full, uncondensed datasets to fully analyze interesting events. This dissertation take the approach of using the condensed datasets as a starting point for finding anomalous events and selectively accessing the full source data as necessary.

1.1.2 Antarctic Land Surface Temperature Extremes

New tools for rapid data exploration and discovery, such as the condensed databases, are worthwhile only insofar as the results actually prove effective at enhancing the scientific research process. The current study demonstrates the use of the condensed databases by using them to answer real-world scientific questions and validating data quality. The data chosen for this purpose, a massive, high-resolution dataset of the land surface temperatures of the Antarctic continent, was based on the author's personal interest in Antarctica.

Land Surface Temperatures (LSTs) are a key variable in understanding the health of Antarctica's ice sheets and ice shelves, and polar regions in general. A land surface temperature is the sensible temperature of the thin, topmost layer of the Earth's surface, often called a "skin temperature". In some cases this layer is only a few molecules thick, but it is a key parameter for understanding climate and the Earth's energy budget. Since LSTs can be measured via satellite, they are sometimes used as a proxy to estimate surface air temperatures (SATs) where ground instrumentation is sparse. (While LSTs and SATs are related, they are not identical. This

topic is discussed at length in Section 1.3.4.) In Antarctica, where weather stations may be hundreds of kilometers apart and poorly maintained, the ability to remotely-sense LSTs is especially useful. Trends in LSTs provide a measure of climate change impacts on the continent and may inform predictions of glacier and ice shelf retreat and subsequent sea level rise. This dissertation uses LSTs as its primary data, with supporting information from ground observations and reanalysis products. Condensed databases of LSTs are used to explore and quantify temperature extremes on the Antarctic continent, looking at the spatial distributions and temporal trends. Teleconnections between LST extremes and external climate processes are assessed, and case studies of two noteworthy extreme temperature events are used to attribute atmospheric conditions that may lead to extreme surface temperatures. And finally, two potential data quality issues are also examined.

Antarctic surface temperature trends have been heavily studied [e.g., Miliaresis, 2014; Comiso, 2000; Shuman and Comiso, 2002; Meyer et al., 2016]. Antarctic temperature *extremes*, however, have been explored only for narrow regions or individual events of interest [e.g., Marshall and King, 1998; Turner et al., 2009]. Yet temperature extremes are also associated with significant melting events on ice sheets, and have been implicated in the sudden destruction of ice shelves [Nghiem et al., 2012; Scambos et al., 2004].

Official record-breaking temperatures in Antarctica occur irregularly but always attract media attention. Recently, on 24 March 2015, the Argentinian base of Esperanza measured a new highest Antarctic surface air temperature of 17.5 °C [Skansi et al., 2017; WMO, 2017]. This record was in the relatively warm “banana belt” of the Antarctic Peninsula. However, the previous official record high temperature, 15 °C, was recorded at Vanda Station, in the

McMurdo Dry Valleys, on 5 January 1974 [Krause and Flood, 1997]. Prior to that, Esperanza had again held the record of 14.4 °C, set on 20 October 1956.

Record low temperatures are rarer. Vostok Station set a record of −89.2 °C on 21 July 1983, exceeding the older Vostok record of −87.4 °C set on 25 August 1958 [Turner et al., 2009; Stepanova, 1959]. The processes that allow for such extreme cold temperatures, and the potential limitations, are still being investigated. Interestingly, a theoretical minimum temperature was proposed during the International Geophysical Year era. Initial estimates put the coldest possible temperature at Vostok around −80 °C with a lapse rate of 0.5 °C per 100 meters of additional elevation [Shliakov, 1958]. Soon after, the radiation balance assumptions used by Shliakov were discredited when Vostok Station went colder than the theory by 5.4 °C [Wexler, 1959]. In this same era, a new theoretical coldest temperature of −200 °C was proposed for South Pole [McCormick, 1958]. However, since the calculations used only the radiation balance as a mechanism for surface energy exchange, brushing aside convective effects, advected air masses, and eddy conduction, it was considered highly speculative. Vostok's −89.2 °C record still stands, but satellite evidence now hints that significantly colder temperatures may be found at higher elevations [Campbell et al., 2015].

The actual range of temperatures across Antarctica is largely unknown because there simply is not enough instrumentation in place to measure it. New highs and lows may be occurring frequently, but without a long record of dense ground-based observations, there isn't a clear picture of the absolute temperature extremes or long-term trends. Long-term trends in extreme temperatures, and the processes which may be controlling the extremes, are generally undescribed. That is the motivation for the data exploration within this dissertation. Using

remotely-sensed thermal images, this study examines Antarctica as a whole, seeking the current state, trends, and causes of surface temperature extremes.

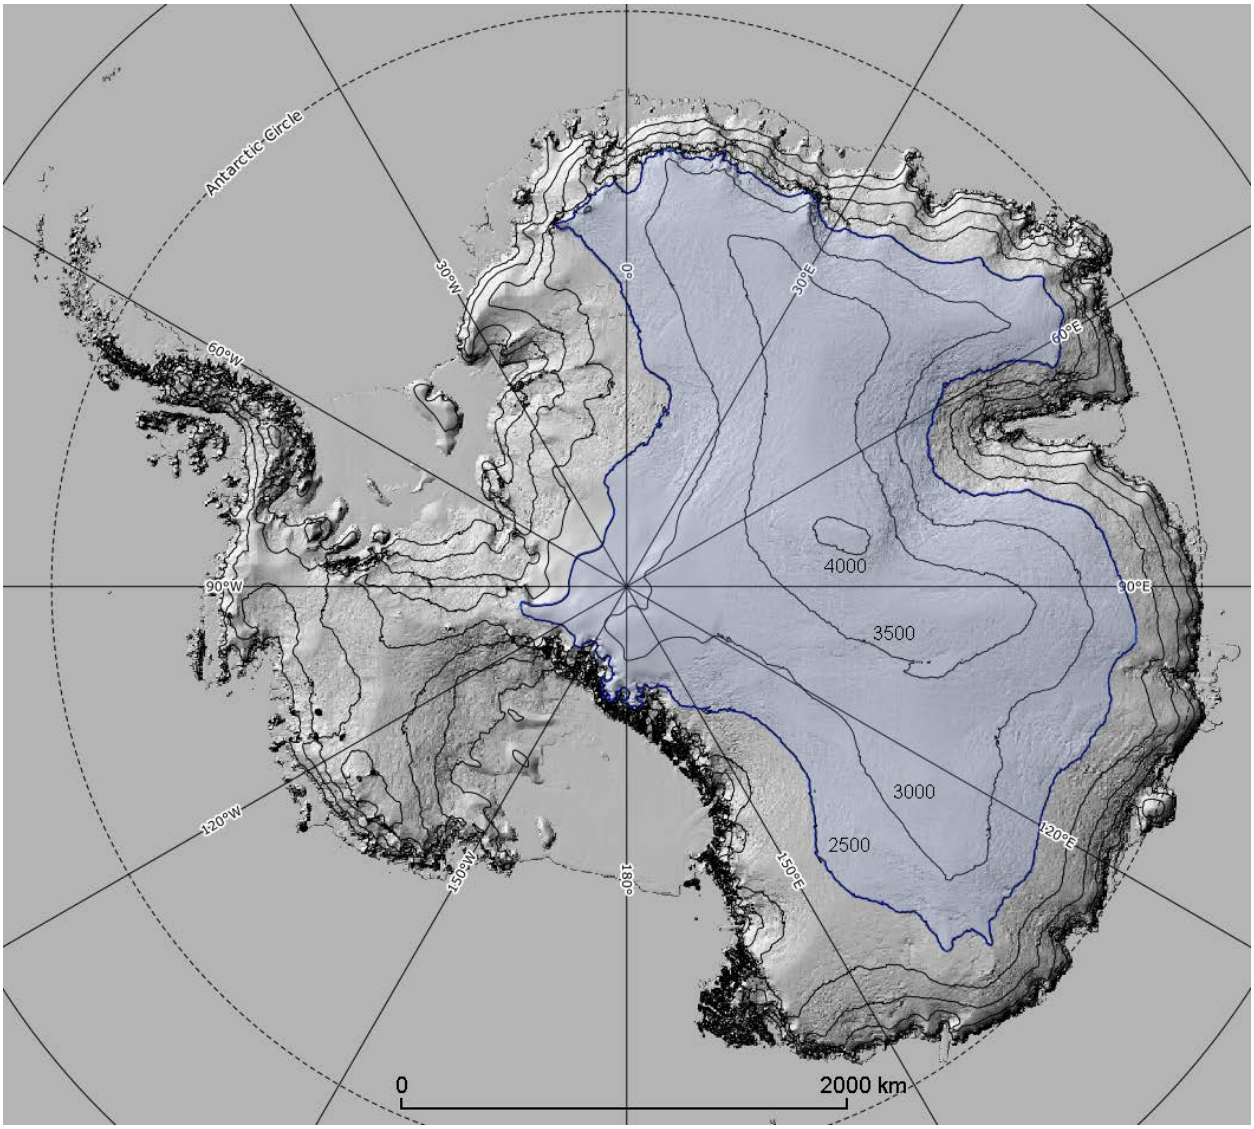

Figure 1.1. The Antarctic continent with elevation contours in 500 m increments [elevations from Bamber et al., 2009]. The blue shaded area is the Antarctic Plateau climate regime, defined by the WMO as areas above 2500 m.

1.1.3 Geography

Ninety eight percent of the land surface of Antarctica is covered by ice. In some studies, floating ice sheets are also considered a land surface; this dissertation only considers grounded ice or bare land. The massive ice sheets in the interior of the continent represent the bulk of the continent, and are roughly divided by the Transantarctic Mountain Range into East Antarctic and West Antarctic regions. These interior regions are largely high-elevation plateaus with very little exposed orography. Not counting floating ice shelf area, the land surface area of Antarctica is approximately 12 million km^2, with a mean elevation of 2450 meters [Drewry et al., 1983] (Figure 1.1).

At the time of this writing, Antarctica has 44 permanent full-year research stations. Most of the stations are located along the coastal areas, with a large collection near the tip of the Peninsula. Many of the stations collect manned weather observations, although the data is not always timely or available. The manned observations are supplemented with approximately 63 automated Weather Stations (AWSs), sprinkled throughout the interior regions as well as the coasts, providing a broader picture of the current weather conditions. If all the weather stations, manned and unmanned, were equally distributed across the continent, there would be roughly one weather observation per 112,000 km^2. The distribution is not nearly that even, of course, with most of the observations occurring near the coasts. The instrumentation that does exist is subject to harsh weather conditions and, if unmanned, may go for extended times without maintenance or calibration.

The low density of weather observations throughout the continent, especially in the interior, makes weather forecasting more uncertain, amplifies the errors in reanalysis products,

and increases the difficulty of assessing any climatic changes that may be occurring. Satellite observations offer a partial solution to overcoming the inadequate spatial sampling. Atmospheric sounders, in particular, can infer conditions through the air column, and thermal infrared (TIR) sensors can estimate the surface temperatures. The combination provides a fuller picture of the weather conditions across the continent. The satellite-based LST measurements used as the primary data in this dissertation are a derivative product from multiple TIR channels.

1.1.4 Objectives

The foremost objective of this dissertation is to demonstrate scientific research using the condensate databases. For the new methods to be considered successful, they must meet several goals: 1) Provide for rapid data exploration, surpassing conventional methods; 2) detect data quality problems; and 3) be readily understandable to a diverse scientific audience. Each of these goals is evaluated in the course of the scientific investigations described below.

Surface temperature extremes are generally short-term phenomena, the result of a set of weather conditions that together create an extreme temperature. If Antarctica's climate is changing, the frequency, magnitude, and locations of extreme temperature events may also be changing. Extreme temperature events are thus one way of gauging the effects of climate change. Antarctica is physically and thermally isolated, and, due to sparse ground-based data collection, the effects of climate change throughout the continent are not entirely clear. By using satellite data rather than interpolated ground-based readings, this study seeks to infill the data gaps and augment the state of knowledge about the effects of climate change in Antarctica.

The specific science objectives are:

a. Baseline Descriptive Statistics. The anomalous temperatures used throughout this study are relative to the 5th and 95th percentiles. Generating basic descriptive statistics was a necessary step in determining anomalous temperatures for the condensation process, but the statistics also provide a look at the overall spatial distribution of Antarctic land surface temperatures. Concurrent with the percentile statistics, the Condensate Database Project also generated additional descriptive statistics, including the LST means, medians, standard deviations, linear regressions through time, correlation coefficients, sample populations, and minima/maxima. The regression slopes provide a simplified basis for comparing LST temperature trends with past studies, and a quick way of identifying areas of Antarctica that may be undergoing the greatest climatic changes. All the statistics are calculated by interannual month, season, and year. Selected charts of the baseline descriptive statistics are provided in the Results and Discussion, Section 4. This effort also reveals data quality problems in the source dataset.

b. Spatial Distribution of Extremes. When studying extreme temperatures, perhaps the most obvious questions are: 1) where are the absolute coldest and warmest spots on the continent? And 2) at any location, what are the coldest or warmest temperatures that have been observed? This effort uses both the baseline LST descriptive statistics and anomaly databases to determine the range and distribution of the most extreme temperatures at each location (minima and maxima) and, overall, the most extreme Antarctic temperatures observed during the dataset's timespan. Additional data quality issues related to cloud contamination become apparent during this analysis.

c. Temporal Distribution of Extremes. The anomaly databases are used to address two questions: 1) Are extreme temperatures becoming more or less common, and 2) if so, in what direction? Using the populations of low (5^{th} percentile) and high (95^{th} percentile) LSTs through time, the prevalence of extreme temperatures is explored. Also presented is a brief look at the yearly and seasonal distribution of LST magnitudes, which are strongly skewed.

d. Teleconnections. Previous studies have found relationships between Antarctic temperatures and patterns of atmospheric circulation, although causality is often inconclusive. Still, these types of studies can indicate possible climate connections between Antarctica and distant regions, pointing the direction for future research. This dissertation demonstrates how similar studies could be conducted using the condensate database information, and evaluates LST extremes in relation to sea ice extent, the Southern Annular Mode (SAM), and the Southern Oscillation Index (SOI).

e. Case Studies. Two case studies of extreme temperature events are presented, with attribution of the weather conditions responsible for each. The analyses are modeled after Turner et al.'s [2009] study of the 1983 extreme cold temperature event at Vostok Station, Antarctica, and build on Haran et al.'s [2016] research into extreme cold spots. In the former study, both satellite and atmospheric reanalysis data were applied to explain the occurrence of the world-record cold temperature. This dissertation uses a similar methodology as Turner et al. to analyze an extreme cold temperature event, and utilizes elevation data and aspect angles to help explain both high and low temperature extremes. Two additional case studies are also presented: In the first, puzzling descriptive statistics prompted a study that examines the long-term trend of South Pole LSTs relative to manned near-surface air temperature observations.

The second case study looks at a discontinuity in the percentage of clouds detected, which has implications for the overall dataset quality.

The significance of detecting quality problems buried deep within massive datasets cannot be overstated. In the current era, where scientific results (especially those related to climate prediction) are increasingly questioned, the veracity of the underlying observational datasets takes on critical importance. Scientists and research groups increasingly rely on data obtained from trusted "data brokers"; few researchers have the time, resources, or expertise to thoroughly vet every dataset they use. Yet their results depend on the data's accuracy, and even an inconspicuous error can propagate from observations to reanalyses, into models, and ultimately affect the final interpretations of study results. More than ever, comprehensive tools are needed to evaluate datasets and detect problems. The tools developed as a part of the *Condensate Database Project*, described in this dissertation, represent a new, holistic method of rapidly assessing the quality of gridded satellite datasets. The techniques scale well, and excel when used at high spatiotemporal resolutions against continental or global-scale datasets. The results demonstrate that even supposedly high-quality datasets may have hidden flaws that can adversely affect scientific results.

1.2 Previous Work

1.2.1 Definition of Extreme Temperatures

When this study uses the term extreme temperatures, it is referring not only to record-setting highs and lows (although those would certainly be included), but "anomalous" measurements that are well outside the norm. A single record temperature may be the result

of a unique combination of weather events, which, by itself, does not explain much. By expanding the dataset to include a larger domain of temperatures in space and time, this study explores the effects of climate change on the Antarctic continent in greater depth. The question then arises, how should an extreme be defined? Statistically speaking, anomalies are at the tails of a temperature distribution. Since each grid point on the Antarctic continent may have a different distribution, this study analyzes each location independently. But because individual anomalies may be due to noise, adjacent extreme events are also considered before concluding that any one anomalous temperature is an extreme. Importantly, sophisticated statistical techniques exist for modeling extremes, such as Extreme Value Theory and variations that address non-stationary sequences [Coles, 2001]. However, these techniques are often computationally expensive, making them difficult to apply across massive multi-decadal satellite datasets. What is needed is a common definition of *extreme* that is applicable across broad spatial areas with varying temperature distributions; the definition must be calculated consistently, and without user intervention. To settle on an exact definition, we turn to scientific conventions and meteorological standards.

Extremes of any climate variable may be absolute or relative. For instance, if anything over 10 °C is declared as an extreme, this would be an example of an absolute threshold. Most scientific papers, however, prefer relative thresholding, using a statistical confidence level relative to a mean or median. Using this latter method, a value is an extreme only if it deviates by a predetermined standard deviation or reaches a threshold percentile. Typical percentile thresholds are 5% and 10%. Indeed, the Intergovernmental Panel on Climate Change (IPCC) outlines specific criteria to define weather extremes using percentile definitions [Hartmann et

al., 2013, Box 2.4, p.221]. Many of the IPCC extreme definitions are based on contiguous daily data and focus on seasonal or annual records, e.g., the maximum or minimum temperature in a given time period. More apropos for this dissertation are definitions based on relative thresholds, such as the number of days or nights that exceed a given threshold. This dissertation uses the 5^{th} and 95^{th} percentile thresholds for determining extremes.

The percentile thresholds are calculated relative to temperature distributions across a variety of time increments. Specifically, the extremes are relative to interannual months, seasons, or all years ("yearly"). For example, a January monthly extreme would be a value that falls below the 5^{th} or above the 95^{th} percentile of all temperature observations during the month of January for the timespan of the dataset. Since the dataset used in this dissertation spans 17 years, a January monthly extreme temperature is relative to 17 years of observations, all of it collected during the month of January. Similar logic applies to the seasonal and yearly extremes. No detrending is performed: Extremes could have been chosen relative to a modeled curve (trend) rather than the actual observations, however, the dataset's noisiness resulted in generally poor models with very low coefficients of determination (R^2). Using modeled curves as a reference would also have obscured interesting events.

Determining extremes is straightforward. Evaluating the *changes* in extremes, whether in terms of magnitude or frequency, is more problematic. Lewis and King [2017] argue that, for a proper evaluation of changes in temperature extremes, analyses must move beyond examining simple shifts in the mean and variance to higher-order statistical moments that reflect a change in symmetry, i.e., skewness and kurtosis. This dissertation agrees with that assessment in principle, and also examines the skewness of the temperature extremes.

When determining extreme values, data quality becomes an important issue. Outlying data points are often due to noise or processing errors, and may be mistakenly identified as extremes. Easterling et al. [2016] emphasizes the importance of addressing climate data discontinuities, inhomogeneities, data gaps, and erroneous values. Much of the work herein focuses on these very problems. Data gaps are handled via monthly, seasonal, and yearly aggregation. Inhomogeneities and discontinuities, which may be the result of multiple different data sources or observation methods, are addressed by using a single data source. The bulk of the data quality assurance efforts ensure that individual outliers are not treated as important events. Specifically, spatially- and temporally-isolated outliers are ignored in favor of larger, clearly defined patterns of temperature extremes.

1.2.2 Climate Regime Subregions

The World Meteorological Organization (WMO) sets the standards for surface temperature measurements and adjudicates record temperatures [WMO, 2008]. Briefly, measurements must be made by calibrated instrumentation at official meteorological recording sites, which may be manned or unmanned. Instrumentation is sparse in Antarctica, and record temperatures are always a hot topic of interest, so there is a detailed validation procedure that occurs when a potential record is observed. Antarctic record temperatures, however, are also controversial because there is a large difference in the climatic conditions between the moderate coastal areas, especially along the Peninsula and sub-Antarctic islands, and the cold continental interior. The WMO's Commission for Climatology (CCI) addresses this issue by delineating different climate regimes based on elevation and latitude [Skansi et al., 2017; WMO, 2107]. While the all-encompassing Antarctic Region includes any area south of 60

degrees south latitude, two special subregions have been defined: Low-lying areas on or immediately adjacent to the Antarctic continent, and regions above 2500 meters elevation. In general, the higher elevation areas are the interior plateau of East Antarctica, although some areas in West Antarctica and the Transantarctic Mountains also meet this definition. The distinction between climate regimes can be important when conveying information to the public. When Esperanza's recent record high temperature was reported, the WMO took pains to point out the regional division; the new record temperature was only valid for the coastal areas, while an inland AWS designated D-80 still holds the record plateau high temperature of −7.0 °C, set on 28 December 1989 [WMO, 2017].

For reasons of data quality and sample size, this dissertation generally concentrates on the higher elevation plateau regions above 2500m, matching with the WMO's definition of inland Antarctic extreme temperatures (see Figure 1.1).

1.2.3 Recent Climate Change Findings

Recent studies have shown that some regions of the Antarctic continent are warming, while others may be cooling. The Peninsula region in particular, which until recently had been showing faster warming than most other places on earth, now appears to be cooling slightly due to atmospheric circulation changes [Vaughan et al., 2003; Turner et al., 2016]. The West Antarctic Ice Sheet (WAIS) is of particular interest because of its significant potential to influence sea levels. For example, Bromwich et al. [2013] have shown that Central West Antarctica is warming, attributed partially to changes in atmospheric circulation patterns over the Amundsen/Ross Sea sector of Antarctica (but still not fully understood). Likewise, Ding et al. [2011] also found warming in West Antarctica, and attributed the warming to changes in

atmospheric circulation patterns, covarying with anomalous tropical pacific sea surface temperatures (SSTs). Schneider et al. [2012] found similar connections between circulation patterns, SSTs, and sea ice extent. Examples of studies that addressed variations in LSTs across the continent as a whole include Thompson and Solomon [2002], Schneider and Steig [2004], Schneider et al. [2006], and Bracegirdle et al. [2008], all showing statistically significant increases in LSTs to some degree, except for East Antarctica. Often, the changes in atmospheric circulation are correlated with various climate indices, including the Southern Annular Mode (SAM), the Pacific South American (PSA) mode, and to a lesser degree the El Niño Southern Oscillation [e.g., Kwok and Comiso, 2002].

In contrast to the temperature trend studies mentioned above, there are few papers that address only extreme Antarctic temperatures. These papers more often focus on the effects of extreme temperatures on biota [e.g., Quayle et al., 2003; Peck et al., 2004]. While the WMO occasionally reports on Antarctic temperature extremes, there are few actual analyses of the weather conditions and climate teleconnections surrounding them. A notable exception is Turner et al.'s [2009] in-depth analysis of the record low 1983 temperature at Vostok Station, Antarctica. Their paper considered the physical principles of longwave radiative transfer and energy advection, and brought to bear multiple data sources. It used a long-term Antarctic climatology dataset, 6-hourly surface observations from Vostok Station, orographic information, and reanalyses from the European Centre for Medium-Range Weather Forecasts (ECMWF) 40-year and National Centers for Environmental Prediction - National Center for Atmospheric Research (NCEP-NCAR) datasets. The study's conclusions were corroborated by TIR data from NOAA/AVHRR satellite passes. In their summary, Turner et al. describe how the

extreme low temperature was the result of a combination of near-optimum events, especially the persistence of a cold mid-tropospheric vortex, low surface advection, and clear skies. The Case Studies section of this dissertation builds on their methodology but uses LSTs as the primary data source instead of surface air temperature (SAT).

Most studies that examine Antarctic surface temperatures use SATs rather than satellite-based LSTs. The advantage of SAT is that it provides a continuous, calibrated measurement on the ground, and is unaffected by clouds, viewing angles, or infrequent satellite revisit schedules. Surface air temperatures, however, are limited by the lack of spatially dense measurements. The sparseness of *in situ* measurements presents one of the most challenging problems for studies of polar regions. A variety of statistical techniques have been used to interpolate between distant measurement sites and extrapolate Antarctic SATs, including simple linear interpolations, weighted spatial regressions, and distance-based semi-variance techniques. Monaghan et al. [2008] interpolated Antarctic near-surface temperatures across the entire continent using readings from just 15 manned stations. In their study, they used a "Kriging-like method" to identify similar climatic zones across the continent, and then interpolate surface temperatures across those zones to reconstruct Antarctic temperature trends. In comparison, Chapman and Walsh [2007] interpolated sparse surface air temperatures over correlation length scales exceeding 1000 km. Steig et al. [2009] used a combination of satellite TIR data with Automated Weather Station observations to interpolate the spatial covariance of surface temperatures. This latter method used custom software that accounts for the temporal changes in the spatial covariance patterns. The results were used to validate temperature reconstructions and estimate regional ice sheet warming since 1957.

However, their method of interpolating backwards in time was unsupported by corroborating data, such as from ice cores or manned stations.

Perhaps the most thorough effort at spatially reconstructing surface air temperatures has been the Reference Antarctic Data for Environmental Research (READER) project [Turner et al., 2004]. READER is a continuously updated database of Antarctic climatology from manned and automated weather stations. The project participants recognized from the outset that *in situ* measurements of Antarctica are sparse, spatially and temporally discontinuous, and may not always be properly calibrated. READER is a collective effort to build a consistent climatology data record, and provides a quality-controlled basis for studies that may want to interpolate data records to other locations on the continent [Monaghan et al., 2008].

1.3 Physical Basis

Satellite-based LST measurements are estimated from thermal infrared (TIR) radiances in the 8 μm – 14 μm electromagnetic spectrum. A cloud-free atmosphere is mostly transparent in this spectral band, and the measured TIR radiances are largely unaffected by sunlight, which is predominately composed of light in the shorter-wavelength visible spectrum. This combination makes LST measurements possible year-round, during both day and night. Clouds, however, pose a significant problem because they can block the satellite's view of the surface. When this happens, the LST measurements are said to be "cloud contaminated" and composed of, at least in part, the temperatures of the cloud tops and not the earth's surface. Thus, the elimination of cloud-contaminated measurements is important for any study that uses LSTs, a topic discussed throughout this dissertation.

Thermal infrared remote sensing of surface temperatures starts with a basic principle of physics: that every object radiates electromagnetic waves in relation to its thermal energy [Petty, 2004]. For earth surfaces, emissions are predominately in the infrared spectrum, from roughly 4 µm to 50 µm, typically peaking around 15 µm. Not all surface types radiate uniformly throughout the spectrum. The monochromatic radiant intensity (I_λ) of a surface is a product of its spectral emissivity (ε_λ) and the theoretical spectral radiant intensity (B_λ) that a blackbody at temperature *T* would produce:

$$I_\lambda = \varepsilon_\lambda B_\lambda(T) \qquad \text{(A)}$$

where the radiant intensity, B_λ, is calculated via the Planck Function:

$$B_\lambda(T) = \frac{2hc^2/\lambda^5}{e^{hc/\lambda kT} - 1} \qquad \text{(B)}$$

where *h* is Planck's constant, *c* is the speed of light, λ is the wavelength of light, and *k* is the Boltzmann constant [Petty, 2006]. Equation (B) states that the intensity of emissions at any wavelength is directly related to the temperature. Using this equation, the actual temperature can be estimated by remotely sensing the radiant emissions from surfaces.

A plot of Equation (B) shows that, for polar surfaces of snow and ice near ~255 K, the peak of theoretical black body emissions occurs in the TIR band at approximately 11 µm (Figure 1.2). In the near-IR region, below about 3 µm, the emissions are insignificant. The emissivity of many surfaces in the TIR spectrum is almost constant and near unity, as if they were black bodies; as a first-order estimate, this assumption of near-black body emissivity also applies to snow and ice. Thus, it is reasonably straightforward to calculate what the observed *brightness*

temperature (*Tb*) of a snow surface would be using the observed spectral radiance and the inverse of Planck's Function:

$$Tb = B_\lambda^{-1}(I_\lambda) \quad (C)$$

Combining Equations (A) and (C) produces a measure of brightness temperature solely in terms of the radiant intensity (Planck's Function), the surface temperature, and the surface emissivity:

$$Tb = B_\lambda^{-1}[\varepsilon_\lambda B_\lambda(T)] \quad (D)$$

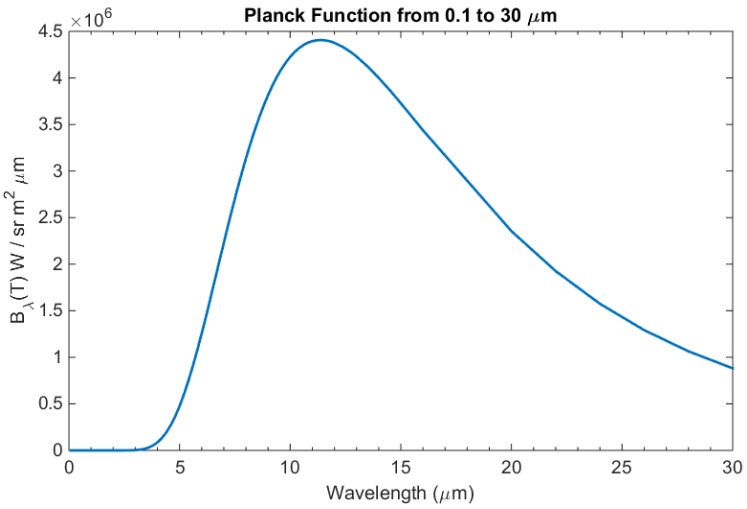

Figure 1.2. The Planck Function, Equation (B), plotted for a surface temperature of 255K (-18 °C) and emissivity of 1.0. A surface at this temperature radiates primarily in the longwave infrared region, peaking at a TIR wavelength of about 11 µm, with virtually no emissions below 3 µm.

If the emissivity is near unity and there are no external sources of emission or absorption, Equation (D) tells us that the brightness temperature would be a close estimate of the actual surface temperature. If only things were that easy. To convert this brightness temperature into a true surface temperature, estimates must account for any intervening atmospheric absorption, emissions, reflections, and scattering in the line of sight, and then integrate across the entire spectrum while accounting for the spectral sensitivity of the sensor.

Most satellite-based infrared (IR) sensors have on-board systems for calibrating thermal radiance measurements, largely eliminating sensor calibration as a significant issue [Xiong and Barnes, 2006]. For precise calculations, however, unity can no longer be assumed for surface emissivity, or that the emissivity is uniform in all directions. Fortunately, under most circumstances, atmospheric scattering in the thermal infrared bands is negligible, as are reflections. When viewing through clear skies, this leaves atmospheric absorption and emissions, and off-nadir changes in emissivity as the prime contributors to LST uncertainty.

The sensing of snow and ice, especially that of ice sheets, has a distinct advantage over other surfaces in that the surface is typically smooth (at the scales measured by satellite sensors), uniform over broad spatial areas, and the spectral emissivity is well-characterized. This simplifies the assumptions compared to varied surfaces, such as forest canopies.

Broadly, then, the steps for remotely sensing the actual surface temperature are:

1) Measure surface radiance at wavelengths where atmospheric absorption is minimal.
2) Avoid cloudy days when the satellite would be measuring the emissions of clouds rather than the surface (weighted by the clouds' optical opacity).
3) Estimate the amount of atmospheric absorption that occurs.
4) Compensate for sensor-specific calibration, transforms, and sensor degradation.
5) Estimate and apply the emissivity of the surface.

Of these steps, accurate cloud detection over ice sheets is the most error-prone. There are also assumptions being made about the surface type. Some aspects of snow and ice surfaces, including snow grain size and ablated bare ice surfaces, can influence emissivity, yet these are seldom fully addressed in most surface temperature algorithms. This study uses LSTs derived

from the MODerate Resolution Imaging Spectrometer (MODIS), and the MODIS processing does not distinguish between snow surface types. Arguably, although the differences in emissivity tend to be small, a better surface temperature could be calculated if the surface types were described more accurately.

The next sections address each of the challenges in measuring LSTs from space.

1.3.1 Remote Sensing of LSTs

Within the IR spectrum, the majority of atmospheric absorption is due to water vapor. There are, however, "window" regions where IR emissions transfer mostly unattenuated: In the mid-IR region of 3 μm – 5 μm, and the thermal IR window at 8 μm – 12 μm. For ice surfaces, the 8 μm – 12 μm region is used because it is here that the thermal emissions are greatest, improving the signal-to-noise ratio. Ideally, if the amount of atmospheric absorption is known, a single IR channel may be used to derive the surface temperature. Indeed, some older algorithms do exactly that, using pre-estimated atmospheric absorption coefficients. However, the uncertainty is large when using the single-channel technique because water vapor content is highly variable in the atmospheric column. Now, most multi-spectral and hyper-spectral sensing systems use a so-called "split window" technique, wherein the sensor simultaneously measures radiances centered near 11 μm and 12 μm [e.g., Price, 1984; Wan, 2014; Wan and Dozier, 1996]. This study uses MODIS LSTs generated via a split window algorithm, produced from the MODIS sensor's TIR bands 31 and 32, centered at 11 μm and 12 μm [Wan et al., 2004].

The split window technique relies on the fact that atmospheric absorption due to water vapor is less at 11 μm than at 12 μm, and the ratio of absorption between the wavelengths is well-characterized. Comparing the difference in the measured surface radiances at these two

wavelengths produces in a robust estimate of atmospheric absorption without knowing, *a priori*, the amount of water vapor along the line of sight. The generalized split window algorithm used by MODIS is:

$$T_s = C + \left(A_1 + A_2 \frac{1-\varepsilon}{\varepsilon} + A_3 \frac{\Delta\varepsilon}{\varepsilon^2}\right) \frac{T_{31}+T_{32}}{2} \quad \text{(E)}$$

$$+ \left(B_1 + B_2 \frac{1-\varepsilon}{\varepsilon} + B_3 \frac{\Delta\varepsilon}{\varepsilon^2}\right) \frac{T_{31}+T_{32}}{2}$$

[from Wang et al., 2008] where:

$$\varepsilon = \frac{(\varepsilon_{31}+\varepsilon_{32})}{2} \quad \text{(F)}$$

$$\Delta\varepsilon = \varepsilon_{31} - \varepsilon_{32} \quad \text{(G)}$$

and T_s is the surface temperature, T_{31} and T_{32} are the brightness temperatures from MODIS bands 31 and 32, regression coefficients A_n and B_n are from multi-dimensional look-up tables based on view angle, surface type, atmospheric conditions, and time of day, and ε_{31} and ε_{32} are the surface emissivities in the respective bands. Assuming the surface is viewed through a clear sky, the spectral surface emissivity is known, and that any calibration issues or uncertainties in the physical sensor system have been accounted for, T_s is our calculated surface temperature.

1.3.2 Energy Balance at the Surface

To gain a better understanding the conditions that can create temperature extremes, it is helpful review the vertical components of the surface energy balance and consider their effects. The surface energy balance can be expressed as:

$$R_{net} + G + H_S + H_{LE} = 0 \quad \text{(H)}$$

where R_{net} is the net radiation, G is the snow heat flux, H_s is the turbulent sensible heat flux, and H_{LE} is the turbulent latent heat flux, or latent heat of evaporation. Any horizontal components, at least in the context of Antarctic ice sheets, are considered negligible and not included in Equation (H).

To fully appreciate the energy exchange, it helps to break down the radiation term even further. The net radiation is the sum of the downward (or "downwelling") and upward ("upwelling") shortwave and longwave radiation:

$$R_{net} = S\downarrow + S\uparrow + L\downarrow + L\uparrow \quad \text{(I)}$$

where S is the shortwave solar radiation, L is the longwave radiation, and the arrows represent downwelling (↓) and upwelling (↑) fluxes. For the purposes of this discussion, we define shortwave radiation as wavelengths in the range of 100 nm (0.1 μm) to 3000 nm (3 μm), which encompasses the ultraviolet, visible, and near-IR bands.[2] Solar longwave radiation in wavelengths above 3 μm is negligible compared to the shortwave energy. Any downwelling solar longwave energy is absorbed in the atmosphere long before reaching the surface, so it is not a factor in these energy balance calculations. By convention, the signs of all terms in Equations (H) and (I) are downward positive, implying that $S\uparrow$ and $L\uparrow$ must always be negative. Shortwave energy exchange is only present during daylight hours, thus $S\uparrow$ and $S\downarrow$ are essentially zero during the winter months when much of Antarctica is under 24 hours of darkness. When the sun is up, the net shortwave radiation is controlled by the surface albedo (reflectivity):

$$S\downarrow + S\uparrow = R_{ss}\cdot(1-\alpha) \quad \text{(J)}$$

[2] Definitions of the spectral bands vary.

where R_{ss} is the total solar shortwave radiation, and α is the surface albedo. The right-hand term of Equation (J) is thus the net shortwave energy absorbed by the surface. Clean, dry snow and ice surfaces will typically have an albedo around 0.85, perhaps reaching 0.9 for very fresh snow, and down to 0.8 for old snow [Warren, 1982]. Bare, ablated ice can be somewhat lower, although these areas are infrequent on the Antarctic Plateau. Melting can drastically reduce snow's albedo, to 0.5 or less. The plateau region of Antarctica rarely, if ever, experiences melting, so it is safe to assume that the albedo stays at 0.8 or above in most high elevation areas. The practical take-away from Equation (J) is that the majority of daytime solar flux is reflected away, but enough of it may be absorbed to warm the surface. Cloud cover, with its high albedo, can significantly reduce the downwelling shortwave flux. Although the loss of shortwave flux will reduce one source of surface warming, the presence of clouds may intensify the incoming longwave flux.

The downwelling longwave flux (L↓) on the right side of Equation (I) is governed by the air column temperature, cloud base temperature, and atmospheric constituents (which affect the air's emissivity). Longwave energy from clouds and water vapor can substantially warm the surface, although the effects of clouds should be minimal when considering clear-sky LSTs. Accurate computation of $L\downarrow$ can be challenging, and is usually accomplished using 1-D radiative transfer models or direct measurement.

The upwelling longwave radiation (L↑) is a product of the surface temperature. Snow and ice surfaces reflect very little of the downwelling flux, but absorb it efficiently and then re-emit it again. The total power radiated by a surface is related to the temperature, as specified by the Stefan-Boltzmann law:

$$j = \varepsilon\sigma T_s^4 = L\uparrow \quad \text{(K)}$$

where j is the radiant flux energy per unit surface area, ε is the broadband emissivity of the surface, σ is the Stephan-Boltzmann constant, and T_s is the surface temperature in degrees Kelvin. At the temperatures found on the Antarctic ice sheet, the Planck Function (Equation (B) shows this energy will peak in the TIR band at around 11 μm, with no appreciable emissions below about 3 μm (see Figure 1.2). Thus, none of the radiant flux given by Equation (K) is included in the shortwave emissions (S↑) of Equation (I).

To understand the total net radiation budget, Equations (I), (J) and (K) can now be combined:

$$R_{net} = R_s\cdot(1-\alpha) + L\downarrow + \varepsilon\sigma T_s^4 \quad \text{(L)}$$

During the dark Antarctic winter months, R_s drops to zero and Equation (L) is dominated by the longwave fluxes. These are generally negative (cooling the surface), although a low, warm cloud cover can equilibrate the energy balance or even send it positive [Serreze and Barry, 2014]. Thus, it is likely that the most extreme warm surface temperatures during the winter months will not be viewable by satellite because of cloud cover. Nighttime cold extremes, on the other hand, which depend heavily on outgoing longwave radiation under clear skies, should be detectable.

Returning to the energy balance Equation (H), the ground heat flux G represents the energy coming from the snow under the surface. The rate of the ground heat flux can be estimated from the equation:

$$G = -k_s \left.\frac{dT}{dz}\right|_{z=0} \quad \text{(M)}$$

where dT/dz is the temperature gradient at the surface, k_s is the thermal conductivity of the snow, and the equation is evaluated at the surface (z=0). Note that, since all the terms in Equation (H) are "downward positive", a positive value of G would indicate that warmer temperatures are flowing downward and the skin is losing energy to the underlying snowpack. To keep the equation balanced and the flux direction consistent, the sign of the ground heat flux is reversed in Equation (M). There is rarely empirical data available for dT/dz, and instead it must be modeled. Models for ground heat flux typically assume the instantaneous surface temperature ($T_{z=0}$) is the same as T_{air} at an elevation of 0 meters. If a time series of surface temperatures is available at the same location, dT/dz – and thus G – can be estimated using a model that incorporates the thermal diffusivity of snow.

The turbulent components of the energy balance equation, H_s and H_{LE}, are the sensible and latent heat fluxes. These fluxes are weather dependent and much more uncertain. Specifically, H_s is controlled by conduction with the overlying air and turbulent eddies. The turbulence, in turn, is created by convection and advection of the air. Surface heat conduction through sensible heat flux (H_s) is highly efficient, and sometimes the surface temperature (T_s) is assumed to be the same as the surface air temperature (T_{air}). While close, this latter assumption may not always be correct, with T_{air} commonly a few degrees warmer than T_s [e.g., Comiso, 2003]. Section 1.3.4 explores the differences between LSTs and SATs in greater depth.

The latent heat flux, H_{LE}, is the energy exchanged due to evaporation or, as would typically be the case on the higher elevation parts of the Antarctic ice sheet, sublimation. Like H_s, the rate of H_{LE} is also controlled by turbulent eddies. Turbulence causes an increase in the humidity gradient between the laminar air layer (immediately on top of the snow surface) and

near-surface air layer, and thus an increase in evaporation or sublimation, cooling the surface. In very cold environments, however, the sensible heat flux often far outweighs the heat fluxes due to sublimation, and H_{LE} may be ignored [King and Anderson, 1994]. A possible exception to this generality might when the sensible and latent heat fluxes have opposite signs, as might be the case when the air is warmer than the surface (H_s is positive) and the surface is evaporating (H_{LE} is negative); this situation might occur on a very warm day. Another possibility would be when the air is colder than the surface ($T_{air} < T_s$) and water vapor is either freezing or condensing on the surface, e.g. frost due to adiabatic cooling of a rising air mass. In this latter case, H_s would be negative and H_{LE} positive, but LSTs are unlikely to be measureable due to fog and cloud cover from the condensing water vapor.

In summary, the surface energy balance controls the surface temperature. Computing each component of the energy balance (Equation (H) relies on a variety of assumptions, the most uncertain being the atmospheric conditions that drive the downwelling longwave radiation. Depending on the surface temperature extremes being observed and the atmospheric conditions, some terms in the equation may be negligible. Other terms must be modeled, or else estimated using weather observations or reanalysis data.

1.3.3 Uncertainties and Validation

Much of the uncertainty surrounding satellite-based measurements of LSTs centers on the problems of quantifying surface emissivity, pixel contamination due to undetected clouds, and accurate determination of atmospheric optical depth. There are also geometrical uncertainties caused by off-nadir viewing angles; these can often be reduced to problems related to anisotropic emissivity or atmospheric absorption due to an increased optical path.

The uncertainties inherent in the atmospheric optical depth are addressed by the split-window technique described in the previous section. The following paragraphs address emissivity uncertainties and the cloud contamination problems that cause an inherent *cold bias* in satellite-based LST measurements.

When calculating LSTs, the surface type may strongly affect the infrared emissivity, adversely influencing the temperature measurements. Antarctica's "land" surface is almost entirely ice and snow, with few surface features other than small-scale roughness and sastrugi. The texturally homogeneous surface simplifies IR remote sensing of snow and ice sheets, reducing the possibility of mixed pixel problems where multiple different surface types are observed in a single pixel's field of view. NASA's MODIS processing uses a table of known surface types for each location to estimate emissivity, and categorizes the bulk of Antarctica as a combined snow/ice surface.

If the surface type changes, the LSTs may be erroneous. In particular, the dynamic margins of the Antarctic continent may have newly exposed open ocean, bare rock, or melt pools that may still be classified as snow/ice, increasing the uncertainty in those observations. Throughout this dissertation, the surface is assumed to be ice, snow, or firn, and occasionally bare rock. Temperature extremes in the coastal regions are not fully explored here, in part due to the possible emissivity variance, but largely because the coastal areas are often cloudy, making data density thin and satellite-based conclusions difficult.

The emissivity of dry snow and ice in the longwave IR spectrum is well-characterized, with $\varepsilon \approx 0.99$ for most snow surfaces [Hori et al., 2006; Dozier and Warren, 1982]. The emissivity of bare rock is more variable but tends to be high as well, above 0.95 for dark surfaces [e.g.,

Valor and Caselles, 1995]. Snow grain size is a factor only outside of the 11 µm to 12 µm region used by the LST calculations. Off-nadir emissivity only starts to fall significantly when the viewing angle exceeds about 60°, where ε may drop to 0.98 and fall rapidly at greater viewing angles [Warren, 1982]. Many satellite-based thermal remote sensors naturally avoid the oblique-angle viewing problem by having a narrow field of view centered on nadir. For example, MODIS sensors have a field of view (FOV) of ±55 degrees [NASA, 2016]. Moreover, MODIS LSTs are derived from observations that use the smallest nadir viewing angles from the available cloud-free observations, substantially reducing uncertainties due to off-nadir viewing angles. Even so, the MODIS algorithms do account for differences in attenuation due to the viewing angle, even if they are small [e.g., Key et al., 1997; Wan, 1999; Wang et al., 2008]. Of more importance is attenuation due to clouds in the FOV.

Accurate satellite-based surface temperature measurements rely on clear-sky conditions. MODIS LSTs are generated for all land surfaces regardless of cloudiness, and quality flags are included in a separate SDS indicating the confidence of clear sky conditions at each pixel. Using these flags, no additional cloud mask is necessary. The quality flags can thus provide a kind of coarse frequency and distribution for cloud coverage. In this study and most others, cloud contaminated LSTs are rejected. Thus, surface warming due to longwave emissions from clouds may never be observed. In polar regions this effect typically manifests itself as a cold bias in the data [e.g., Koenig and Hall, 2010; Hall et al., 2008; Hall et al., 2012; Wan et al., 2002]. The cold bias problem may also be compounded by imperfect cloud detection algorithms, where LSTs are contaminated with temperatures from cold cloud tops.

Determining whether the sky is clear can be a difficult problem. Due to the low contrast between clouds and snow/ice surfaces, detecting clouds over ice sheets is challenging. Over the years, numerous cloud detection algorithms have been implemented and refined, yet all still have rather high rates of error over ice sheets [e.g., Ackerman et al., 1998; Ackerman et al. 2008]. The problem is compounded at night when the near-IR spectral reflectance bands cannot be used [Liu et al., 2004]. As a result, although this dissertation used only the highest quality LST values (uncertainty <=1 K), some cloud contamination was unavoidable. In an attempt to mitigate the cloud detection errors, the MODIS cloud masking was augmented with a spatial adjacency test wherein anomalous extreme pixels are ignored unless there are additional extreme pixels in the vicinity. Furthermore, the extreme temperature events used in the case studies were selected based on broad temporal and spatial patterns of extreme temperatures, and only after manual inspection of visible surface imagery to ensure clear skies (see Sections 4.5.2 and 4.5.3).

Referring back to Equation (K), note that temperature is raised to the fourth power. The implication of this exponent is that colder surfaces radiate much less than warmer surfaces, reducing a sensor's ability to measure the incoming radiation over the noise floor and thus increasing the uncertainty. The problem becomes especially acute when trying to measure cold ice sheets. To overcome this limitation, satellite sensors often use an expanded bandwidth for measuring TIR radiances. This effectively gathers more light, increasing the signal to noise ratio. However, integrating over a larger bandwidth also decreases the precision of the measurements (while surface emissivities may be relatively constant within each channel, there are often significant variations in the spectral absorption due to atmospheric constituents).

Importantly, the LSTs have been calibrated and validated against known surface types, typically in the warmer mid-latitudes [Wan et al., 2002, 2004; Wang et al., 2008]. Their accuracy when measuring extremely cold snow surfaces may be substantially different from the stated accuracy, probably worse.

Given all these sources of error and uncertainty, empirical measurements are the most practical method of validating surface temperatures. For MODIS, calibrations are performed using an airborne sensor, identical to the satellite's, but viewing surface temperatures from a lower position in the atmosphere [e.g., Baum et al., 2000; Wan et al., 2002]. This technique significantly reduces the errors due to uncertainties in the atmospheric optical depth. Intercomparisons are then made between the airborne sensor's data and a simultaneous satellite overflight. The primary shortcoming of airborne sensing is that it is intermittent and cannot cover all possible surface types, nor has it been used extensively over cold polar surfaces, especially at night. The airborne calibrations are nonetheless considered accurate, although evidence of sensor drift or degradation may call that assumption into question.

1.3.4 Comparisons between LSTs and Near-Surface Air Temperatures

Most studies that examine surface temperatures are actually using *near-surface* air temperatures (SAT), variable T_{air} when used in equations within this dissertation. These air temperature measurements are typically taken 1.2 m to 2 m above the ground with the sensor in a vented enclosure, sometimes referred to as a "screen" temperature [WMO, 2008]. In areas of the world where surface meteorological measurements are sparse, LSTs may also serve as a proxy for SAT. Although this dissertation is looking specifically at the values and trends for LSTs,

not SATs, a brief explanation is given here to provide a foundation for relating the two measurements and expanding on the role of downwelling longwave radiation.

LSTs are tightly coupled with SATs, and correlations over snow and ice surfaces approach 90% or greater [e.g., Urban et al., 2013]. However, the energy exchange processes have different influences and often result in a distinct bias. Before using an LST as a proxy for an SAT, the biases and uncertainty between the two measurements must be quantified. The scale difference is one consideration: A satellite-based LST measurement represents the mean surface temperature over a large area (1 km^2 for MODIS LSTs) whereas SAT is an instantaneous air temperature at a single point. Fortunately, most of the Antarctic ice sheet has little or no orography that would exacerbate uncertainties due to scale mismatches. Other factors that may also bias satellite-derived LSTs relative to SAT include insolation, heterogeneity of the surface emissivity, thermal conduction through the ground, turbulent air mixing (or inversions occurring due to a lack of turbulent mixing), and the effects of having an atmosphere or thin clouds between the satellite and the surface. *In situ* measurements of SAT, of course, are undeterred by clouds, but comparisons between the two are inevitably problematic because the LSTs cannot be measured through clouds.

Because a snow surface is an excellent absorber of longwave radiation (and air is not), the reasons for the difference between T_{air} and T_s become clear. In the absence of radiative forcing, and especially with high advection and turbulence, the snow surface and the air would be nearly the same temperature. Under clear skies, however, $L\uparrow$ (Equation (I) grows significantly and, without the radiative feedback that occurs under a cloud cover, results in a large loss of energy from the surface. During daylight hours, $S\downarrow$ may also be high but the

snow's albedo reflects much of the energy. The high solar zenith angles (SZA) found in polar regions also contribute to the low effectiveness of sunlight at warming the surface. The theoretical net effect is that remotely-sensed LSTs, which only can only be measured in cloud-free conditions, are often colder than SATs.

Past studies have indeed found that LSTs of snow-covered surfaces are typically lower than the SAT. For example, Comiso [2003] found monthly averaged AVHRR LSTs to be 0.34 °C lower than T_{air} over Arctic sea ice, with some seasonal variation. Turner et al. [2009] reported surface temperatures at Vostok Station, Antarctica, 2 K to 4 K lower than T_{air}, averaging 2.7 K. Using MODIS data, Westermann et al. [2012] found a mean lower bias of 3 °C at a wintertime permafrost site in Svalbard, and Wang et al. [2013] found average monthly mean differences of 2.9 °C to 8.1 °C in the area of the Lambert Glacier Basin, Antarctica, with the greatest variations during the winter months.

2 Data

The *Condensate Database Project* tested multiple different datasets during development. The information gathered while processing each dataset guided the development of the condensation algorithm. Example satellite datasets included: Daily sea ice concentration; Soil Moisture Active/Passive (SMAP) daily data of brightness temperatures and moisture content; AVHRR Polar Pathfinder (APP) daily albedo; and Special Sensor Microwave/Imager (SSMI) daily brightness temperatures. Level-4 data, such as the sea ice concentrations, provided a benchmark for gauging the effectiveness of the condensation algorithm against data without noise. SSM/I data, on the other hand, is notoriously noisy, and challenging to condense. Likewise, albedo data is seasonal, with low noise during the summer months but increasing noise at high solar zenith angles. The combination of test datasets resulted in a robust algorithm that can be tuned easily for varying data quality. However, the datasets listed above were relatively low volume and did not thoroughly test the capacity of the system to ingest, condense, and store massive volumes of data. As a comprehensive demonstration of the condensate process, a final, massive dataset was needed. At a continental scale, MODIS data at a 1 km spatial resolution and twice-daily temporal resolution, fulfills that requirement.

2.1 MODIS/Terra Land Surface Temperatures

This study uses land surface temperatures (LSTs) derived from NASA's Moderate Resolution Imaging Spectroradiometer (MODIS), which flies onboard both the Aqua and Terra satellites [NASA LP DAAC, 2017]. Due to sensor degradation issues with the MODIS sensor on Aqua, Terra is the primary data source used herein. The LSTs are part of the newly-reprocessed version of MODIS data, Collection 6 [Wan, 2014]. The Terra satellite was launched on December

18 1999 and has been collecting data since February of 2000. Terra is a polar orbiter, providing full coverage of the Antarctic continent at least twice each day. The level-3 LST data product, MOD11A1, contains twice-daily calibrated and gridded daily composite LSTs at a 1 km resolution [Wan, 2013]. Seventeen full years of data, 01 March 2000 through 28 February 2017, were downloaded from the Land Processes Distributed Active Archive Center (LP DAAC, https://lpdaac.usgs.gov/). The files are organized by date and packaged in Hierarchical Data Format (HDF), a self-describing file format that also includes metadata and ancillary scientific data sets (SDSs). The SDSs include pixel quality values, viewing angles, and times of observation on a pixel-by-pixel basis.

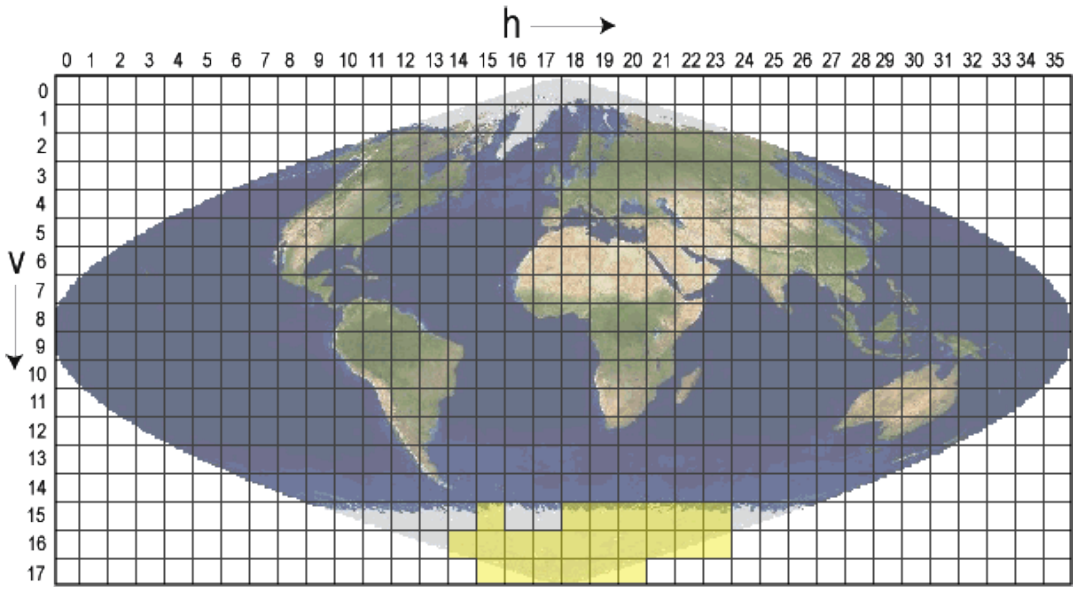

Figure 2.1. Tiles in the MODIS sinusoidal projection of the Earth. Horizontal (h) and vertical (v) tile numbering are indexed from zero in the upper left-hand corner. This study used the twenty three tiles highlighted in yellow, which cover the surface of Antarctica. Image from NASA (https://modis-land.gsfc.nasa.gov/MODLAND_grid.html).

The MODIS climate data record (CDR) breaks the surface of the Earth into a matrix of tiles, each covering a specific geographical area (Figure 2.1). This study uses the 23 tiles that cover the surface area of Antarctica, highlighted in yellow in Figure 2.1. Every tile contains 1200x1200 grid cells in an equal-area sinusoidal projection. The stated spatial resolution of each grid cell is 1 km, although it is actually 0.982 km on an edge [Wan, 2013]. The CDR's sinusoidal projection is generally useful for equatorial and mid-latitude visualizations but inappropriate for polar regions; throughout this dissertation, unless otherwise stated, all maps and charts have been reprojected into a Lamberts equal-area polar projection with the same physical parameters as the NSIDC EASE-Grid 2.0 standard [Brodzik et al., 2012]. Reprojections were performed using *proj.4* tools (http://proj4.org/).

Every daily file contains both daytime and nighttime LST SDSs, as determined by the local solar time of the observations. Not all dates were available for all tiles; of the 6209 potential dates in the time series, roughly 6148 were found for each tile, a success rate of 99%. A single image for one day is termed a "granule", and the LST data point at each grid cell is a pixel. Using twice-daily imagery from 23 tiles, the dataset comprised 141,407 files, or 282,814 granules, for 407,252,160,000 total pixels. The entire dataset size was 434.1 GBytes.

MODIS granules may be composed of pixels from several different orbital passes. NASA processing selects pixels by quality and viewing angle, giving preference to cloud-free observations and look angles close to nadir. Because the images are composites, pixel times may vary throughout any granule. Pixel observation times are included as separate SDSs in each file. The NASA processing converts the UTC times (t_{UTC}) of the observations to local solar time ($t_{LocalSolar}$), where the time shift to local solar time is based on the longitude of the grid cell's

centroid [Wan, personal communication]. This study uses UTC exclusively. Where necessary, LST local solar times have been converted back to UTC using the following sequence of equations:

$$t_{UTC} = t_{LocalSolar} * 0.1 - (\text{Longitude} / 15) \quad \text{(N)}$$

$$(t_{UTC} >= 24) \rightarrow t_{UTC} = t_{UTC} - 24 \quad \text{(O)}$$

$$(t_{UTC} < 0) \rightarrow t_{UTC} = t_{UTC} + 24 \quad \text{(P)}$$

The UTC date of the observations is contained in each file name. When converting the local solar times back to UTC, the observation date remains unmodified.

2.1.1 Cloud Masking Confidence Level

Cloud masking is performed automatically as a part of NASA's LST processing, and cloud-contaminated pixels are identified via a confidence flag. However, the processing uses a different clear-sky confidence level depending on the elevation. Over land and below 2000 m, the confidence level is 95%, dropping to 66% above 2000 m [Wan, 2013]. This variation in clear-sky confidence level creates an inconsistency in cloud contamination, and therefore LSTs spanning the 2000 m elevation contour. The confidence level variation is based on the results of validation testing over high elevation regions, and seeks to overcome a paucity of data under certain conditions in MODIS Collections 4 and prior [Wan, 2008]. However, the 2000 m elevation criterion appears somewhat arbitrary, and its efficacy over polar ice sheets has not been evaluated. Normally, a decrease in temperature is expected with increasing elevation due to adiabatic cooling, but at 2000 m the sudden decrease in average temperature can be attributed to cold cloud tops being tagged as "clear sky" and producing more cloud contaminated LSTs.

LST sample populations around the margins of the continent are often small, in large part because the coastal fringe is often cloudy, but also due to the higher confidence threshold required to declare clear-sky conditions. The small sample sizes make statistical evaluation problematic in some areas. Accordingly, the case studies used in this dissertation are concentrated at elevations above 2500 m, which provides a larger sample size, avoids the 2000 m cloud masking discontinuity, and is compliant with the WMO's 2500 m elevation threshold for Antarctic temperature records in the plateau region.

2.1.2 Discovered Data Errors and Artifacts

The discontinuity in cloud masking confidence level described in Section 2.1.1 is a documented feature of MODIS LST processing. The condensate baseline statistical processing in

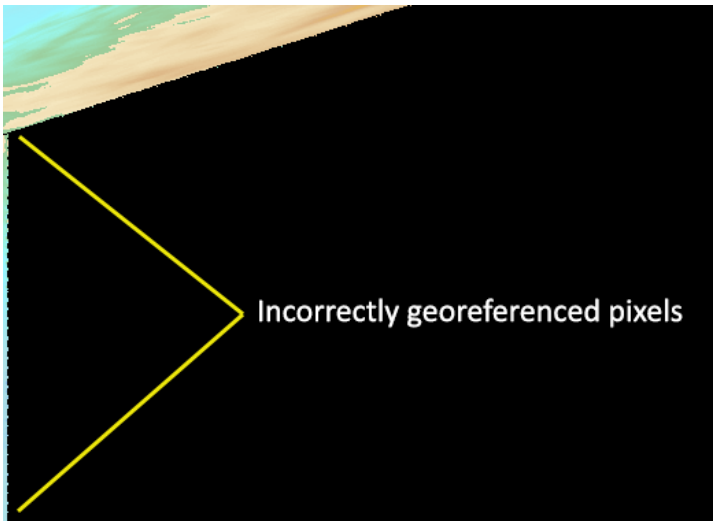

Figure 2.2. An example of pixels incorrectly georeferenced in a MODIS tile. The image shows a portion of tile h19, v17 (refer to Figure 2.1). The area at the top-left is part of East Antarctica, and the black area is located off the Earth's surface. A thin line of pixel data, improperly georeferenced, appears along the left-hand edge. Similar erroneous pixels appear in other tiles.

this dissertation, however, discovered two additional data artifacts, both apparently undocumented.

First, MOD11A1 collection 6 data contains incorrectly georeferenced pixels, an error not seen in the collection 5 CDR (Figure 2.2). The pixels appear to be located in space; it is unclear whether the pixels are real data with incorrect locations, or false data that was inadvertently introduced into the dataset during NASA processing. The condensation algorithm crops-out these pixels, which otherwise would be reprojected incorrectly and appear at unexpected locations on the Earth's surface.

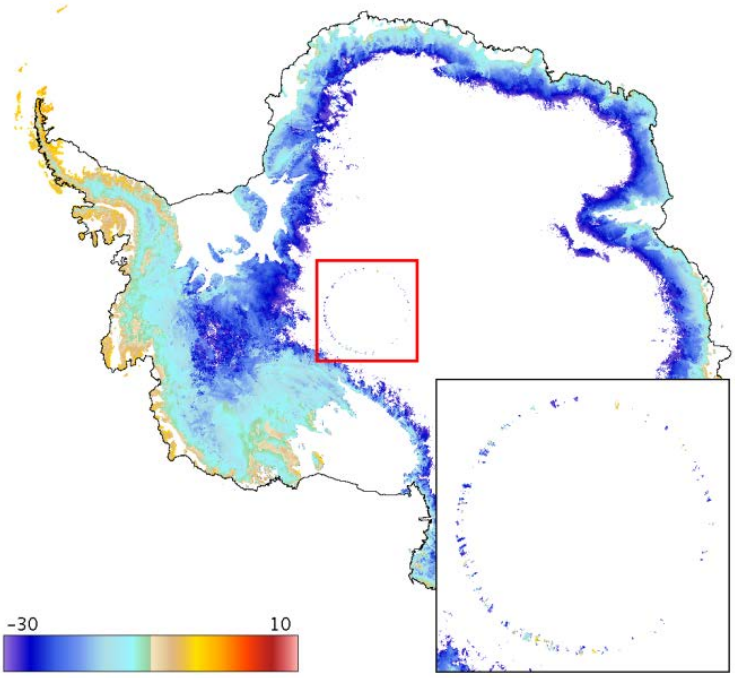

Figure 2.3. A systematic data error in the MOD11A1 dataset. The image shows maximum Antarctic winter temperatures in degrees C, where any temperature less than −30 °C is off-scale and appears white. A ring of spuriously warm pixels appears in the middle of the image at approximately −87 degrees latitude (identified with a red square, see inset for zoom). The cause of the erroneous pixels is undetermined.

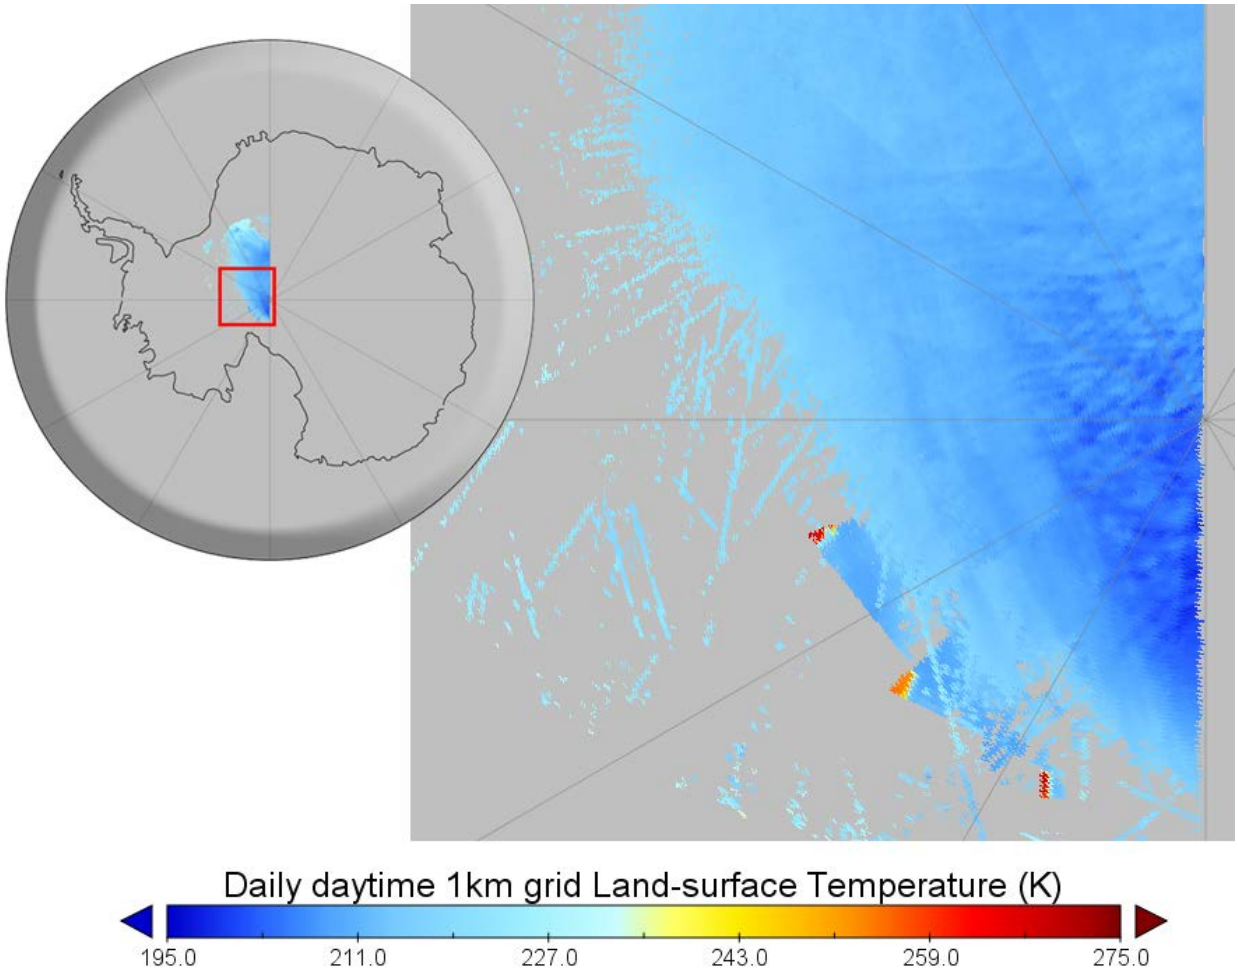

Figure 2.4. A single MODIS tile for 10 July 2010 (inset shows location on the Antarctic Ice Sheet). The tile data is composed of swaths from multiple satellite passes. Poorly calculated LSTs, shown in yellow-red, are evident. This erroneous data is at the farthest off-nadir points along individual swaths, suggesting that a geometrical software bug is leading to a radiometric miscalculation.

In the second form of data artifact, erroneously calculated LSTs appear as a circle of spurious pixels near the −87.1° parallel (Figure 2.3). These pixels, although clearly a systematic error, are within the bounds of reasonable LST values and could not easily be filtered from the cleaned dataset. They are present in the condensed databases, and are most apparent when searching for extreme warm or cold temperatures. Further investigation revealed that the

bright pixels are the result of an error that apparently occurs during the NASA processing when swath data is mosaicked into tiles. For example, Figure 2.4 shows an image of a single tile for one day, where spuriously warm pixels are visible at the extreme cross-track ends of swath data. The error occurs occasionally throughout the 17-year timespan of the data, and is present in both the Terra and Aqua LST datasets, Collections 5 and 6. The tile shown in Figure 2.4 is typical. Notably, no erroneous pixels were found in the level-2 swath source data (MOD11_L2) used to create the reprojected tiles, suggesting that a geometric and/or radiometric processing bug occurs during the mosaicking of the data from level-2 to level-3.

2.2 Reanalysis Datasets

The relationships between temperature extremes and atmospheric weather patterns are examined using two different reanalysis data sets: NASA's Modern-Era Retrospective analysis for Research and Applications, version 2 (MERRA-2), and the European Centre for Medium-Range Weather Forecast's ERA-Interim reanalysis.

2.2.1 MERRA-2

MERRA-2 reanalysis data is produced by NASA's Global Modeling and Assimilation Office (GMAO) and distributed through the Goddard Earth Sciences and Information Services Center (https://disc.gsfc.nasa.gov/uui/search/%22MERRA-2%22). The data used in this study was the hourly time-averaged, single-Level diagnostics assimilation at a 0.5° latitude x 0.625° longitude resolution [Bosilovich et al., 2016; GMAO, 2015].

2.2.2 ERA Interim

ERA Interim reanalysis data is produced by the European Centre for Medium-Range Weather Forecasts and distributed at http://apps.ecmwf.int/datasets/data/interim-full-daily/levtype=sfc/ [Dee, D.P. et al., 2011]. This dissertation uses surface-level atmospheric data at a 0.75° x 0.75° spatial resolution and 6-hourly temporal increments.

2.3 Elevation

Surface elevation data, which is incorporated into the condensate databases, was sourced from the 1 km Antarctic Digital Elevation Model (DEM) produced by Bamber et al. [2009] and distributed by the National Snow and Ice Data Center at: https://nsidc.org/data/docs/daac/nsidc0422_antarctic_1km_dem/#access_tools.

2.4 Surface Observations

Amundsen-Scott South Pole Station has collected weather observations since 1957. The Antarctic Meteorological Resource Center (http://amrc.ssec.wisc.edu/) provides South Pole hourly surface air temperature observations from 1988 onward [AMRC, 2016]. Hourly observations were used only if they indicated clear-sky conditions (code “SKC”).

2.5 Climate Indices

To explore the teleconnections with external climate processes, this study used three climate indices:

1. The Southern Annular Mode (SAM) index, also referred to as the Antarctic Oscillation (AAO), a monthly index available from the National Weather Service Climate Prediction Center [Marshall et al., 2016].

2. The El Niño Southern Oscillation Index (SOI), downloaded from the NOAA Centers for Environmental Information [NCEI, 2017].
3. The Antarctic Sea Ice Index, available for both the continent as a whole and regionally from the National Snow and Ice Data Center [Fetterer et al., 2016]. The index files contain both sea ice area and sea ice extent. This study uses the sea ice extent values.

3 Methods

Software and methods developed as a part of the *Condensate Database Project* performed the extreme LST detection and quantification tasks in this Dissertation. This software, and the associated databases, cataloged the LST extremes and made them available via database query tools and graphic visualizations. For the case studies, the database tools identified events of interest, prompting further investigations that used the full (uncondensed) dataset and reanalysis data.

The following sections detail the algorithms and data flow.

3.1 Anomaly Detection and Condensed Databases

The *Condensate Database Project* was motivated by the massive amounts of remotely-sensed data being generated, and the increasing difficulty of processing or exploring it in an efficient manner. The premise of the work was that, for many scientific questions, much of the data is extraneous and only serves to impede scientific exploration. If the excess data could be eliminated, the process of exploring massive datasets – and detecting anomalous events of great interest – could be accelerated.

The project produced an algorithm that reduces the size of large datasets by approximately two orders of magnitude, subject to the noise level of the data and the desired amount of "condensing". In the case of the MODIS LST data used in this study, the original dataset was condensed to approximately 3% of its original size. The condensed dataset captures both the key characteristics (spatial and temporal norms) and interesting, anomalous data. Note that the software does not use a *compression* algorithm. Compression would require a subsequent decompression step prior to accessing the data, which may obscure sought-after

values and impose a high processing overhead. Similar overhead burdens apply if using bitmap indexing and coding schemes [Li, 2013; Su et al., 2013]. Data aggregation and spatiotemporal subsampling techniques were also avoided because, while effective, they typically degrade transient or localized events. For these reasons, we chose condensation techniques that permit immediate, query-capable access to the actual data at a high resolution, not simply storing the metadata or compressed datasets, and foregoing the overhead of a decompression step.

Condensed datasets may not be able to answer the entire scope of scientific questions that would be posed against the full dataset. The condensed data is intended as a data exploration tool as well as an ancillary method of quality control. Yet, for studies that focus on the extremes of data values, the condensed datasets provide a valuable starting point for rapidly quantifying the spatial and temporal frequency of outlying data points. The full archival-quality dataset, uncondensed, is always available to answer questions beyond the capability of the condensed datasets.

The essence of the condensation algorithm is to identify unneeded data and eliminate it. Figure 3.1 shows a high-level flow diagram of the condensation process. Each step is detailed in the following sections. Critically, the algorithm produces two products: 1) A set of descriptive statistics at every gridded location, which captures the baseline information later used as a basis for data condensation, and 2) the actual condensed dataset. Baseline generation and data condensation can be performed across the entire timeline of data or temporally binned by any desired timespan. In this dissertation, the baselines and condensed datasets were generated for interannual seasons, months, and across the entire dataset (all years). The temporally-

restricted datasets are useful for addressing scientific questions with a seasonal aspect that might be obscured when using full-year data.

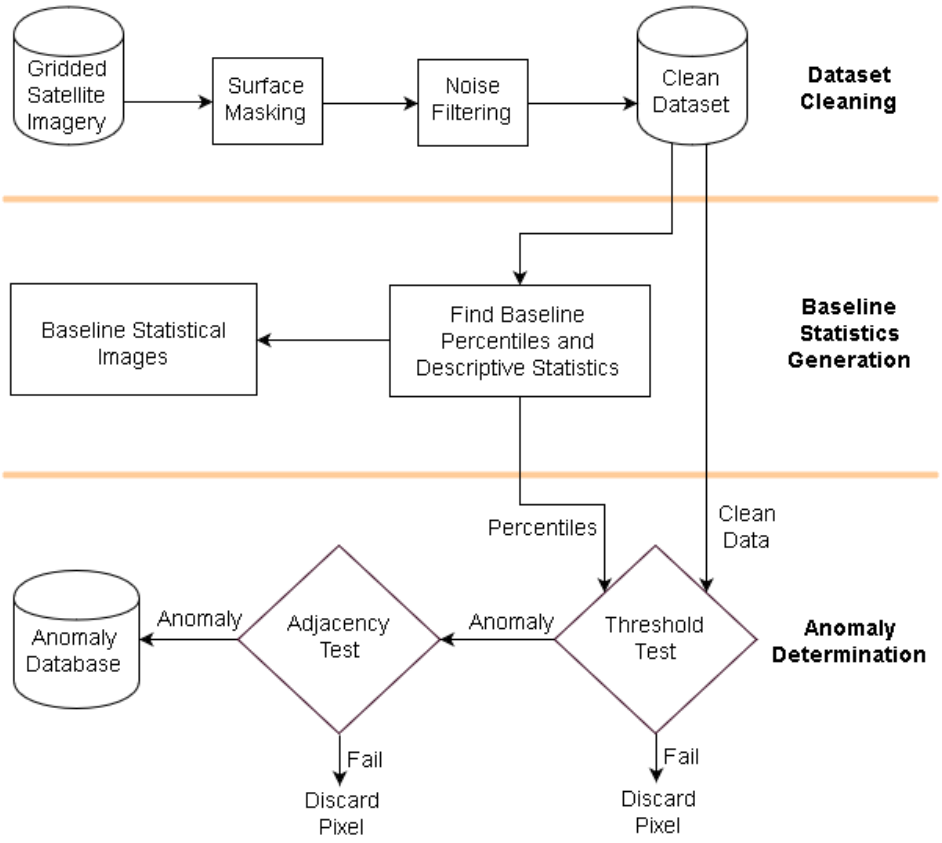

Figure 3.1. A process flow diagram for determining anomalies.

3.1.1 Dataset Cleaning

The source dataset passes through a multi-step process of cleaning before the condensation algorithm begins. Between surface masking and elimination of low-quality or cloudy pixels, the dataset was reduced to approximately 20% of its original size without discarding any valid LSTs.

The details of the cleaning steps are described in the following sections.

3.1.1.1 Surface Masking

Surface masking eliminates data storage for areas that are irrelevant or where no data is possible. For the MODIS MOD11A1 CDR, the dataset does not include surface temperatures for oceans or ice shelves, yet the source HDF data files necessarily come with empty placeholders where that data would exist. Likewise, the sinusoidal projection of the MODIS tiles necessarily includes areas off the globe. The total number of grid cells for the group of 23 MODIS tiles that cover Antarctica is 33,120,000. Of these, 14,278,503, or 43%, represent valid land surface locations. The unused grid cells significantly expand the source data files' sizes; these grid cells have been eliminated from condensed databases.

3.1.1.2 Noise Filtering

At its simplest level, noise filtering consists of determining the valid range of the data and then excluding any pixel that falls outside of it. For example, LSTs above 320 K or below 150 K are unrealistic in Antarctica, yet may be part of the CDR due to satellite telemetry errors or specular reflectance. The noise filtering step removes these from LST dataset.

A bit flag for Quality Control (QC) is also included as an SDS in each the MOD11A1 file. The QC flag reflects the uncertainty level of the temperature calculation, which is typically related to the possibility of cloud contamination. Since clouds can adversely affect the LST accuracy, lower-quality LSTs were deemed unacceptable and filtered from this study. The highest quality LSTs, used exclusively herein, have QC flags indicating uncertainties of <= 1 K. In general, the exclusion of low-quality data adversely affected sample sizes only around the cloudy periphery of the continent where maritime influences are common. In the interior plateau region, which represents the vast majority of Antarctica's surface area, the loss of low-

quality LSTs had minimal impact on sample sizes. Although there are year-to-year variations, the overall percentage of clear-sky pixels was 47%.

3.1.2 Creating Baselines

To determine whether data points are anomalous, the algorithm first determines the 5th and 95th LST percentiles at each grid cell. These baselines are used later for threshold testing (see Section 3.1.3). Other descriptive statistics, generated at the same time, include the means, medians, standard deviations, linear regression coefficients, and minima/maxima. The statistics include population counts for 1) all observations before dataset cleaning, 2) the number of samples remaining after cleaning, and 3) the number of anomalies found within those samples. The baseline descriptive statistics were generated for each month and season (interannually), and across all years without regard to season.

Once completed, the statistical baselines can be re-assembled into images (see Section 4.1 for examples). Although this image-generation step is not necessary for the condensation process, the images provide a wealth of information for users. The images are a quick way of visually exploring the entire dataset and quickly detecting areas of scientific interest or data quality problems.

3.1.3 Anomaly Determination

3.1.3.1 Threshold Test

The threshold test selects pixels from the cleaned dataset that fall below the 5th percentile or above the 95th percentile at each gridded location. Traditionally, percentiles and medians are calculated by sorting data by value and then interpolating the rank of the desired

percentile. Since there may be 10,000 or more observations at each of the 14,278,503 possible locations, sorting methods were too computationally expensive. Instead, the percentiles were estimated by creating histograms of the observational values at each location and then selecting the bins that matched the 5% and 95% ranks (and 50% for the medians). Each observation in the cleaned dataset was then compared against these percentile thresholds. If the observation was below the 5^{th} percentile or above the 95^{th} percentile, the pixel was identified as a possible anomaly. All other pixels falling inside the percentile range were discarded. Temporally invariant grid cells, where $\sigma=0$, are automatically excluded using this method.

3.1.3.2 Spatial Adjacency Test

This test checks for spatial adjacency with other anomalies. Specifically, for a candidate anomaly to be "real" it must be adjacent to *n* or more other anomalies. In practice, this convolution acts as an additional noise filter, eliminating spurious outliers that were not caught during the cleaning process. For the MODIS LST anomaly databases, candidate pixels were convolved against a 3x3 grid of surrounding pixels and the value of *n* was tuned to 1, requiring at least one other anomalous pixel to be adjacent to any candidate anomaly.

3.1.4 Creating and Using Anomaly Databases

The anomalous, condensed pixels may be stored in flat files or databases. The MODIS LST anomalies were hosted in a series of PostgreSQL databases, each database containing the condensed time series of LSTs from a single MOD11A1 tile, and further subdivided into databases of low (5^{th} percentile) and high (95^{th} percentile) anomalies. All necessary metadata,

time tags, and locations were also stored in the databases, enabling sophisticated Structured Query Language (SQL) queries. Appendix B provides examples of SQL queries used in the process of creating this dissertation. The baseline descriptive statistics were also stored in similar databases, allowing for sophisticated query cross-referencing. For instance, a user could ask for all pixels within a geographical bounding box, above a specified elevation, at least two standard deviations above or below the mean, and for any desired time interval. The database schema is shown in Figure 3.2.

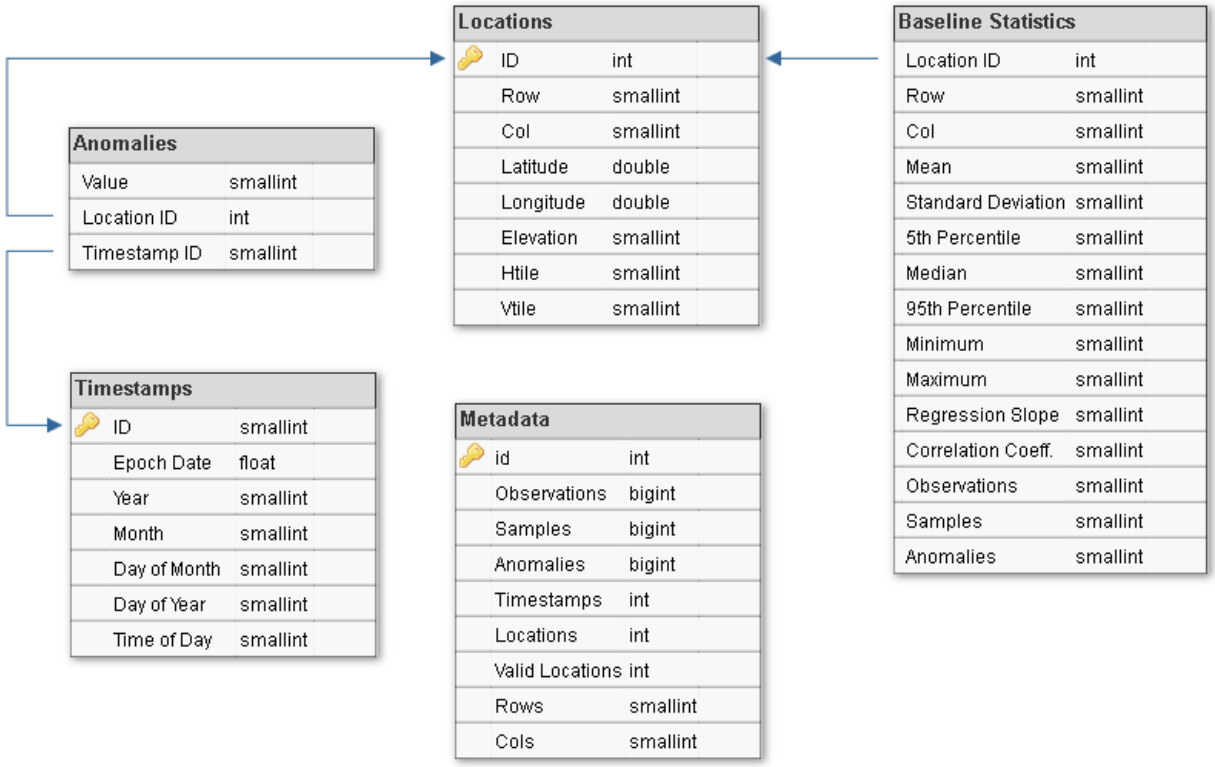

Figure 3.2. The database schema. An individual database contains the anomalies and statistical baselines for one tile and one percentile (5th or 95th), for the designated time period (interannual month, season, or all years).

One limitation of this architecture is the subdivision by MODIS tile. Because there are 23 tiles covering the land surface of Antarctica and two percentile ranges per tile, this produces 46

databases per time period. For seasonal databases, with four seasons per percentile, the result is 184 total databases; for monthly databases, this expands to 552. Administratively, the proliferation of databases poses little problem, and in fact is more efficient in terms of resource utilization: not all databases need to be cached in memory at the same time. If the querying application is intelligent, the subdivision by tile may be transparent to the user. Moreover, if the application only queries the databases needed, the response times are much faster than if all the data was hosted in a single database. For best performance, the database server should handle multi-threaded queries and have a large amount of RAM.

Appendix A provides further details about the database design, as well as performance characteristics, internal formatting, and algorithmic tradeoffs.

3.2 Ancillary Analysis Programs

Although not strictly part of the *Condensate Database Project* as it was originally proposed, two additional programs became necessary tools for the analyses in this dissertation. The first was a simple, command-line-based tool that allows the user to query one or more of the databases using SQL syntax. The tile, time division, and percentile can be specified, and any of parameter may be left to default. Defaulting parameters will result in the program using a wildcard and executing the query across all appropriate databases. The resulting anomalous pixels are returned in a reprojected image of Antarctica, with coastlines, and a comma-separated variable (CSV) file containing the LST values, locations, times, and tile numbers. If multiple anomalies are found at the same location through time, the program can be instructed to generate an image of the mean, min, max, or last LST value at every location. Additionally, the program generates an image displaying the population of LST anomalies at each location.

The second program assists with retrieving the full dataset, not simply the anomalies. As described in previous sections, the condensate databases are a tool for data exploration and quality assurance. Once an event has been found, however, it may be necessary to return to the original source data for further analysis at the location of interest. This second program, also command-line driven, retrieves any time series of LST observations from the original source MOD11A1 dataset for any tile, bounding box, and time span combination. Many of the analyses in Section 4 relied on this capability. On-the-fly dataset cleaning (elimination of cloudy pixels and out-of-range values) is optional.

4 Results and Discussion

The condensate databases, along with their baseline statistics, are tools for data exploration. This section demonstrates the use of these tools to identify areas and events of interest. Where data quality problems are indicated, these too are discussed. The regions and events of greatest interest are then analyzed in greater depth.

4.1 Baseline Statistics

The LST baselines are descriptive statistics of the LSTs at each grid cell. The statistics are generated as a part of the condensation algorithm process, and the imagery produced can reveal important information and provide a starting point for further analyses.

The following sections present a sampling of the available baseline imagery.

4.1.1 Sample Population

The total sample population, following dataset cleaning, is shown in Figure 4.1. For 17 years of twice-daily satellite observations at every grid cell, the theoretical maximum population size at any one grid cell is 12,418. Only in the clear, dry areas of East Antarctica did the sample size approach the maximum, where total samples were around 10,000. Some coastal regions, especially along the Peninsula, were so cloudy that only a few hundred clear samples were observed during the 17 year period. The scant temporal sampling precludes any reasonable analyses in these areas. However, these same regions also tend to be the most heavily sampled using traditional manned and unmanned *in situ* weather observations. The two data collection methods would appear to complement each other, if the differences between the measurement techniques can be reconciled.

The seasonal sample populations are shown in Figure 4.2. For 17 years of twice-daily observations, the theoretical maximum number of samples for any season is approximately 3100. There were substantially more fall (Figure 4.2b) and wintertime (Figure 4.2c) LST observations in East Antarctica. The heightened fall and winter sample populations may be due, in part, to the inherent problems with nighttime cloud detection over ice sheets.

The sample population images reveal two systematic problems in the source MODIS/Terra CDR. First, because the LST data quality is based on cloud detection confidence, and that confidence level drops from 95% to 66% at elevations above 2000m, there is a substantially greater sample population in the high elevation Plateau region of East Antarctica. This discontinuity is apparent in Figure 4.1 where the smaller, low-elevation sample sizes ring the margin of the continent and lager samples are found in the interior. The color scheme used in the image highlights the difference, and the warmer colors (yellows and reds) closely matching the 2000 m contour line (compare with Figure 1.1). Secondly, a reduction in sample population is apparent around the South Pole. The sampling deficiency most pronounced in the summertime, where the sampling distribution forms a kind of target-shaped pattern (Figure 4.2a). This effect may be traced to a combination of factors, including increased solar zenith angles (reducing cloud detection effectiveness), oblique viewing angles, reprojection artifacts, and natural topography. Speculatively, oblique viewing angles are the most likely cause. NASA's MODIS processing preferentially selects observations closest to nadir. However, the orbital inclination of the MODIS/Terra satellite is 98.2°, and at no time does the satellite ever pass directly over the pole. The oblique viewing geometry increases the atmospheric optical depth, which may increase cloud detections and result in frequent pixel rejection. During the dark

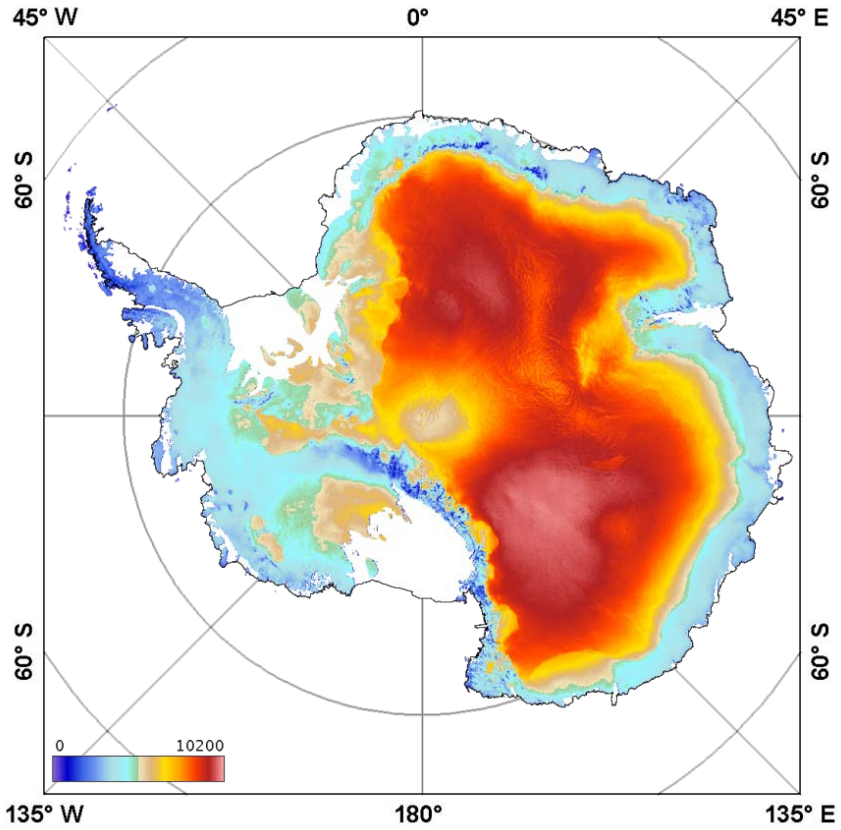

Figure 4.1. Sample population for the cleaned MODIS LST CDR, using twice daily observations, 2000-2017.

winter months, when the summertime near-IR cloud detection tests are unavailable, the sampling rate picks up again – but at the expense of increased cloud contamination [Ackerman et al., 1998].

4.1.2 Percentiles

The 5th and 95th LST percentiles were used as initial thresholds for determining anomalously low and high temperatures (Figure 4.3). While not as extreme as the absolute minima and maxima, the percentiles are much less influenced by cloud contamination. Various

surface features are easily recognizable based on their temperature differences with surrounding areas. Interestingly, the region grid-southeast from the South Pole, called Titan Dome, shows 5th percentile LSTs nearly as cold as East Antarctic Ridge, which is much higher in elevation.

4.1.3 Median, Mean, and Standard Deviation

The median, mean and standard deviations of the interannual baselines are shown in Figure 4.4 and Figure 4.5. As expected, the coastal areas have the highest mean temperatures with the lowest variance, exhibiting the strong maritime influence of the weather patterns. The high-elevation Plateau regions have the lowest mean temperatures and the highest variance. The median tends to be cooler than the mean at most points, likely due to warm air masses positively skewing the distribution of surface temperatures (see Section 4.3.1 for more information on the data distributions and skewness).

The seasonal means and standard deviations are shown in Figure 4.6 and Figure 4.7. Most areas experience a relatively low seasonal variance except for the high elevation Plateau region during the austral spring (Figure 4.7d). The Transantarctic Mountains show the highest temperature variance of any region, occurring during the winter season (see inset in Figure 4.7c). The large variance suggests cold, katabatic winds from the Plateau may be alternating with warm maritime air masses infiltrating across the Ross Ice Shelf.

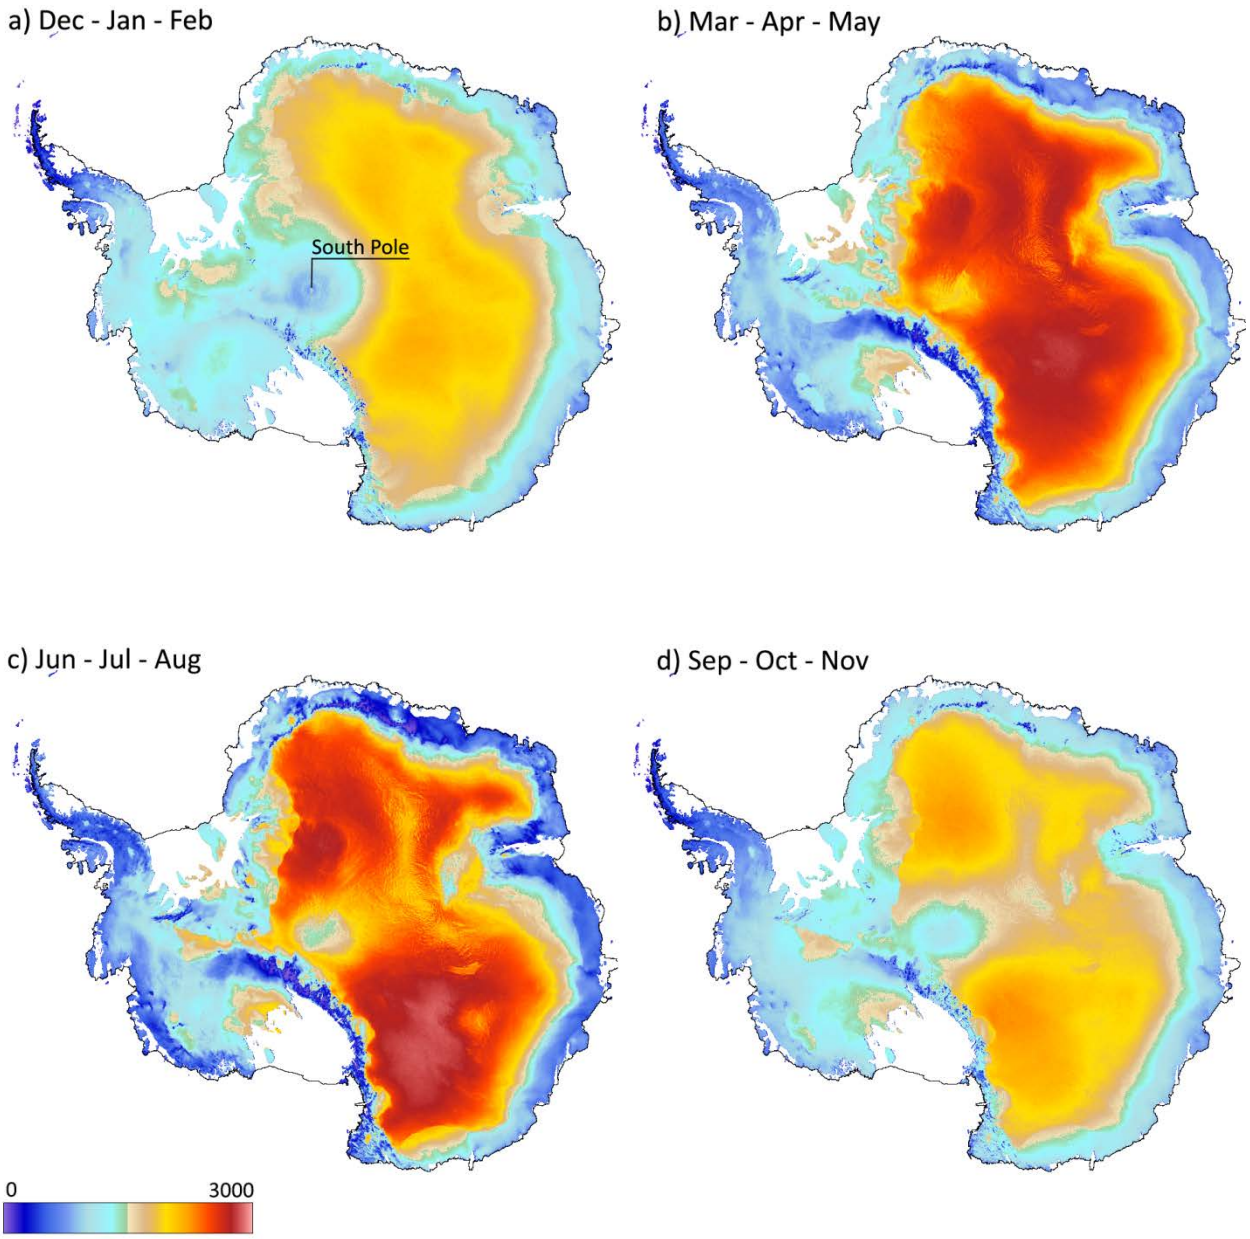

Figure 4.2. Seasonal sample populations for MODIS LSTs, from twice daily observations, 2000 – 2017.

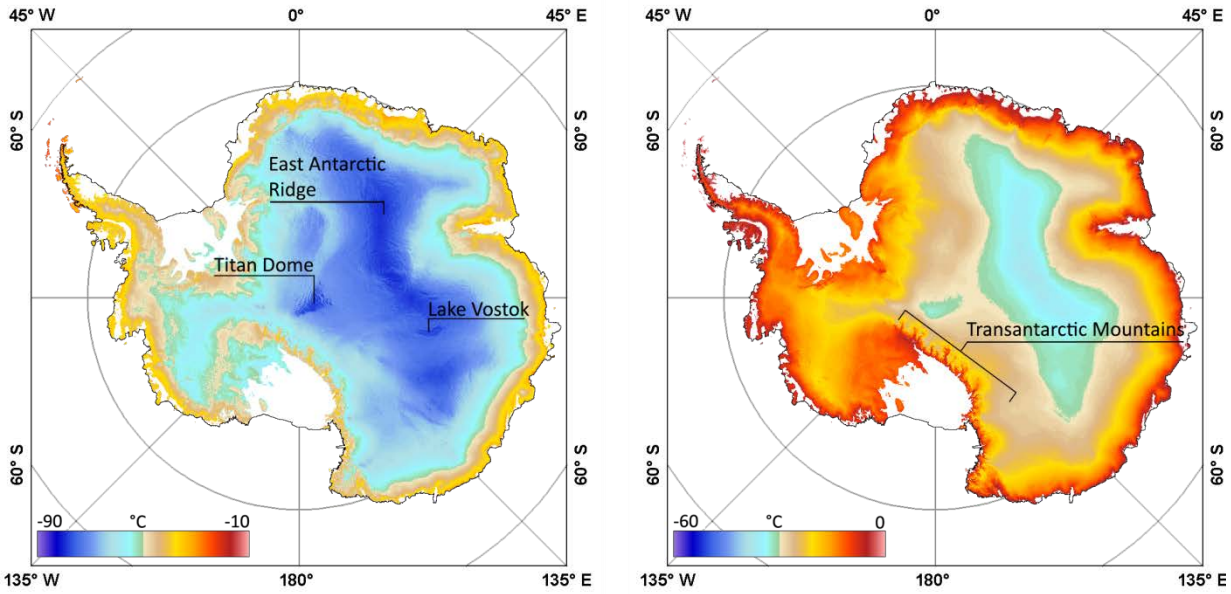

Figure 4.3. The 5th (left) and 95th (right) percentiles of land surface temperatures, using all dates in the full 17 year dataset. Similar percentile baselines were generated for interannual months and seasons (not shown). Some surface features are identifiable based on temperature signatures.

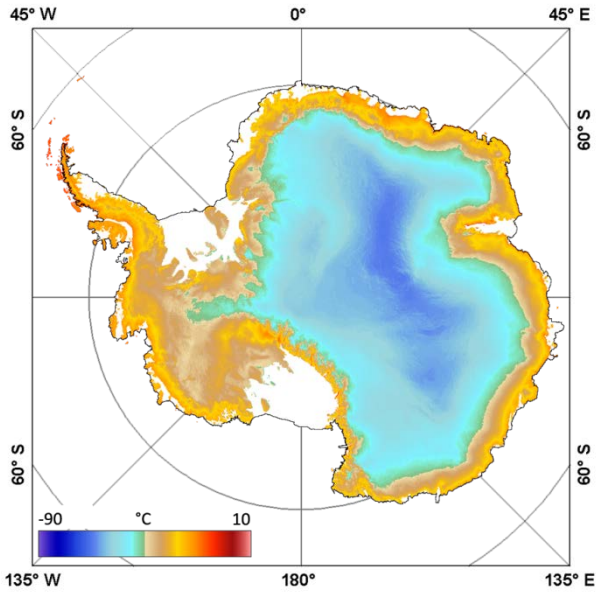

Figure 4.4. The median interannual land surface temperatures, 2000-2017, in degrees C.

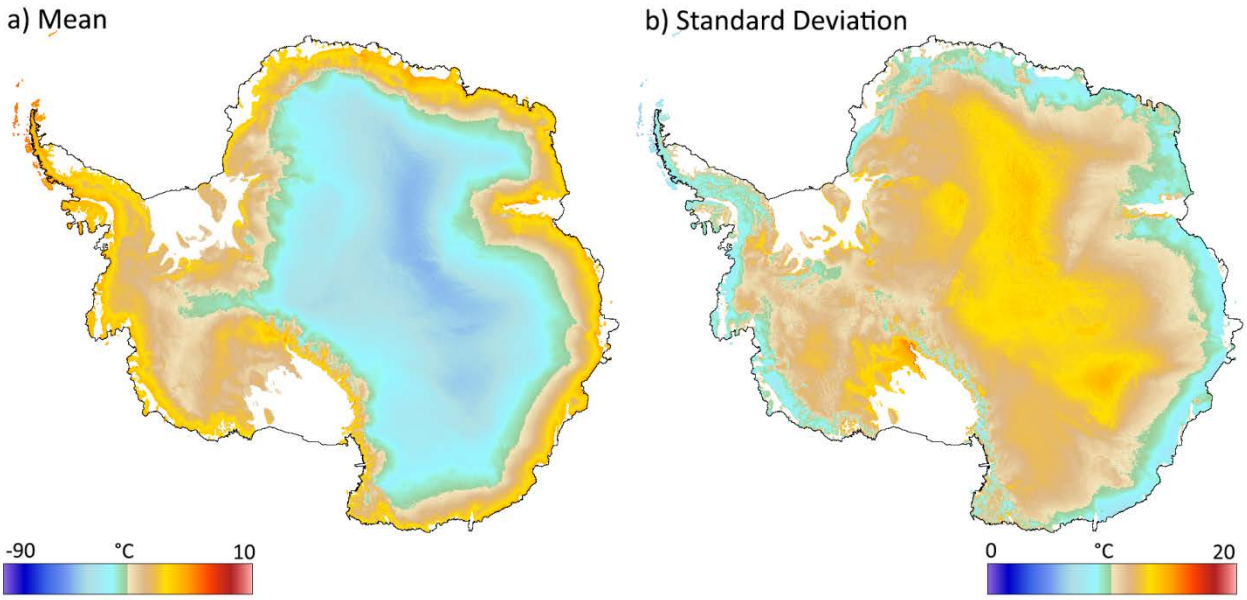

Figure 4.5. The mean interannual land surface temperatures for 2000-2017 (a) and the standard deviations (b).

4.1.4 Linear Regressions

Linear regression calculations are not a necessary part of the condensate process or anomaly algorithms; however, the baseline statistical data processing provides an opportunity to produce simple Ordinary Least Squares (OLS) linear regressions at each grid point (Figure 4.8). Users of these images are cautioned that the sample populations may not be normally distributed, homoscedastic, or that the sampling is evenly distributed through time. Linear models are therefore unlikely to be the best fit or statistically significant. A better model would use customized functions at each location to account for uneven sampling distributions, transformations to adjust for non-normality and/or heteroscedasticity, and statistical

significance checks. These latter types of transformational methods are computationally expensive and generally require manual intervention, and were thus infeasible on the massive scales used in this study. Nevertheless, the simple regressions provide a quick visual tool for exploring possible trends. For instance, Figure 4.8a suggests that much of the continent may be warming. In general, the correlation coefficients (Figure 4.8b) at each location are less than 0.3, and in most cases substantially smaller, suggesting that the linear regression models are not effective at capturing the trends and perhaps not statistically significant. The low correlation coefficients may also be attributed to the short timeline and inherently noisy dataset. Subsequent analysis (Section 4.5.4) has also revealed a potential discontinuity in cloud detection, potentially contributing to the apparent upward temperature trend.

A few regions show interesting phenomena worthy of further investigation. When viewed seasonally (Figure 4.9), East Antarctica appears to be warming substantially during the Austral spring (Figure 4.9d). In particular, the mountainous areas surrounding the Amery Ice Shelf exhibit some of the greatest warming and largest correlations with time. In contrast, the South Pole is the sole area with a negative slope (cooling), albeit with very low correlation coefficients. The possibility of a negative LST trend near the South Pole is explored in greater depth in Section 4.5.1.

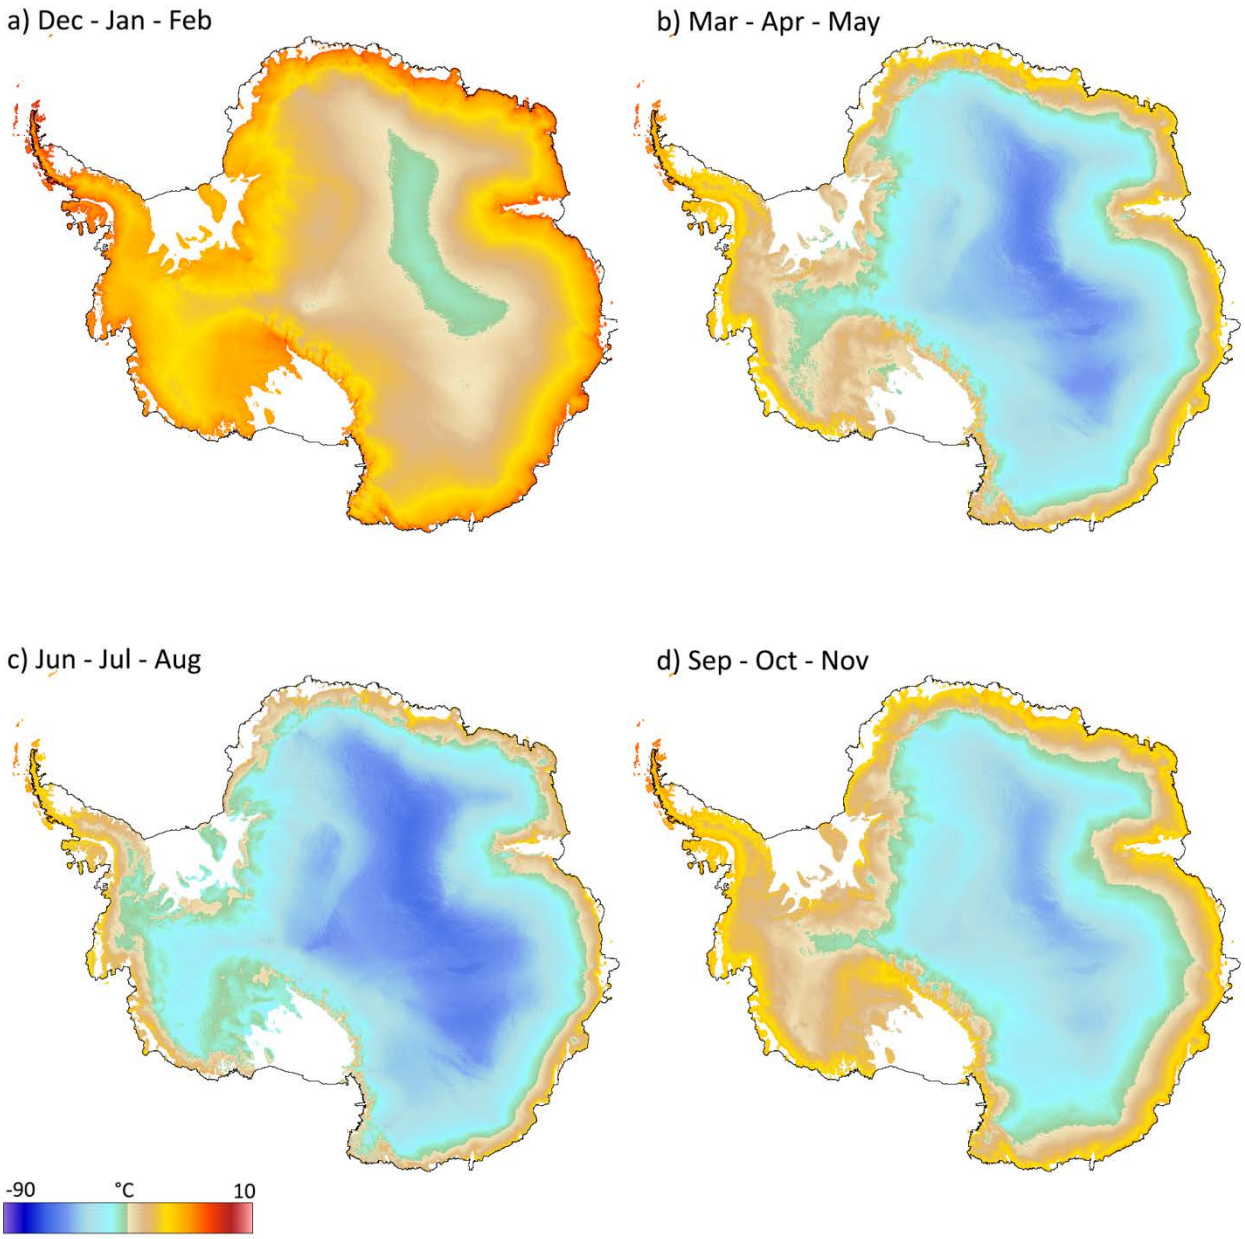

Figure 4.6. Interannual seasonal mean land surface temperatures, 2000 - 2017.

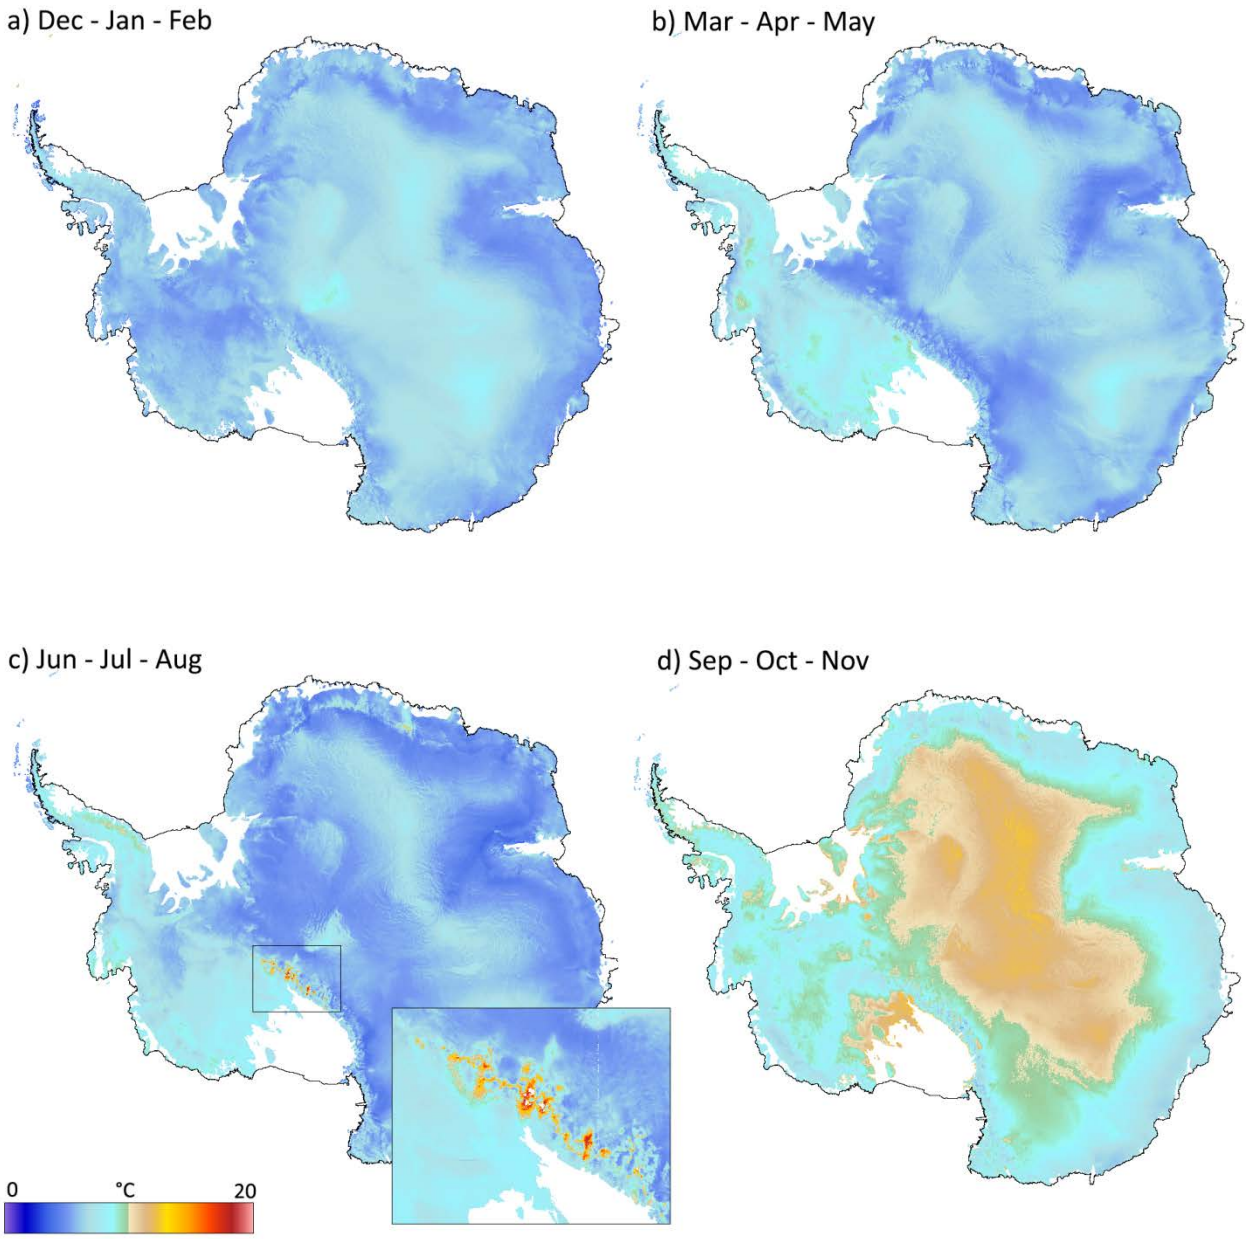

Figure 4.7. Seasonal standard deviations of land surface temperatures, 2000 – 2017. The winter season, Panel (c), includes a zoomed inset of the Transantarctic Mountain Range.

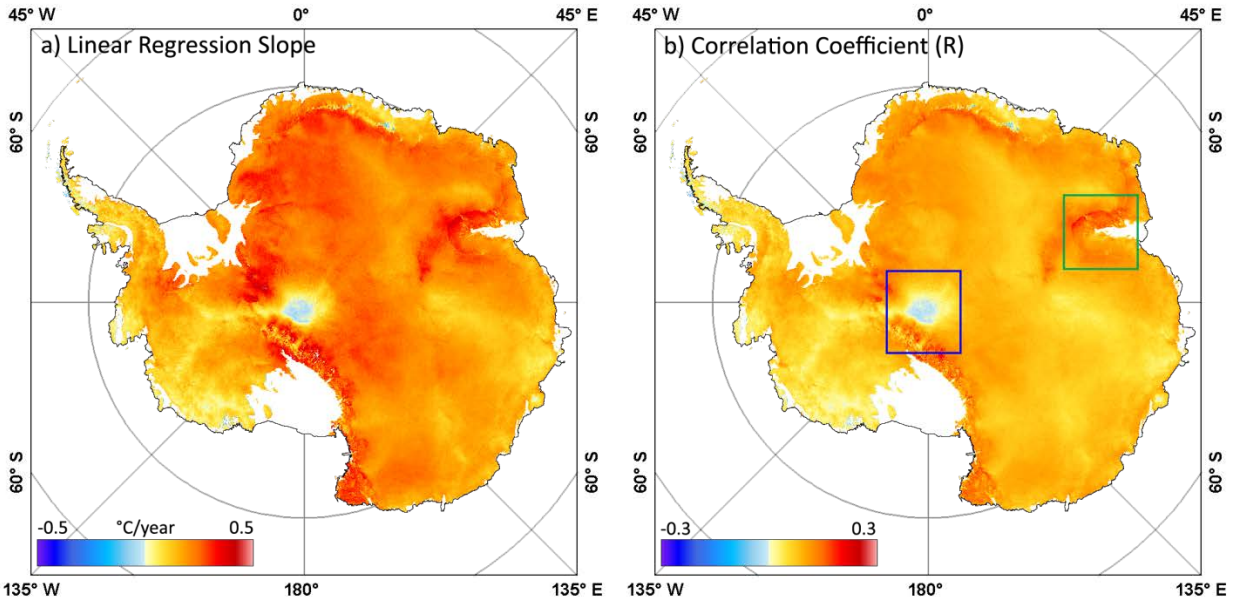

Figure 4.8. Slopes (a) and correlation coefficients (b) of the interannual OLS linear regressions of LSTs versus time, 2000 – 2017. In Panel (b), the green box identifies the Amery Ice Shelf region, and the blue box the South Pole.

4.1.5 Absolute Maxima and Minima

The absolute maximum and minimum temperatures, by season, are shown in Figure 4.10 and Figure 4.11. Because these temperatures represent only a single observation at any location, and without the benefit of the adjacent anomaly filtering algorithm, they are highly susceptible to cloud contamination errors. Cloud contamination is especially apparent during the winter season when accurate cloud detection is difficult (e.g., Figure 4.10c). Although the absolute temperatures may be equivocal, the images provide a rapid method identifying the most extreme temperatures around the continent and potentially the coldest and warmest locations.

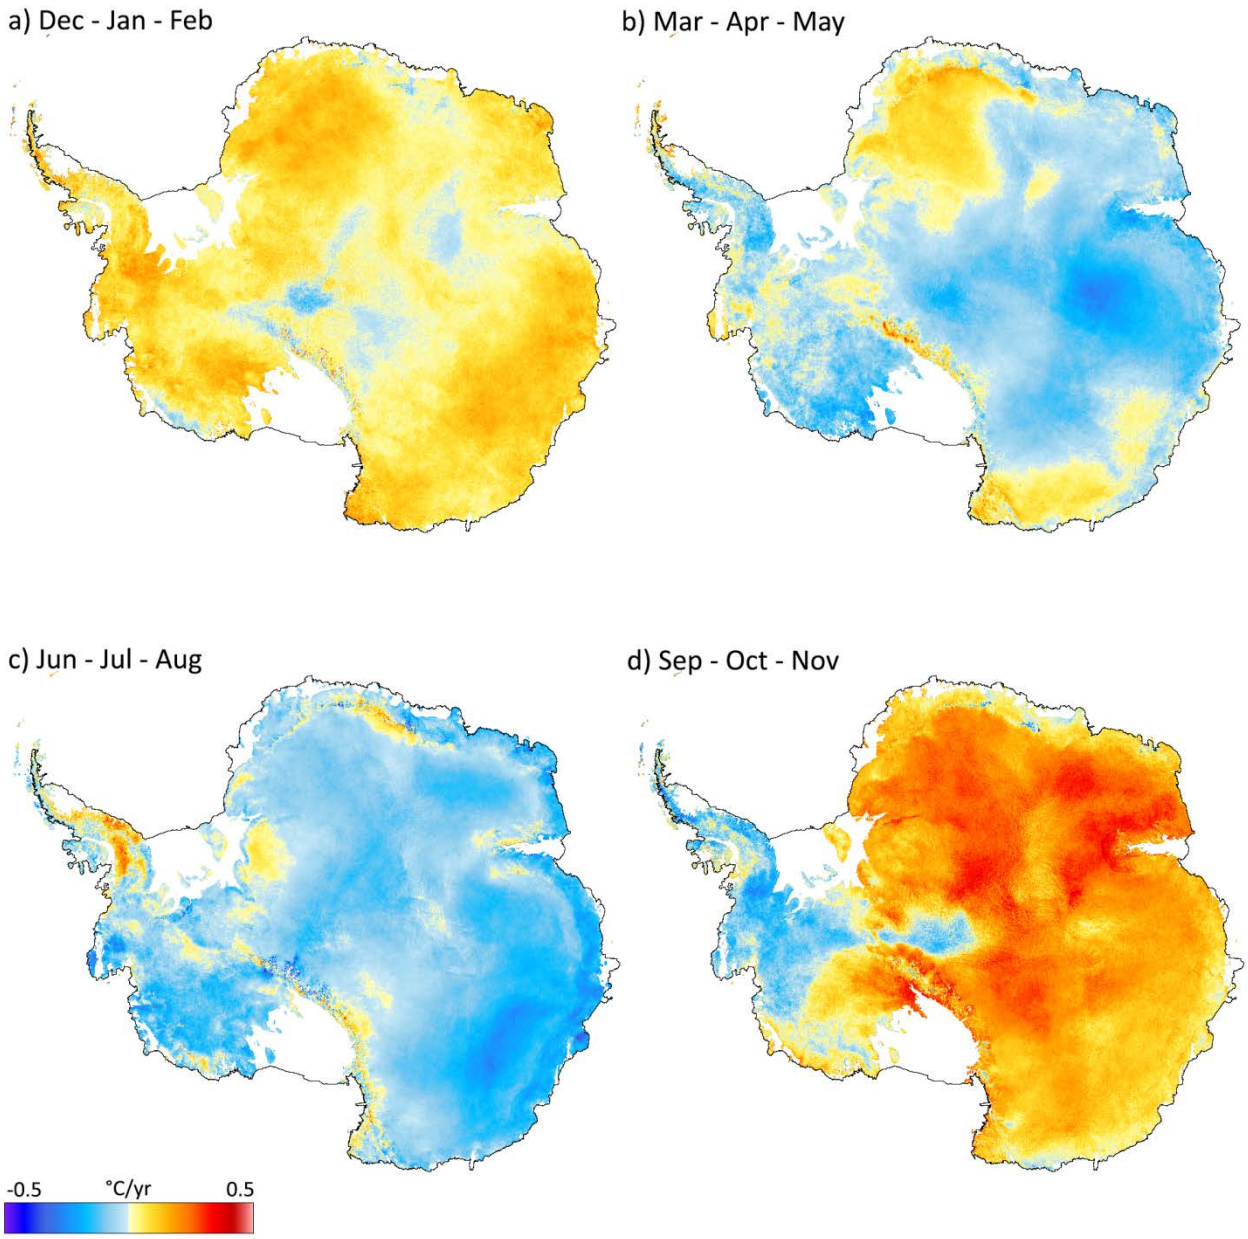

Figure 4.9. Linear regression slopes by season, 2000 – 2017, in degrees C per year.

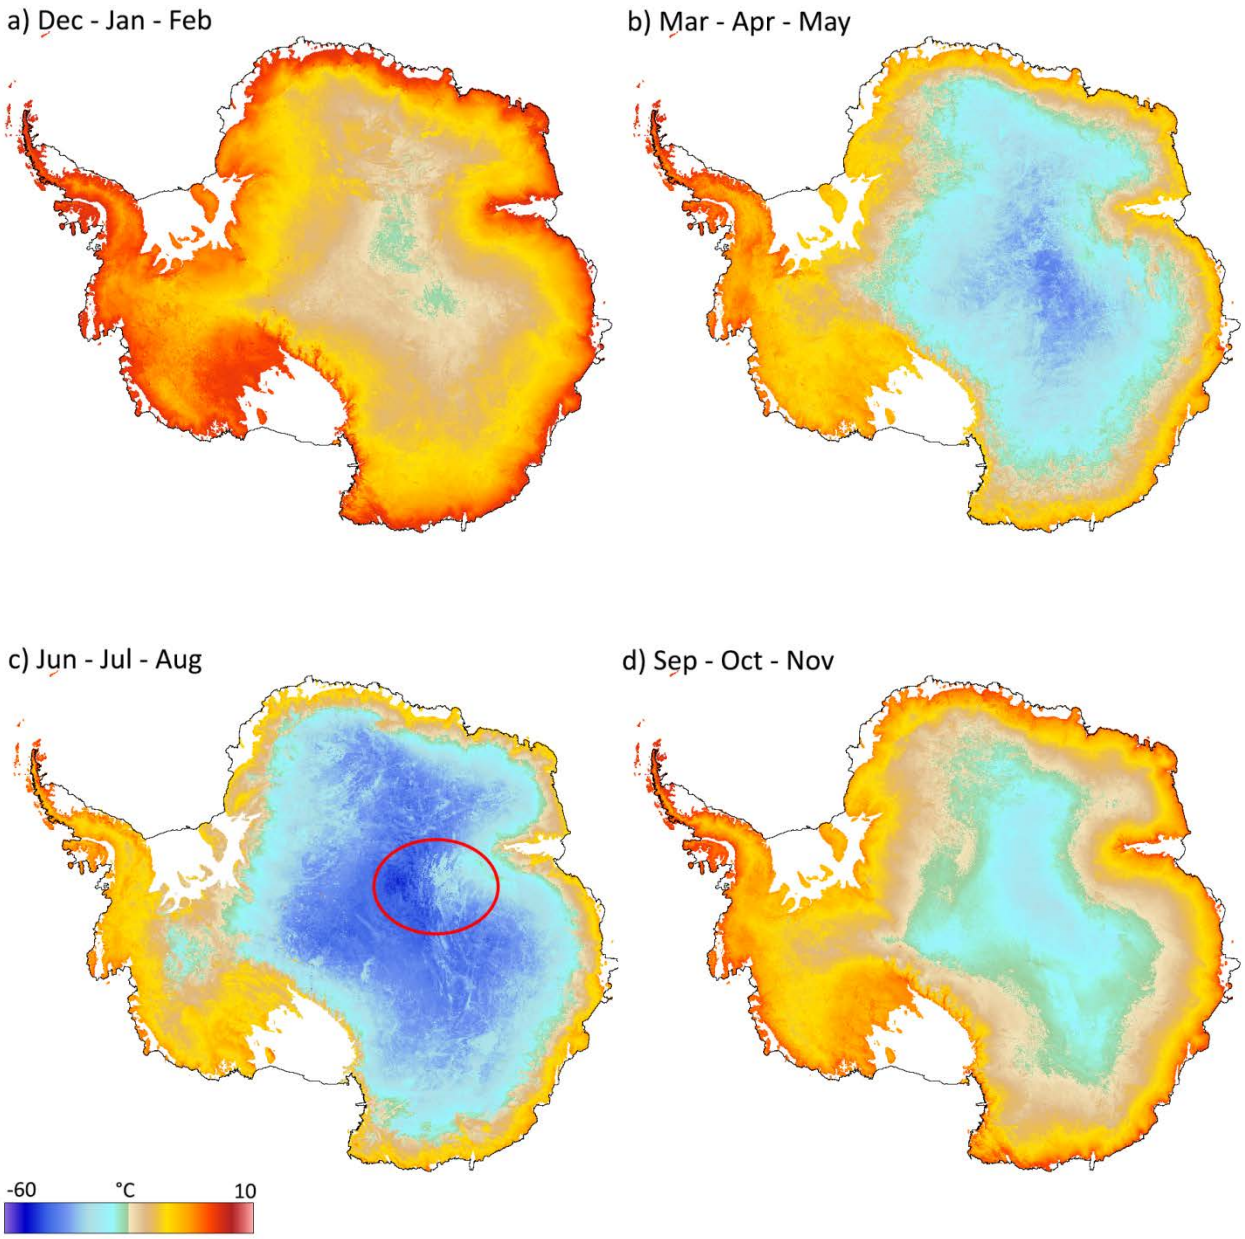

Figure 4.10. Absolute maximum land surface temperatures, by season, 2000 – 2017. Cloud contamination is visually evident in the data, especially during June-July-August. The red circle in Panel (c) identifies irregular data caused by cloud contamination.

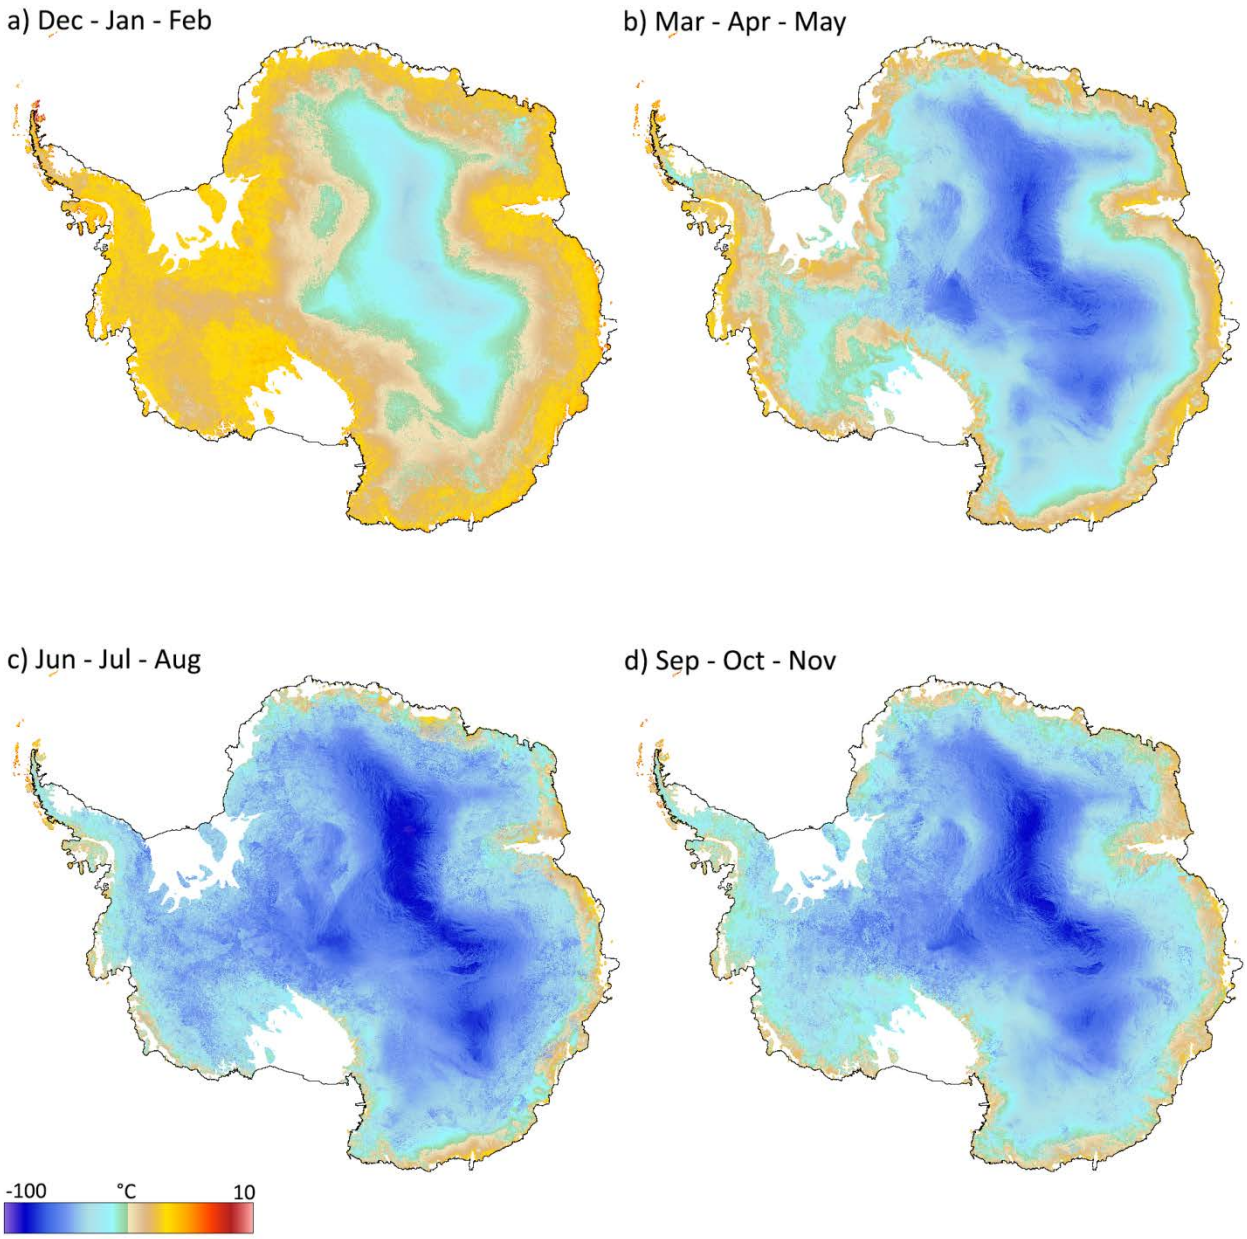

Figure 4.11. Absolute minimum land surface temperatures, by season, 2000-2017.

4.2 Spatial Patterns of LST Extremes

4.2.1 Highest Temperatures

Figure 4.12 shows the 95th percentile of LSTs for the summer season of December through February, typically the warmest months in Antarctica. Argentinian station Esperanza, the current high temperature record holder for the low-elevation Antarctic sub-region, is shown at the northern tip of the Antarctic Peninsula (Figure 4.12a). Esperanza set the current temperature record, 17.5 °C, on 24 March 2015. Visible images show the area to be overcast on the record-setting date, and indeed no LSTs were observed in the area on that day. Analysis indicates that the high temperature was observed during a period of warm air advection, combined with warming from foehn winds from the leeward side of the Peninsula Mountains [Skansi et al., 2017]. On the tip of the Peninsula, warm air advected from the Southern Ocean is typically moist and cloudy, explaining the overcast conditions. This event exemplifies a limitation when using LSTs to detect temperature extremes: In some regions, the same conditions that result in the most extreme temperatures often prevent direct satellite observation. It is highly likely that other, even higher temperatures have occurred, yet they remain unknown due to the lack of *in situ* instrumentation.

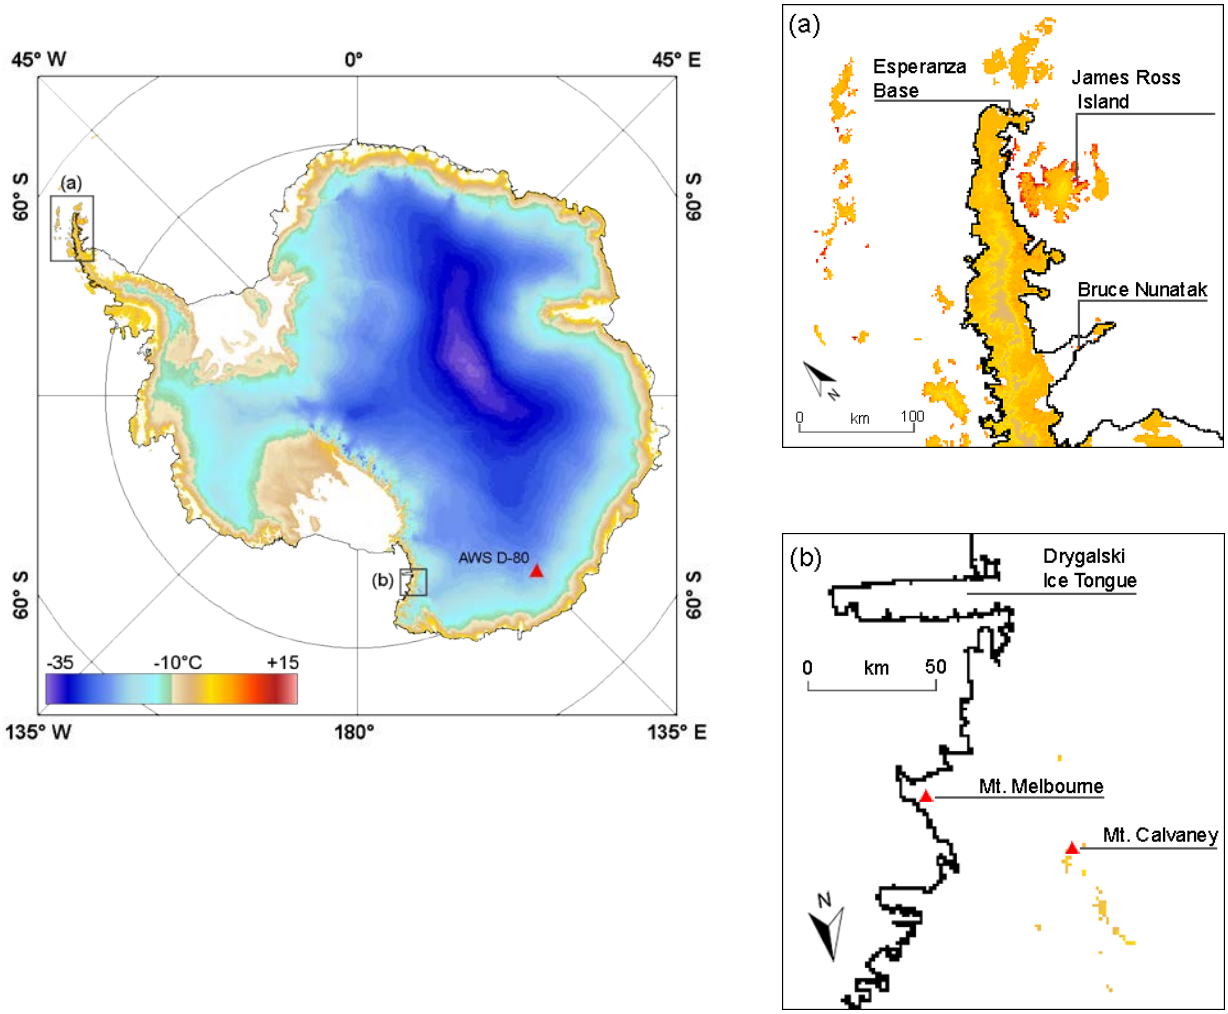

Figure 4.12. The 95th percentile LSTs for the austral summer, December-January-February, 2000-2017. Inset (a) shows the locations of extremes along the Antarctic Peninsula. Inset (b) shows a portion of the Ross Sea coastline; white areas are either ocean or below 2500 m elevation, and yellow-orange pixels are anomalously high LSTs above 2500 m elevation.

From the perspective of a satellite thermal sensor, some of the warmest temperatures occur not at the northern tip of the Peninsula but farther south, where the surface type and aspect angle combine to warm the earth through efficient sunlight absorption. Combined with the latent heat release from strong foehn winds, these conditions may result in extreme warm temperatures. James Ross Island (64.17S, 57.75W) has the most consistent high temperatures, especially along the coastal areas exposed to the north (Figure 4.12a). The island is partially

glaciated but much of the coastline has seasonally exposed bare soil and rock. The exposed rock and soil has a very low-albedo. Observed LSTs frequently rise above 10 °C during the summer time. Following five days of generally clear skies, on 08 February 2017 the east coast of James Ross Island reached an LST of 15.1 °C. MODIS true-color reflectance imagery confirms cloud-free conditions. While this temperature does not exceed the Esperanza record (17.5 °C), it suggests that higher temperatures may be possible. Slightly farther south is Bruce Nunatak (65.08S, 60.25W), part of the Seal Nunataks chain. Bruce Nunatak juts from an unnamed island adjacent to the northern edge of the former Larsen B ice shelf, where north-facing steep-sided bluffs of dark rock receive a considerable insolation. Using the LST databases, two notable extreme events were observed here: First, on 05 January 2006, the land surface temperature reached 17.4 °C. Then on 01 January 2009 the satellite observations exceeded that of Esperanza's air temperature record, with an LST of 17.8 °C. Given that LSTs are commonly 2 to 3 °C less than the 2 m air temperatures, it is possible that T_{air} at this location may sometimes exceed 20 °C. Visible imagery confirms that the sky was clear during both extreme events (not shown).

Because the cloudy conditions in low elevation coastal areas may preclude observations of extreme high temperatures, this dissertation sought the most likely areas above 2500 m where new records may occur frequently but remain unrecorded. This search was easily performed using simple queries against the anomaly databases (see Appendix A for examples).

The record high air temperature above 2500m, −7.0 °C, was set on 28 December 1989 at an automated weather station called D-80 (70.0S, 134.7E). Data indicates that both warm air advection and clear-sky solar heating played roles in the D-80 record [Skansi et al., EOS 2017].

The drier conditions on the plateau, with clearer skies and greater insolation relative to the coasts, suggest that record extreme warm temperatures at higher elevations may be more observable using satellite-based sensors. However, data exploration using the condensate databases revealed that most of the extreme high LST temperature observations were in fact cloud contaminated.

One area above 2500 m stands out as a frequent site of clear-sky extreme land surface temperatures: Inland and northwest from Mt. Melbourne, along the Ross Sea coastline of the Transantarctic Range, is a high elevation area of north-aspect mountain slopes (Figure 4.12b). This region exhibits frequent extreme warm temperatures rivaling that of the D-80 record. One such extreme event is further described in the Case Studies part of this dissertation, Section 4.5.3.

4.2.2 Lowest Temperatures

Recent studies have shown that Antarctica's most extreme cold temperatures are associated with prolonged periods of clear weather, when longwave heat loss can continually cool the surface [Campbell et al., 2013; Haran et al., 2016]. A working hypothesis is that strong surface inversions and local topography then combine to concentrate pockets of cold air at topographic low points. The conditions along the East Antarctic Divide, a high ridge between Dome Argus and Dome Fuji, are ideal for creating these winter extremes. Figure 4.13 shows the 5^{th} percentile threshold for winter (JJA) LST extremes, with local topography clearly visible due to these pockets of cold air. Vostok Station itself, with the official record cold temperature of −89.2 °C, is situated downslope from Dome Argus and collects cold air in the low basin created by subglacial Lake Vostok. Data exploration using the anomaly databases indicates that extreme

cold temperatures are common near the East Antarctic Ridge, and often colder than the Vostok Station record, with land surface temperatures frequently falling below −90 °C. One extreme cold event case study is detailed in Section 4.5.2.

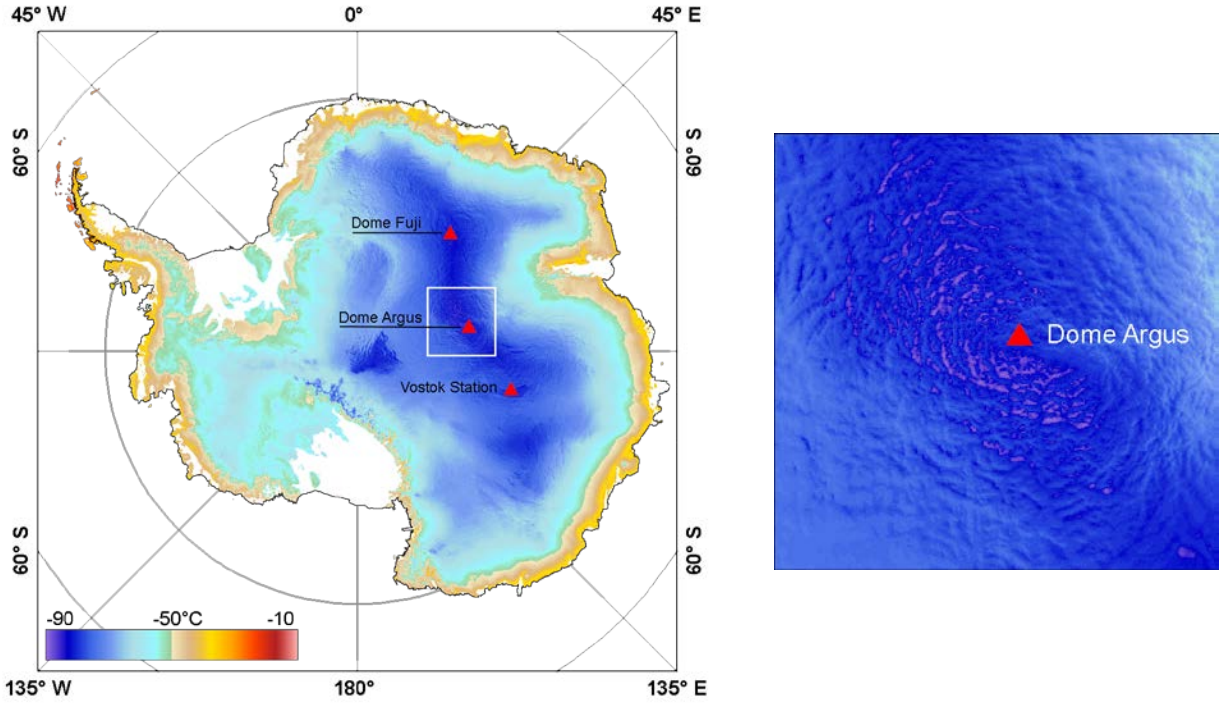

Figure 4.13. The lowest 5th percentile (left) of LSTs for JJA 2000-2016. The East Antarctic Ridge, site of many extreme cold temperatures, runs from Dome Fuj to Dome Argus. Cold air collects at low topographic areas, revealing the surface texture near Dome Argus in a zoomed image (right).

The East Antarctic Ridge and domes are part of the high elevation Plateau region of East Antarctica. The ridge is entirely above the 2500 m threshold for WMO's high elevation temperature regime. Although low temperature LST extremes below 2500 m are not thoroughly explored in this dissertation, it appears that feeder glaciers to the Ross Ice Shelf, at the extreme southern part of the Transantarctic Range, may experience the lowest temperatures at lower elevations. The extremes are likely caused by cold katabatic winds

flowing downslope from the Plateau. The low temperatures and high variances in the Transantarctic Range are visible in Figure 4.7c and Figure 4.13.

4.3 Temporal Patterns of LST Extremes

This analysis seeks answers to two basic questions: First, are LST extremes becoming colder or warmer over time? And is the frequency of extremes changing? Since the LST data is not detrended, and the percentile thresholds remained static during this study, changes to the anomaly temperatures and populations may be visible through time. First, we look at the yearly LST anomaly medians, medians, and standard deviations.

Queries against the anomaly databases returned all anomalies for individual years, and the anomalies were then assembled into graphs of descriptive statistics through time. Figure 4.14 shows the medians, means, and standard deviations of the anomalies by year, using full

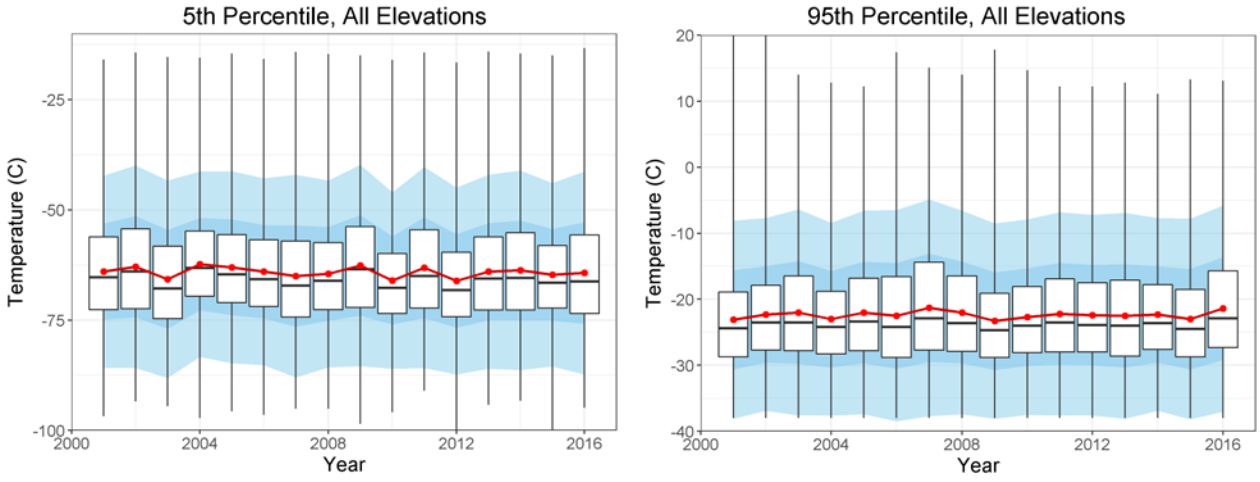

Figure 4.14. Box and whisker plots of the anomalous LSTs for the lower 5th percentile (left) and upper 95th percentile (right) by year, for all elevations. Also shown are the means of the anomalous LSTs (red points and lines) and the variance around those means (dark blue is one standard deviation, light blue is two standard deviations).

years from 2001 through 2016; the data used for the figure includes all elevations and locations, resulting in a large variance.

For both the lower and upper percentiles, the mean temperature is generally 1° to 2° C higher than the median, reflecting greater variance in the upper end of the temperature range. The variance is likely due to warm fronts intruding into the continent, with some cloud contamination. The cold tails of the extremes tend to vary less. As a result, the temperature distributions are positively skewed in both the upper and lower percentile anomalies.

Table 1. A search for temporal trends in the 5th and 95th anomaly percentiles. Linear regressions for the medians of the anomalous LSTs, by percentile, elevation, and season are shown. None of the regressions are statistically significant for $\alpha = 0.05$ or $\alpha = 0.10$.

Elevation	Season	Median of 5th Percentile, Slope: °C/year	R^2	p (α = .05)	Median of, 95th Percentile, Slope: °C/year	R^2	P (α = .05)
<= 2500 m	DJF	0.00	0.03	0.92	-0.13	0.11	0.18
	MAM	-0.03	0.01	0.79	-0.07	0.90	0.24
	JJA	0.04	0.02	0.58	0.06	0.08	0.28
	SON	0.00	0.00	0.99	0.01	0.00	0.91
> 2500m	DJF	0.03	0.02	0.58	0.07	0.13	0.14
	MAM	-0.10	0.11	0.18	0.05	0.16	0.12
	JJA	-0.11	0.11	0.19	0.03	0.12	0.17
	SON	-0.10	0.23	0.25	0.02	0.03	0.48

A look at linear regressions of the medians through time show that the 5th percentile has a negative slope ($\beta_1 = -0.09$ °C/year) and the 95th percentile has a slightly positive slope ($\beta_1 = 0.01$ °C/year), suggesting that the warm extremes are slowly becoming warmer while the cold extremes are becoming colder (not shown). None of the regression slopes are statistically significant; in all cases, p >= 0.28, and is often above 0.5. The same test applied to the interannual means produces similar, inconclusive results. Likewise, no significant trends were found when the medians were subdivided by elevation and season (Table 1). Overall, the dataset appears to be too noisy and the timeline too short to conclude there are any

meaningful extreme LST value trends, at least when using the current methodology and without further subdividing the continent.

For another perspective on the extreme LST time series, here we examine the sample populations of "clear sky" (ostensibly cloud-free) observations that fall within the 5^{th} and 95^{th} percentiles (Figure 4.15). The noisiness of the dataset is clearly visible here, with substantial interannual variance. Interpretation of this data also is confounded because the anomaly sample populations are necessarily a product of multiple variables and parameters: The surface temperatures, percentile temperature thresholds, the number of grid cells on the Antarctic continent, and the population of cloud-free observations in the source data. Fortunately, byproducts of the condensate baseline generation process include a population count of the total cloud-free observations for a given time span (e.g., seasonal or yearly), without regard to the percentile thresholds, and the total number of valid grid-cells (i.e., land surfaces as opposed to ocean or space). Using this information, the percentile sample populations can be decomposed into the fraction of total land surface observations that were cloud-free, and the percentage of those observations that met the percentile threshold criteria. These components can answer two basic questions: 1) Is the observed cloud cover changing over time, and 2) Within the cloud-free observations, is the percentage of surface temperatures that exceed a percentile threshold changing?

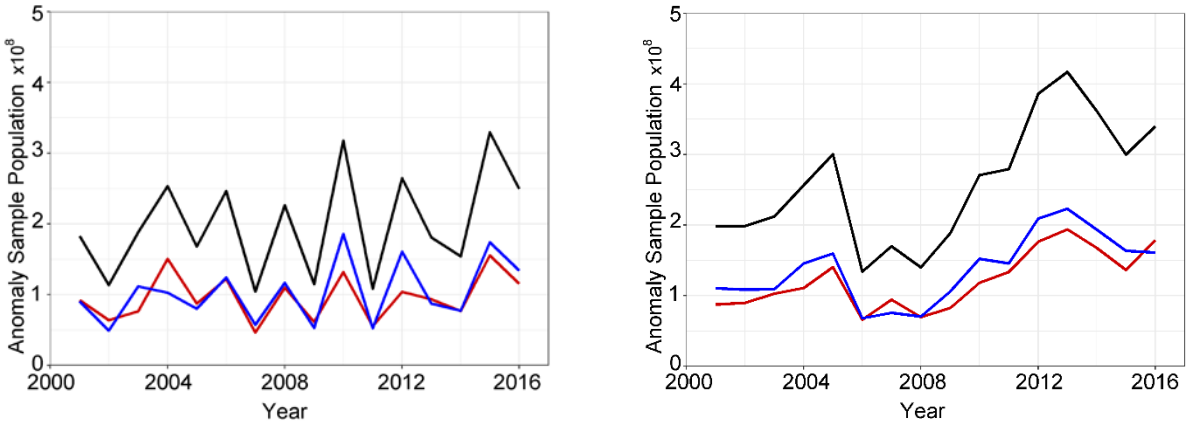

Figure 4.15. The population of LST anomalies each year, 5th percentiles (left) and 95th percentiles (right). Elevations greater than 2500 m are shown in blue, below 2500 m in red, and all elevations combined in black.

The fraction of cloud-free observations (OBS_{clear}) is found by normalizing the number of clear observations each year (N_{clear}) to the total Antarctic land surface grid cells observed each day (P_{total} = 28557006 for twice-daily observations) and the time span used (e.g., t = 365 or 366 days for the yearly databases):

$$OBS_{clear} = (N_{clear} / P_{total}) / t \qquad (Q)$$

A graph of the fraction of yearly cloud-free observations is shown in Figure 4.16. Aside from a jump of approximately 7% starting in 2011, the fraction of cloud-free observations is remarkably unvarying. The sudden increase, however, may reflect a data quality problem. Section 4.5.4 takes a deeper look at this topic.

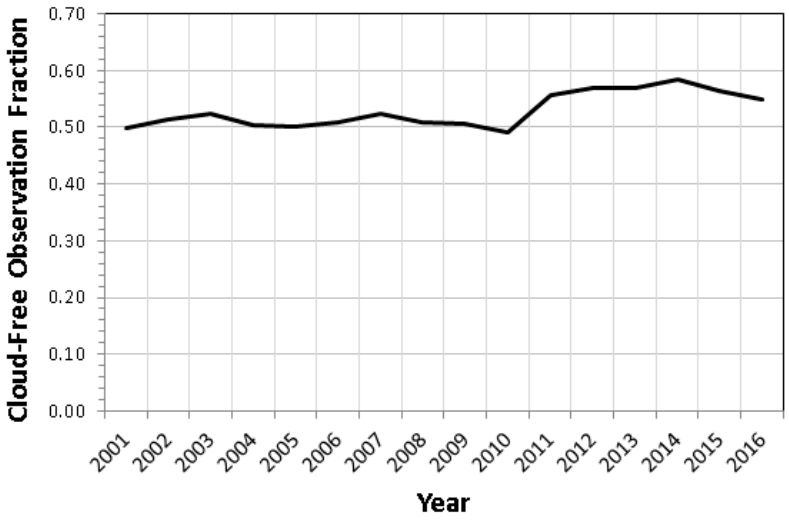

Figure 4.16. The fraction of all land surface observations that were cloud-free, by year.

Dividing the anomaly sample population (from Figure 4.15) by the total number of cloud-free observations yields the percentage of cloud-free observations that met the condensate algorithm's anomaly percentile thresholds and adjacency test (Figure 4.17).

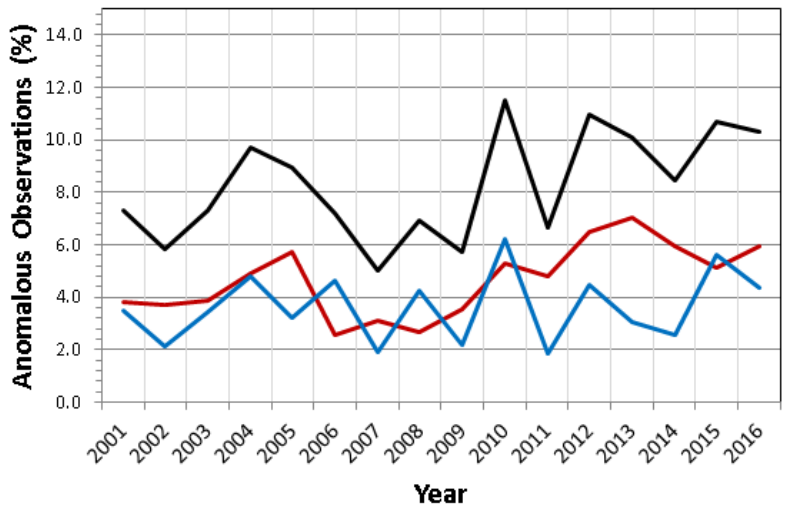

Figure 4.17. The percentage of cloud-free observations that were determined to be anomalous. The blue line is anomalies in the 5th percentile, the red line is the 95th percentile, and the black line is both combined.

In Figure 4.17 the high interannual variability of extreme LST sample populations is again visible, but with the variations due to cloud cover removed. Extreme conditions appear more pronounced in some years than others, and the figure suggests the likelihood of an "extreme year" may be independent of the cloud coverage. For instance, the number of anomalous observations took a sharp upward jump in 2010, with the number of 5^{th} percentile observations tripling compared to the previous and following years, yet the total number of cloud free observations (see Figure 4.16) was slightly less than the preceding years. Similar but less dramatic jumps in the number of extreme LSTs occurred in 2004 and 2012.

The percentage of anomalous observations can also act as a benchmark for evaluating the efficiency of the condensation process. In this case, the lowest 5^{th} percentile anomaly database contains 3.6% of the cleaned, cloud-free source observations, and the 95^{th} percentile database contains 4.7%. Although the use of 5^{th} and 95^{th} percentile thresholds would normally imply that exactly 5% of all cloud-free observations should be included in each percentile's database, the sample population is reduced because of the addition of the adjacent anomaly filter (see Sections 3.1.3.1 and 3.1.3.2).

In summary, the year-to-year anomaly sample populations have a high variance, hampering any meaningful conclusions. Statistically significant changes over time are not apparent. Similarly, changes to the median and mean LSTs in both the 5^{th} and 95^{th} percentiles are virtually zero and trends are not statistically significant. When factoring-in elevation and season, LST anomaly trends are still unremarkable.

These results may appear to run counter to the exploratory baseline linear regressions shown in Section 4.1.4. However, the baselines represent the entire cloud-free LST population, not only the 5th and 95th percentiles. The implication is that, while the extreme percentiles (both temperatures and sample populations) show no significant changes, the center 90% of LST samples may still be changing.

4.3.1 Skewness

Initially, the condensate database algorithm determined LST anomalies by using a threshold of two standard deviations from the mean. The analysis phase of this dissertation, however, found that some regions of Antarctica had very large LST variances with strongly skewed distributions. In most of East Antarctica, the full-year interannual LST data distribution is positively skewed, where the mean temperatures are heavily influenced by warm outliers. Using the two standard deviation threshold, the skewness in these areas often resulted in few anomalously low data points (or none at all). The problem was especially prevalent at the highest elevations, where the coldest temperatures were common. Figure 4.18 examines one such location, near Lake Vostok, that typifies the problem. This finding motivated a reassessment of the condensate algorithms and a redefinition of extreme temperatures. The result was a change to using percentiles rather than standard deviations as criteria for determining extremes, ensuring that the lowest and highest temperatures at any location would still be identified as extremes. The mean and standard deviation statistics, however, were still deemed useful for analytical purposes, and retained in the databases' statistics tables.

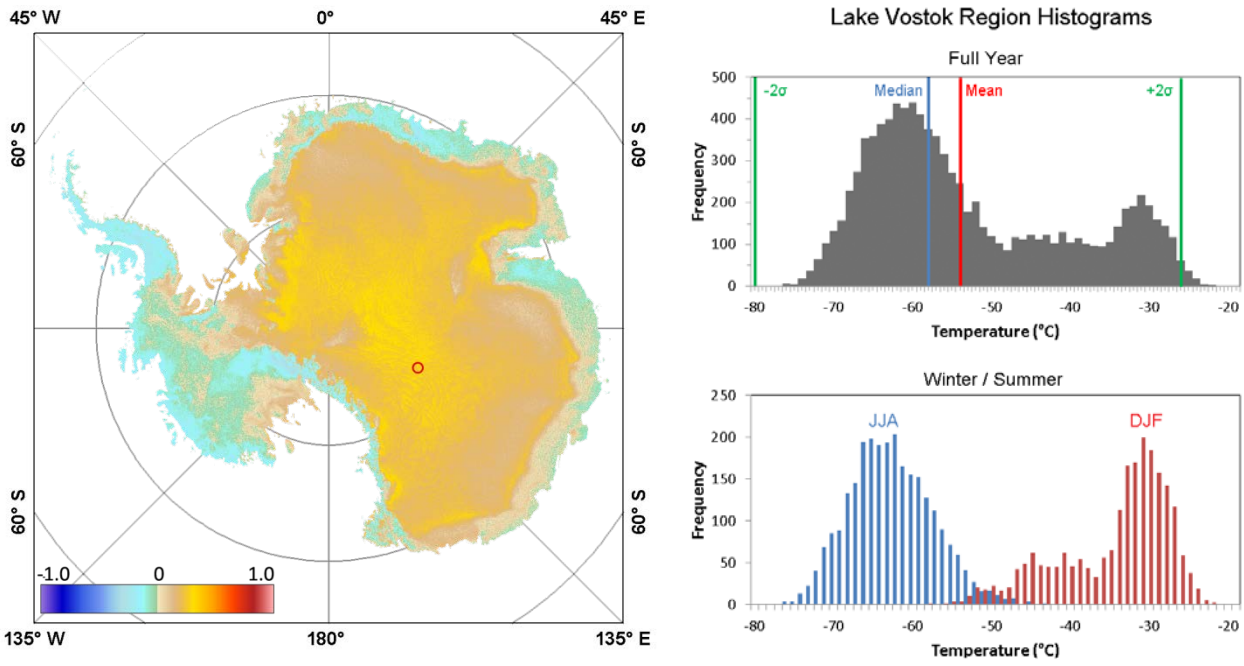

Figure 4.18. Non-parametric skewness of the temperature distribution at each grid cell (left), using all observations from 2000-2017. The temperature distributions in the high elevation regions are positively skewed. A closer examination of one area, southeast of Lake Vostok (red circle), reveals the cause of the skewness. Histograms (right panels) of the interannual yearly (top) and seasonal (summer and winter, bottom) show the LST distributions for the Vostok location. The full-year histogram is positively skewed with a bimodal distribution, the result of combining multiple seasonal distributions. In particular, the summer (DJF) and winter (JJA) temperature distributions (lower histogram) drive the bimodal yearly distribution. When using two standard deviations as criteria for determining extreme values, the high variance and skewness of the full-year distributions often excludes the lowest temperatures. Similar cropping may occur with the seasonal distributions.

4.4 Teleconnections to Climate Indices

Section 4.3 examined the trends in LST extremes, including the annual number of cloud-free samples and the frequency at which those samples met either the 5th or 95th percentile threshold. Although no statistically significant trends were found in the magnitude of the temperature extremes or the sample populations, there was considerable variance from year to year. This section looks deeper into possible associations between the variance in the sample

populations of Antarctic LST extremes and climate cycles. Many studies have looked at possible teleconnections between Antarctic weather conditions and climate modes, such as the Southern Annular Oscillation, Southern Oscillation Index, or sea ice extent. The goal here is not to attribute causation or necessarily conclude there are teleconnections, but to demonstrate the ease of replicating similar analyses when using condensate databases.

The climate of Antarctica, or regions within it, may be influenced by a variety of climatic cycles with periods of days to years, as well as aperiodic patterns. Cause and effect are not always clear, and there may be a time lag between an influential climate mode and changes to Antarctic weather conditions. A small time lag demands a higher temporal resolution, so here a series of condensate databases is used that contain extreme LSTs binned by month. The higher temporal resolution enables statistical tests for time lags that might otherwise be obscured when using longer time periods.

In the following subsections, the monthly extreme LST sample populations are correlated with three possible climate indices: the Southern Annular Mode (SAM) index, Southern Oscillation Index (SOI), and Sea Ice Extent. Any statistically significant connections would necessitate further investigation before concluding attribution.

The primary tool used here for examining correlations is the Cross Correlation Function (CCF). A CCF asks the question: What is the relationship between two time series variables, X_t and Y_t? The formal definition computes the set of correlations between X_{t+h} and Y_t, for time lags of h = 0, ±1, ±2, etc., to some maximum time offset. For example, in Section 4.4.1 the Southern Annular Mode (SAM) index is used as the CCF X time series and the sample population of LST anomalies is Y. A common output from a CCF is a plot of possible correlations against time lags

(e.g., Figure 4.19). Interpreting a CCF plot is straight forward, where the strength of the correlation is represented by the vertical lines. Statistically significant correlations will exceed one of the two critical values (blue horizontal dashed lines). Positive correlations imply that a positive change in the SAM index is associated with a positive change in the anomaly population; negative correlations indicate an inverse relationship. A correlation to the left side of time *t*=0 indicates a leading relationship, i.e., a change in the SAM index *leads* a change the number of anomalies by *h* months. Conversely, correlations to right side of *t*=0 indicate a lagging relationship, where changes in the SAM index *lag* a change in the anomaly population by *h* months. Each critical value is at the 5% level, a function of the sample population and a hypothesis of zero correlation. The critical values are approximately $\pm 2/\sqrt{n}$, where the sample population (*n*) is the number of time increments. For the LST dataset *n*=204 months, from 01 March 2000 through 28 February 2017. Throughout this study, *h* was chosen to be ±15 months, ensuring that any plausible relationship between the two variables would be visible. Notably, with a total of 31 possible lags (*h* = ±15 lags, plus *t* = 0) and a 5% critical value (α = 0.10 for a two-tailed test), there may be approximately 3 statistically significant correlations in each CCF simply through random chance. Thus, interpretation of small or isolated, but statistically significant, correlations should be circumspect.

4.4.1 Southern Annual Mode Index

The Southern Annular Mode (SAM) is a hemispheric pattern of climate variability, driven by atmospheric dynamics in the middle latitudes. The SAM, and its northern hemispheric equivalent, the Northern Annular Mode (NAM), describe the variability of *anomalous* atmospheric flow. These variations in flow are considered anomalous because they are not

seasonal patterns. The SAM index is based on the zonal pressure difference between 40° south and 65° south latitude [Gong and Wang, 1999]. Formally, the index is defined as:

$$SAM = P^*_{40°S} - P^*_{65°S} \quad (R)$$

where $P^*_{40°S}$ and $P^*_{65°S}$ are the monthly normalized zonal mean sea level pressures at 40° south and 65° south latitude. This pressure difference can be visualized as changes in atmospheric

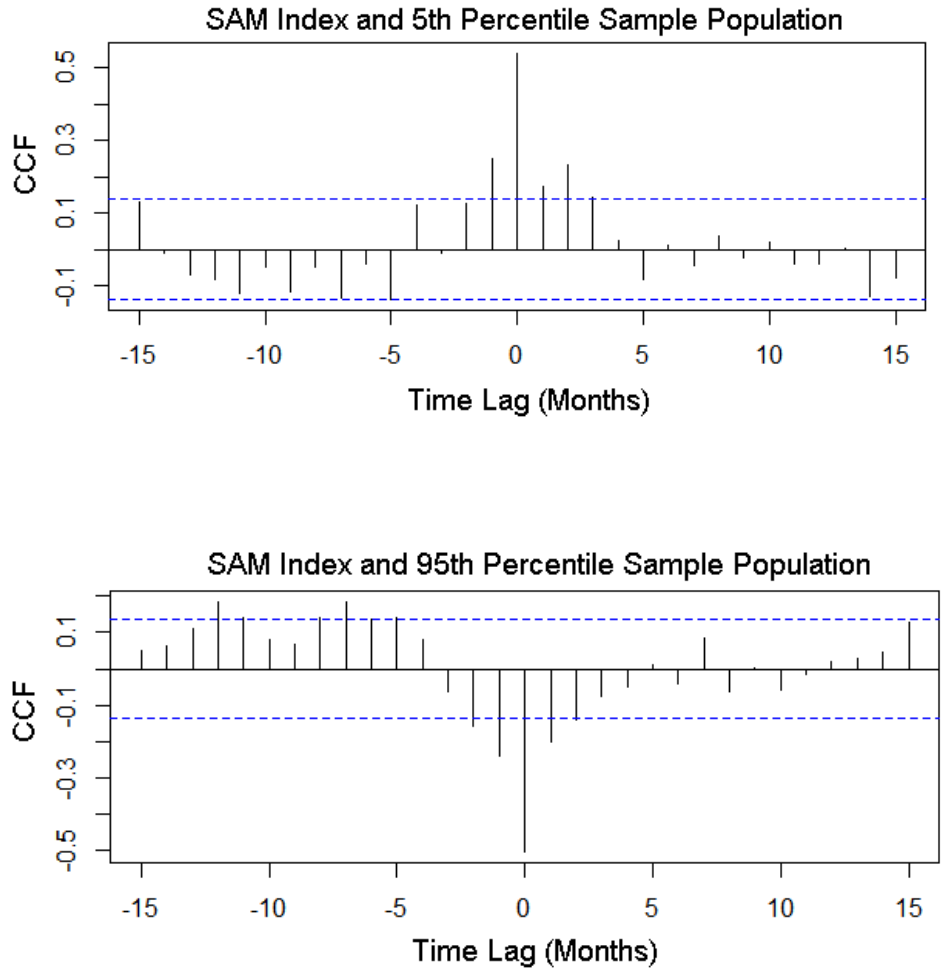

Figure 4.19. Cross Correlation Function plots for the SAM index and LST anomaly sample populations. Shown are correlations with the 5th percentile LST anomaly populations (top) and 95th percentile (bottom), with critical value thresholds (blue dashed lines). Correlations to the left of time t=0 indicate that changes to the SAM index lead changes in the sample population.

mass as air moves from the mid-latitudes to more polar regions. Such changes occur on time scales of weeks to years. A monthly SAM index provides a practical temporal resolution, smoothing out much of the noise without losing sight of the bigger atmospheric changes. Positive SAM indices are associated with above-average west winds at higher latitudes (50° south to 70° south) and weaker west winds in the mid-latitudes (30° south to 50° south). The mechanisms whereby the SAM affects Antarctic temperatures are still up for debate and often regionally defined. For instance, Marshall et al. [2006] suggest that the stronger westerly winds associated with a positive SAM index are pushing over the orography of the Antarctic Peninsula, creating, among other effects, warming föhn winds on the lee side of the mountain range. Farther south, the correlation is less clear. In Marshall [2007], which examines correlations between interior surface temperatures and SAM polarities, causality is not thoroughly explained; the study suggests using General Circulation Models for further investigation.

This dissertation found a strong, statistically significant positive correlation between the SAM index and the number of LST anomalies around t=0 for the 5th percentile of LST anomalies, and a strong negative correlation with 95th percentile anomaly populations (Figure 4.19). Rephrased, a positive SAM index is associated with more cold LST extremes and fewer warm LST extremes. In general, the relationships neither lead nor lag. Minor correlations with 5, 7, and 12 month lags are also apparent, but these are only slightly significant (or not at all) and perhaps a seasonal effect in the data.

Previous studies have found relationships between the SAM index and near-surface air temperatures, but the direction and magnitude of the correlation varies by location and season. Marshall [2007] found significant correlations, generally positive in the Peninsula region and

negative in East Antarctica, but varying seasonally. Kwok and Comiso [2002] also found relationships, notably "positive polarities of the SAM are associated with cold anomalies over most of Antarctica." This latter finding matches the results shown in Figure 4.19, where the population of anomalous cold temperatures (5^{th} percentile) is positively correlated with a positive SAM index and warm anomalies are negatively correlated with a positive SAM index. Importantly, and supporting this dissertation's purpose of demonstrating the efficacy of the condensed database system, these teleconnections were identified quickly using a simple set of SQL queries and a short R script to perform the CCF.

4.4.2 Southern Oscillation Index

The Southern Oscillation Index (SOI) measures the difference in sea level air pressure between Tahiti and Darwin, and is one index for gauging the strength of El Niño and La Niña events. Although the index is based on the atmospheric pressure in Australia and the South Pacific, these conditions may have an impact on global weather. A positive SOI above +8 may indicate a La Niña event while values below –8 indicate an El Niño. The SOI correlates well with warm ocean temperatures in the eastern tropical Pacific. The ocean temperatures, in turn, affect atmospheric circulation, and create a potential climatic teleconnection with Antarctic temperatures. Since the pressure difference can vary considerably due to short term weather conditions, the SOI is calculated on monthly scales or longer. This dissertation uses a monthly SOI for comparisons with the monthly population of anomalous Antarctic temperatures.

The SOI has no clear teleconnection with the 5^{th} percentile (coldest) anomaly population (Figure 4.20, upper plot). There appears to be a slight negative leading correlation, and a small positive lagging correlation, but the relationships are not distinctive or statistically significant.

However, the SOI does have a distinctly positive, statistically significant leading correlation with the 95th percentile (warmest) anomaly population (Figure 4.20, lower plot). The implication is that, 5 to 6 months after a positive SOI event, more warm temperature anomalies may be expected in Antarctica.

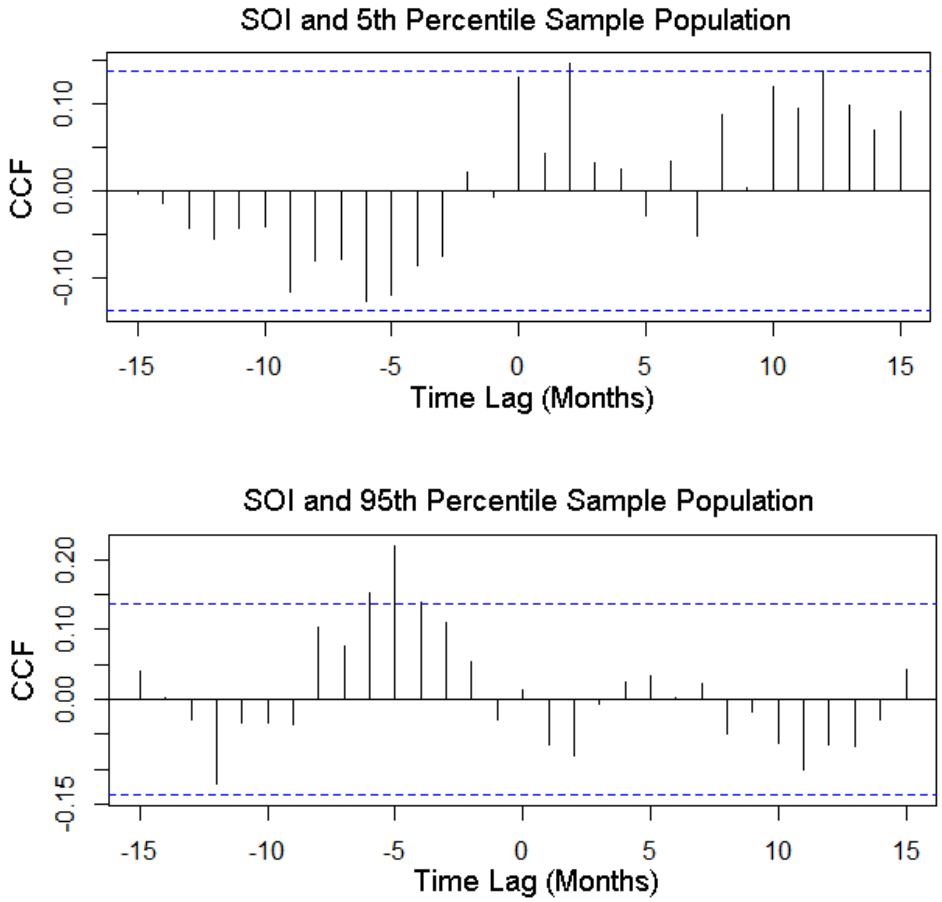

Figure 4.20. Cross Correlation Function plots for the SOI and LST anomaly sample populations. Shown are correlations with the 5th percentile LST anomaly populations (top) and 95th percentile (bottom), with critical value thresholds (blue dashed lines).

The causes of this connection are not necessarily clear. While zonal wind patterns appear to have a relationship with Antarctic temperatures, studies of the teleconnections between the El Niño Southern Oscillation (ENSO) index and Antarctic weather patterns are generally inconsistent or not robust over long time spans [Turner, 2004]. This suggests that other factors, unidentified ocean and atmospheric processes, may be playing important roles. With the exception of the 5 to 6 month leading correlation between the SOI and LST warm anomalies, the teleconnections between the SOI and LST extremes are less definitive than those of the SAM index. Smith and Stearns [1993] also looked at anomalous Antarctic near-surface air temperatures (using observations from 10 surface stations, mostly coastal) surrounding SOI minimums and found conflicting correlations as well, both negative and positive, with a variety of lag and lead times. The findings in this dissertation suggest that broader spatiotemporal coverage, using satellites instead of surface observations, may reveal a positive leading correlation between the SOI and high surface temperature anomalies.

4.4.3 Sea Ice Extent

Sea ice extent defines the ocean as either ice-covered or not ice-covered based on a threshold for coverage. The extent data used in this dissertation uses a 15 percent threshold, where the ice concentration within a grid cell must be 15 percent or greater. Sea ice extent is, of course, a highly seasonal index with possible trend components. Since unusual extent events may be lost in the natural (or forced) seasonal variations, this dissertation instead uses the *residuals* between the detrended, modeled sea ice extent and the observed monthly average extent (Figure 4.21).

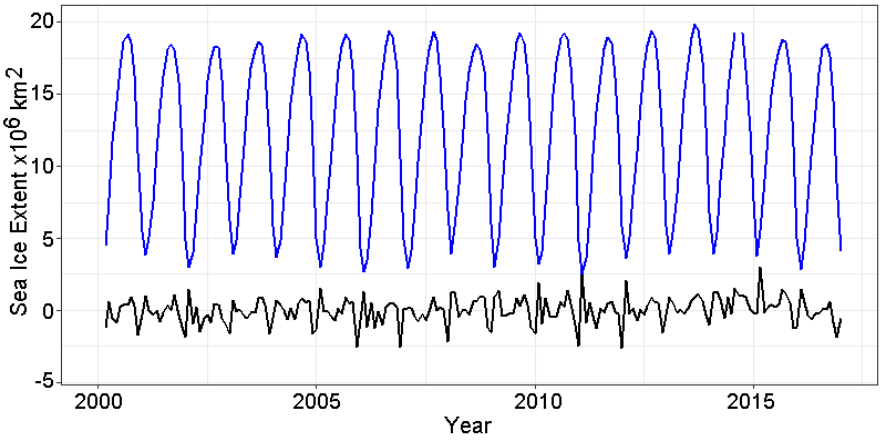

Figure 4.21. The observed sea ice extent (blue) and residuals from the modeled extent (black).

Correlations between the sea ice extent and anomaly sample populations rise to the statistically significant level only once, with an 11 month lag, in the 95th percentile of LST anomalies. The correlation is isolated and appears to be due to random chance (Figure 4.22). A similar study using sea ice *area* residuals rather than extent had equally uncorrelated results (not shown).

Past work has suggested a negative correlation between air temperatures and sea ice area with a one to two season lag, but the effect is highly regional [Weatherly et al., 1991]. The current study does not subdivide by region and is thus missing regional relationships. Although not demonstrated here, SQL queries against the condensate databases can be easily modified to select local sectors, latitude/longitude bounding boxes, individual MODIS tiles, or individual grid cells, potentially duplicating or expanding on Weatherly et al.'s work.

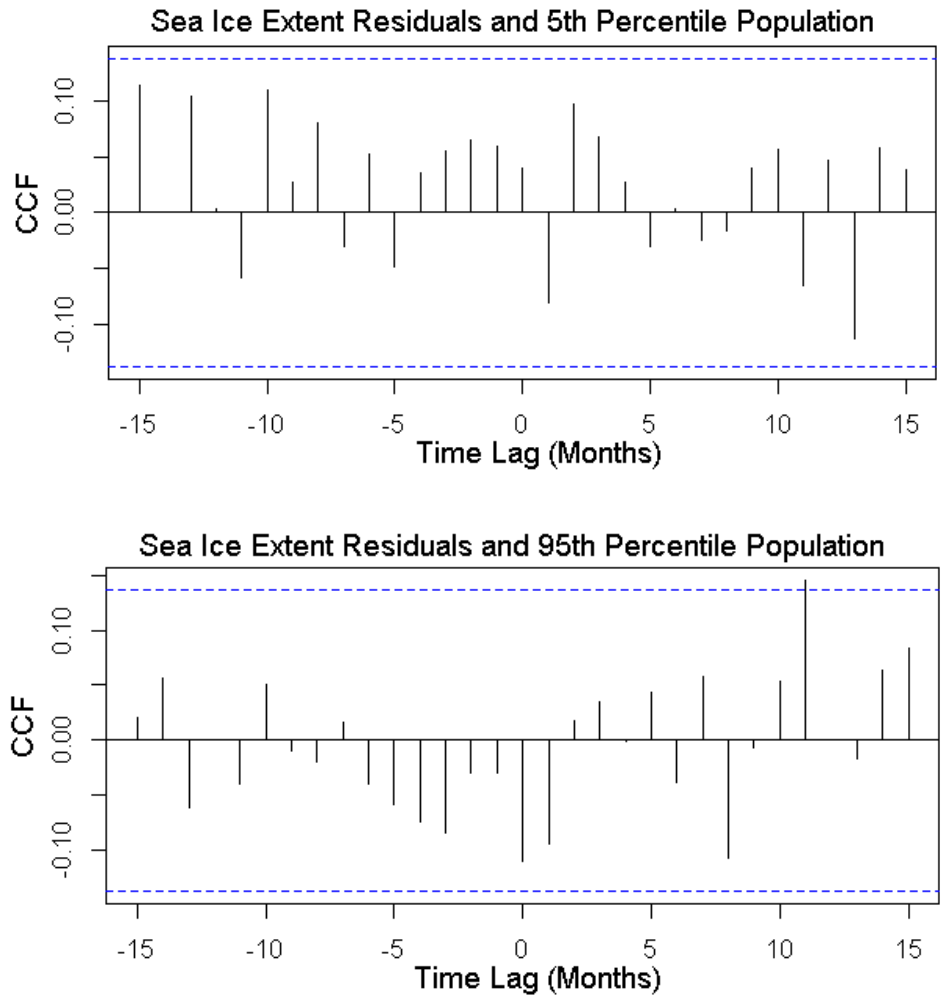

Figure 4.22. Cross Correlation Function plots for Antarctic sea ice extent residuals and LST anomaly sample populations. Shown are correlations with the 5th percentile LST anomaly populations (top) and 95th percentile (bottom), with critical value thresholds (blue dashed lines).

4.5 Case Studies

Previous sections explored the baseline LST descriptive statistics, the spatial and temporal patterns of Antarctic LSTs extremes, and looked for possible teleconnections between LST extremes and climate indices. The following case studies dig deeper into a few of the most intriguing discoveries, and provide demonstrations of end-to-end scientific analyses, from

hypotheses to conclusions. Since the anomaly databases were intended as tools for data exploration and QC, not necessarily for comprehensive scientific analyses, these case studies are augmented by the full MODIS LST CDR and ancillary datasets.

4.5.1 South Pole Temperature Trend

Figure 4.8a displayed a possible cooling trend in the South Pole region. The following case study asks the question: 1) Is the observed trend real, and not an artifact of the crude linear regression? And 2) if the LST trend is statistically significant, do ground air temperature observations at the South Pole show a similar trend?

To answer the first question, this analysis starts by referring back to the full, uncondensed LST data set. Figure 4.23a shows the cleaned time series of South Pole LST observations. To produce this time series, two grid cells adjacent to the South Pole were aggregated and data gaps interpolated, generating a monthly averaged LST. Since the LST time series does not show, nor do we expect to see, exponential growth over time, and the variance is approximately constant, the analysis may use a classical decomposition of the time series (Figure 4.23b-d). This decomposition includes the time series trend and enables a test of the hypothesis: *Is the observed cooling trend statistically significant?*

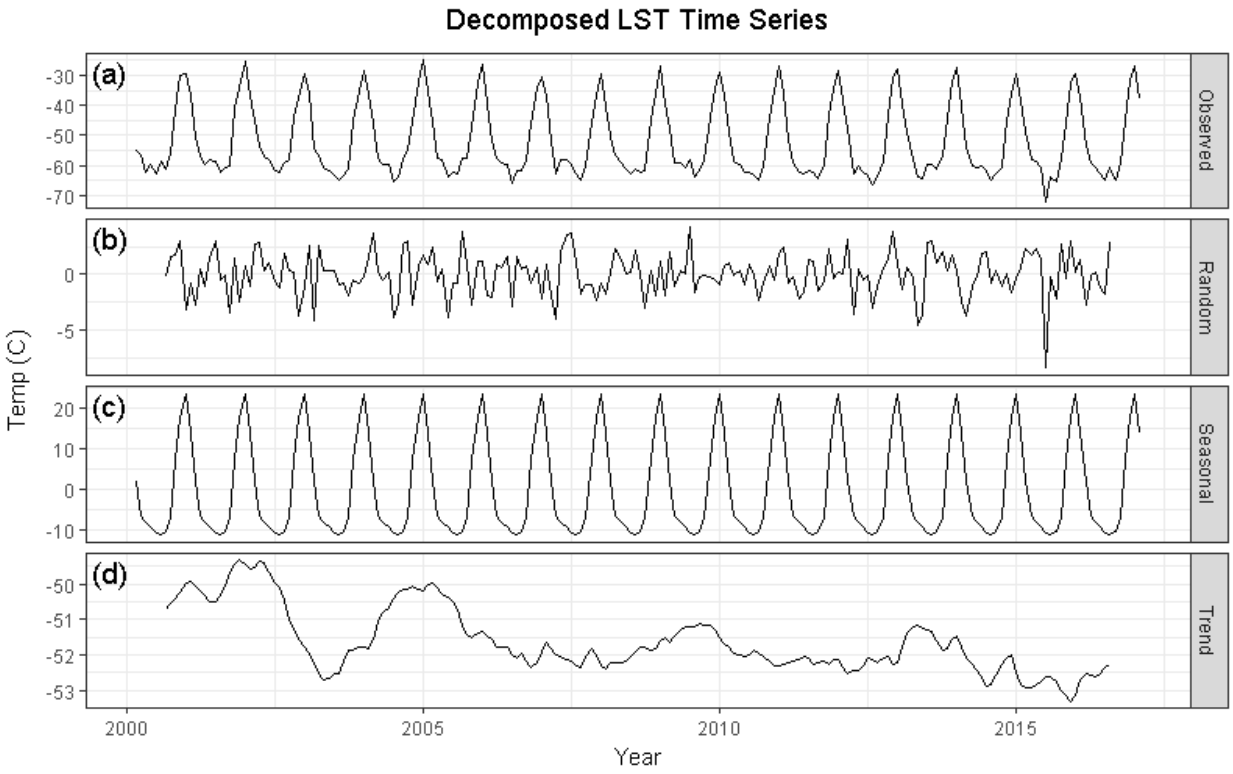

Figure 4.23. A classical additive decomposition of the time series of LST observations near to the South Pole. Panel (a) shows the actual observations, (b) the random component, (c) modeled seasonal cycle, and (d) the trend component.

Before testing the trend for statistical significance, the observation time series must be normally distributed, or reasonably so; here, normality is confirmed using the de-trended residuals from the seasonally modeled observations using a quantile-quantile plot (Figure 4.24a). The residuals are generally normally distributed, although a slight "fat tail" of cold temperatures is present, most likely due to occasional cloud contamination. A single outlier is present at the most extreme quantile, corresponding to unusually cold weather during July 2015. A cross-check against infrared MODIS imagery confirms clear skies for the LST observations during that July, and the South Pole's manned surface observations approximately match the LSTs. Therefore, the extreme LST is "real", but prominent because many days were

cloud contaminated (possibly warmer) and the monthly sample size was small. The outlier is clearly visible late in the random error plot (Figure 4.23b). The otherwise-normal distribution of residuals allows us to perform a linear regression on the trend component (significantly non-normal residuals would imply a skewed or heteroskedastic data distribution, requiring a transform to normality before performing a linear regression).

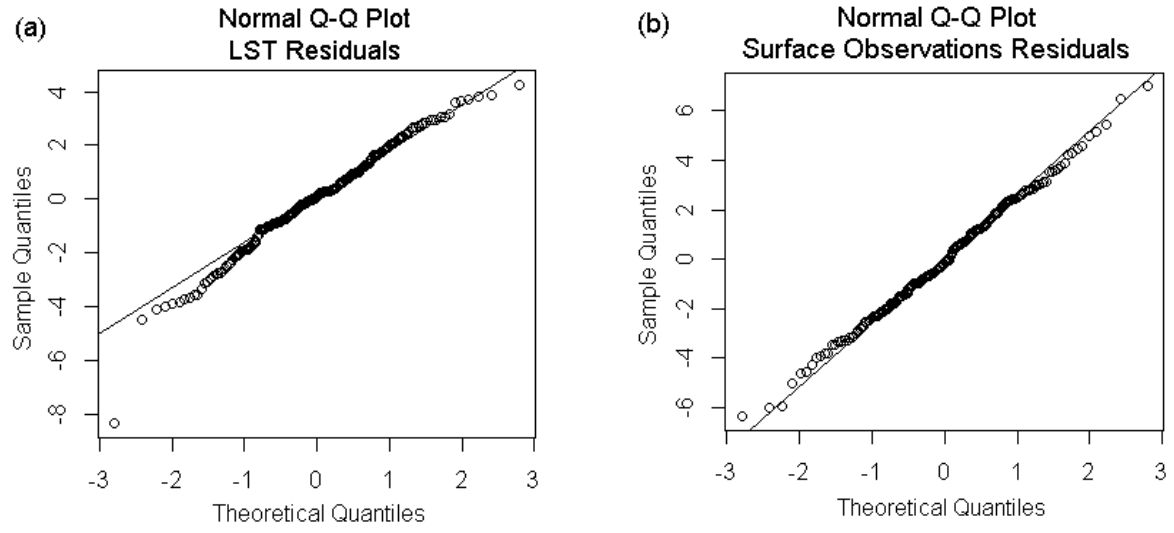

Figure 4.24. Quantile-quantile plots of the residuals from the detrended and modeled South Pole LST observations (a) and the manned surface observations (b).

Retaining the outlier in the data, the decadal trend is approximately −0.2 °C/decade but not statistically significant ($p = 0.92$). With the outlier removed, the trend is substantially less, −0.05 °C/decade and statistical significance decreases ($p = 0.98$). Therefore, we do not have sufficient evidence to conclude the trend is “real”.

For completeness, a similar decomposition of South Pole’s manned surface observations is also given here (Figure 4.24b and Figure 4.25). The results show a positive slope for T_{air}, 1.7 °C/decade, but it too is not statistically significant ($p = 0.35$). Interestingly, this positive trend

roughly matches the coarse linear regressions for much of East Antarctica outside of the area directly around the South Pole (see Figure 4.8).

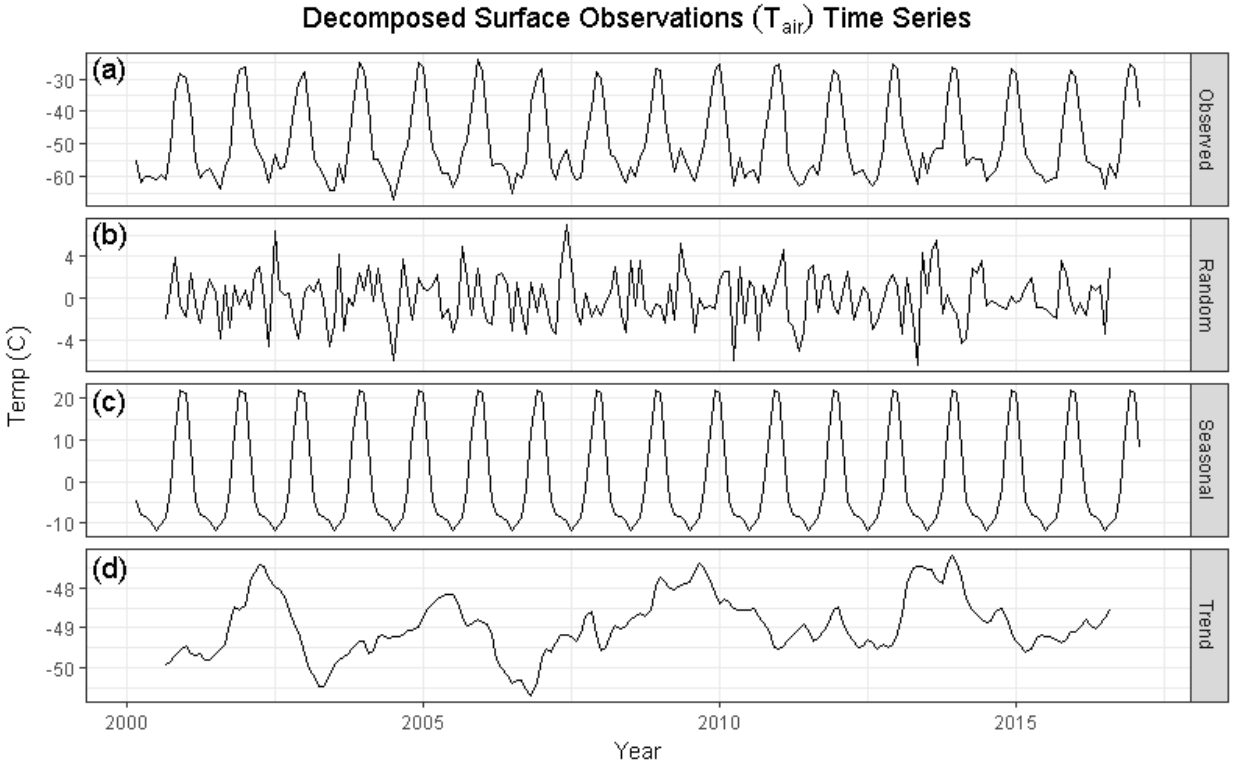

Figure 4.25. A classical additive decomposition of the time series of manned air temperature observations at South Pole Station. Panel (a) shows the actual observations, (b) the random component, (c) modeled seasonal cycle, and (d) the trend component.

In summary, the hypothesis of an LST trend toward colder temperatures around the South Pole was not confirmed. On the other hand, surface observations suggest the temperatures are actually warming, but the variability is so high that it does not show statistical significance. Left for later investigation is another possibility, that a seasonal breakdown of LST trends may reveal statistically significant results.

4.5.2 Extreme Cold event

The rapid data exploration capabilities provided by the condensed anomaly databases provide a method of demonstrating the hunt for Antarctic temperature extremes. Queries

against the anomaly databases reveal that the most extreme cold temperatures occur along the East Antarctic Ridge, a high elevation feature dominating East Antarctica. Here, land surface temperatures are consistently colder than any other area (Figure 4.26). Locations near and down-slope from Dome Fuji and Dome Argus show the most extreme cold temperatures. This case study looks at one such extreme, where on 23 July 2004, after 10 days of decreasing temperatures, the observed land surface temperature dropped to −97.2 °C (−143 °F). The grid cell location is 81.85° south 59.37° east, downhill from the 4,093 m (13,428ft) Dome Argus. This site is not unique; several nearby locations show LSTs nearly as cold. Aggregating the 3x3 grid cell matrix (9 square kilometers) surrounding the location produces an average LST of −95.4 °C on the same date.

The relatively coarse spatial scales of both the LST observations (1 km^2) and reanalysis data (0.5° x 0.625° for MERRA-2 and 0.75° x 0.75° for ERA Interim) tend to smooth over topographical features that enhance cold air accumulation. The ice sheet, although visually flat, has depressions where the coldest air accumulates. These depressions are the locations where record cold temperatures occur. The orographic effect is visible in Figure 4.27, where the coldest monthly land surface temperatures occur in topographic low spots downhill from the ridge. Once the air pools in these locations, clear skies permit continued heat loss through longwave radiation. Without significant air advection and mixing, strong inversions can form, enhancing the conditions for the formation of extreme lows. Lacking incoming shortwave radiation and with only minimal downwelling longwave radiation, and in the absence of turbulent heat flux or latent energy exchange through evaporation or condensation, the only other source of energy is the ground heat flux (see Equation (H). The initial rapid decrease in

temperature on day 14 (Figure 4.28) is likely due to the top layer of snow/ice coming into thermal equilibrium with the air temperature. Subsequent surface cooling relies on the continuation of favorable weather conditions, especially the absence of clouds or warm advected air masses.

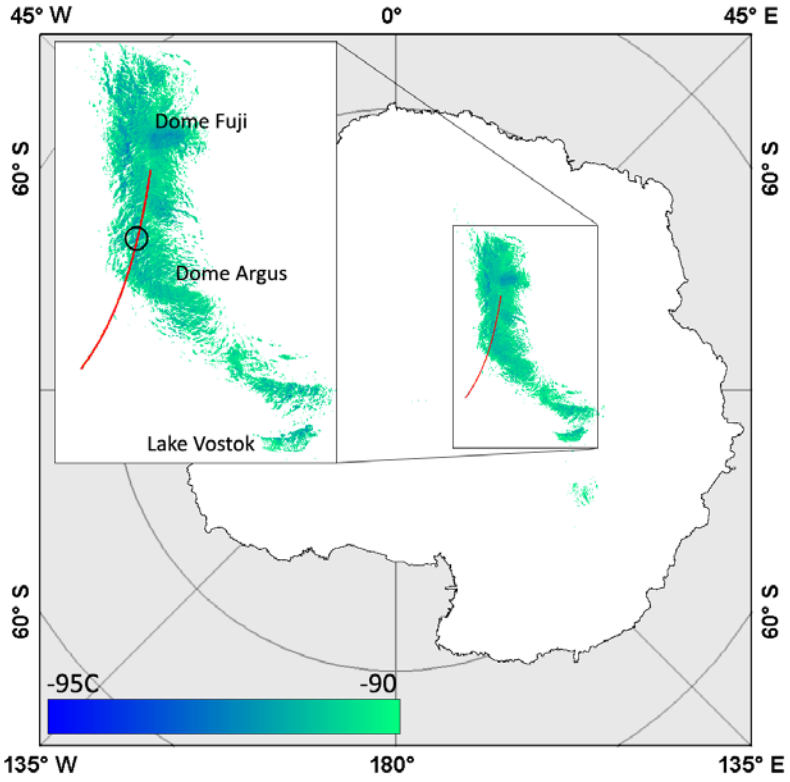

Figure 4.26. Extreme cold temperatures along the East Antarctic Ridge on 23 July 2004, and the location of the extreme cold temperature event (inset, black circle). A profile of the elevations and temperatures (red line) is graphed in Figure 4.27.

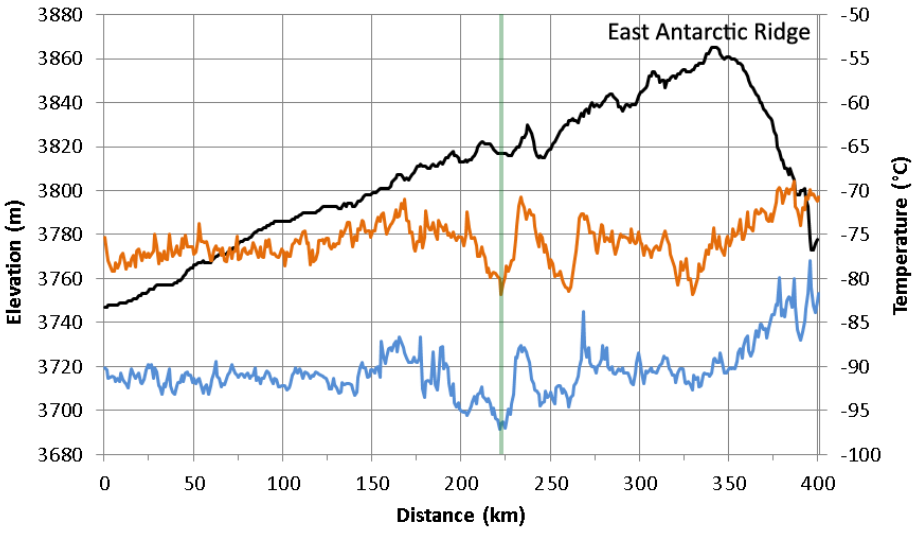

Figure 4.27. A southeast-to-northwest transect (path shown in Figure 4.26) of elevation (black line), with the average monthly temperature (orange line) and minimum temperature (blue line), for July 2004, and the location of the coldest temperature on July 23 (vertical green line).

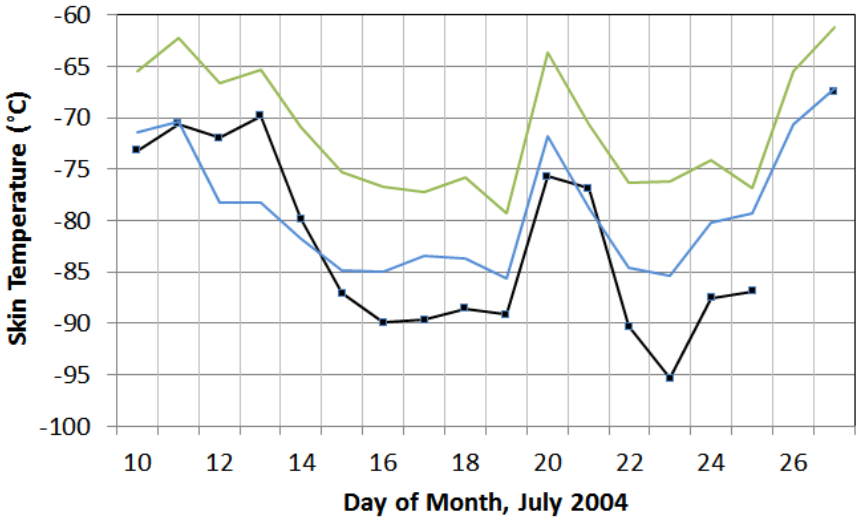

Figure 4.28. The time series of temperatures leading to the extreme cold event. The graphs are an average of the 3x3 group of grid cells at the location of the observed coldest LST (black), temperatures from the MERRA-2 reanalysis dataset (green), and the ERA Interim reanalysis data (blue).

Manual inspection of infrared imagery during the 2004 event confirms clear skies on most days. The anomaly LSTs are high quality and do not show evidence of cloud contamination. A time series plot of the local LSTs shows rapidly decreasing temperatures starting on July 14th and tapering off, then continuing, with a short interruption on the 20th and 21st, until the coldest LST occurred on July 23rd (Figure 4.28). The brief rise in temperatures coincided with light cloudiness that moved through, obscuring the coldest center grid cell for two days. Comparisons with surface skin temperatures from NASA's MERRA-2 and ERA Interim reanalysis datasets show similar patterns, although both reanalysis datasets are biased warmer, likely due to their coarser spatial resolution (Figure 4.29).

The weather conditions during this period show light winds, coinciding with a cold high pressure system settling over the area (Figure 4.30 through Figure 4.32). Between July 15th and 23rd, central Antarctica is dominated by a poorly organized mid- to high-pressure system, with the highest pressure settling in over East Antarctica on July 23rd. (Figure 4.30). A brief intrusion of lower pressure, coincident with higher winds and thin clouds, raises temperatures on the 20th and 21st, but then the higher pressure reasserts itself bringing a sharp drop in temperatures and wind speeds. On the 23rd, the winds remain mild and the temperature drops to its low point (Figure 4.31 and Figure 4.32).[3] The decreasing temperatures during an extended period of light winds, high pressure, and clear skies are consistent with previous ultra-cold temperature events [e.g., Turner et al., 2009]. In each case, land surface (skin) temperatures drop due to lengthy long-wave radiation cooling combined with light air advection. Near-

[3] While the reanalysis datasets generally show the same patterns of temperature, pressure, and winds, this study found persistent differences between the two. For example, ERA Interim consistently reports wind speeds approximately 2 m/s faster than MERRA-2, and in light winds the directions may be 180° apart. Similarly, MERRA-2 consistently reports skin temperatures 5 °C to 10 °C warmer than ERA Interim.

surface air, flowing down from higher elevations and pooling in topographic valleys, contributes to the land surface temperature cooling as well.

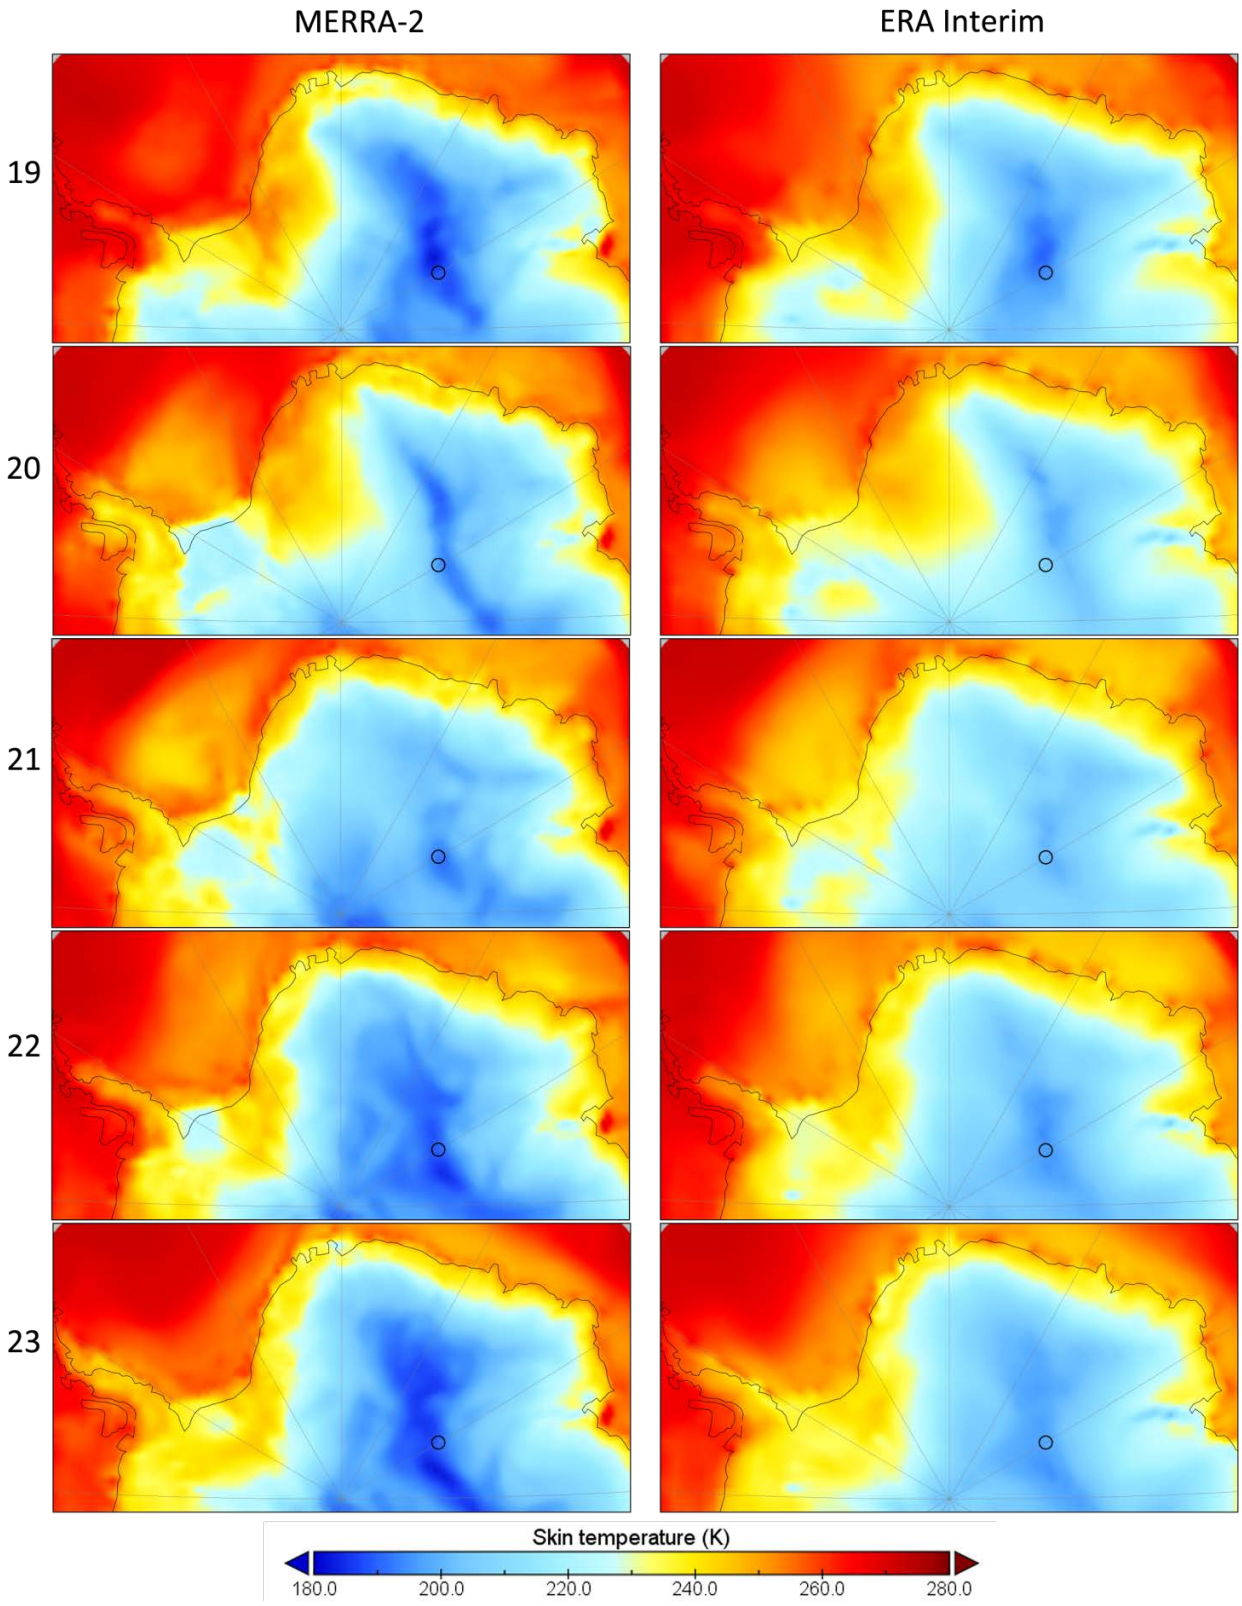

Figure 4.29. Antarctic skin temperatures for July 19-23, 2004, from MERRA-2 (left column) and ERA Interim (right column). The location of the coldest point is shown in the black circles.

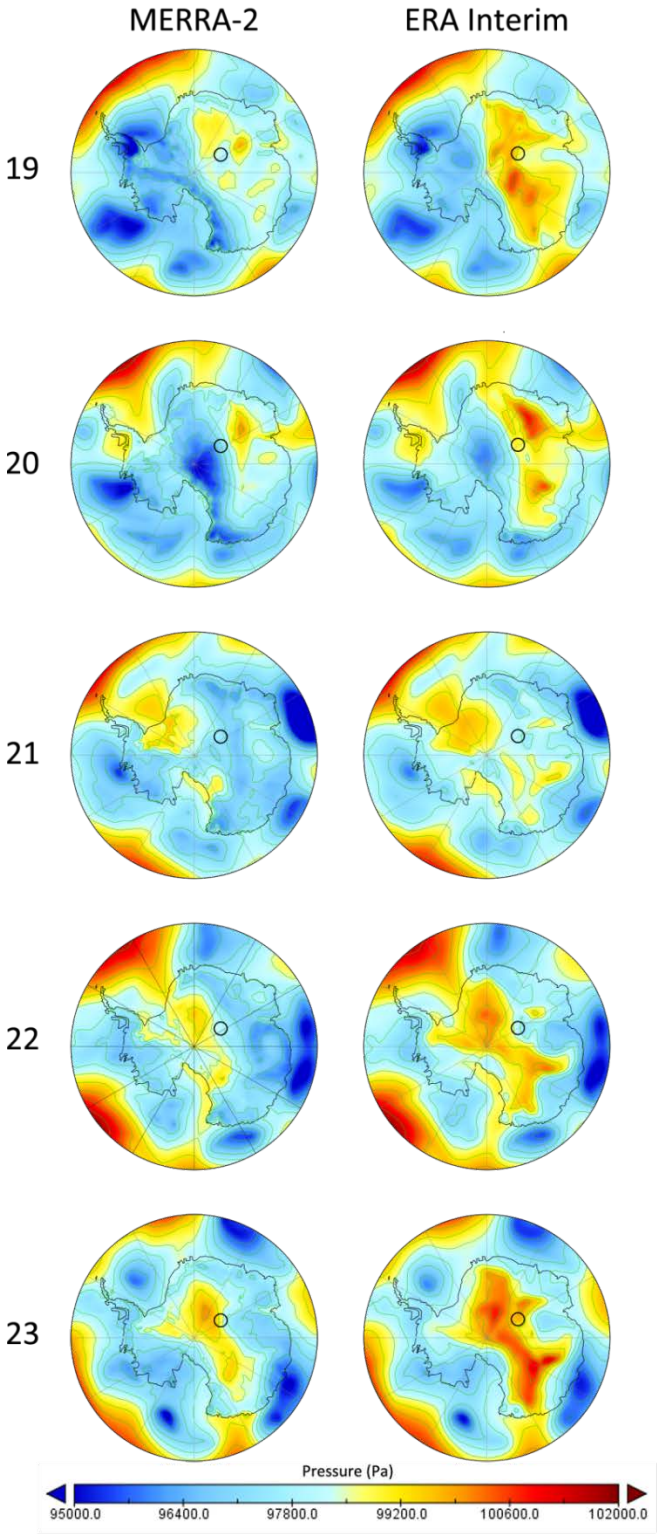

Figure 4.30. Antarctic sea-level-adjusted pressure for July 19-23, 2004, from MERRA-2 (left column) and ERA Interim (right column). The location of the coldest point is shown in the black circles.

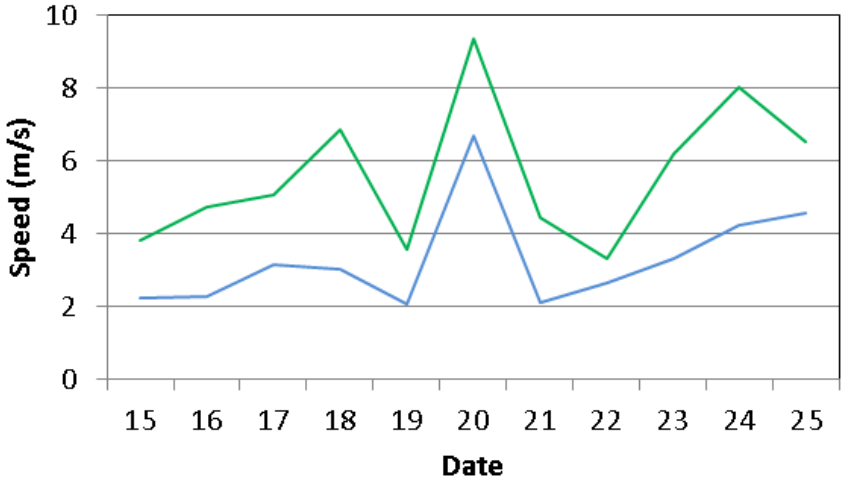

Figure 4.31. Surface wind at the extreme cold event location from 15 to 25 July 2004 from MERRA-2 (blue line) and ERA Interim (green line).

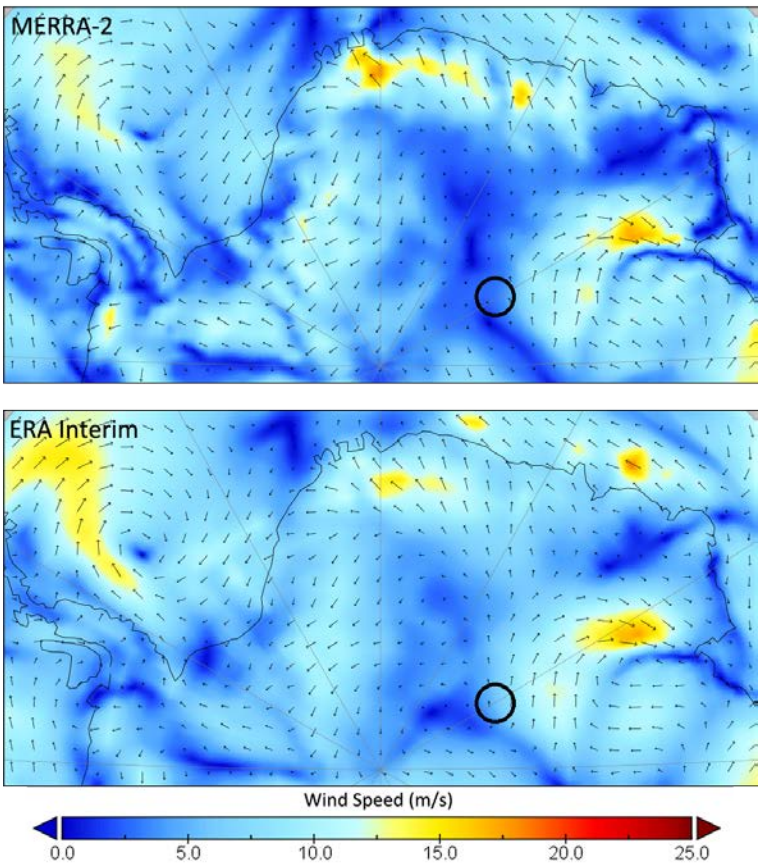

Figure 4.32. Antarctic 2 m wind vectors on 23 July 2004 from MERRA-2 (top) and ERA Interim (bottom). The location of the cold temperature event is within the black circles.

4.5.3 Extreme Warm event

Unlike extreme low surface temperatures, the conditions causing extreme high temperatures are more complex. A key factor is often solar exposure. Returning to Equation (I), the downwelling solar radiation (S↓) is no longer zero, and the net shortwave radiation (S↓ + S↑) will be positive. Cloud cover and advected air masses affect the net longwave radiation, but if R_{net} becomes positive this can easily lead to a warming surface.

In Equation (H), when R_{net} is positive the sign of the ground heat flux may also go positive as the subsurface absorbs energy. The turbulent heat fluxes, H_S and H_{LE}, gain influence due to advected warm air and increased latent heat exchanges. Near the coasts, the H_S term becomes especially potent due to warm air masses infiltrating inland from the ocean. For example, the 24 March 2015 record high near-surface air temperature at Base Esperanza is attributed to the influx of a warm maritime air mass and foehn winds [Skansi et al., 2017; WMO, 2017]. This record came in the early Antarctic autumn, after the equinox, limiting the influence of solar radiation; indeed, visible satellite imagery of the region confirms overcast skies during that time. Thus, under certain conditions H_s may overwhelm R_{net}. For satellite-based LST observations, this confluence of conditions is problematic because the cloud cover precludes measurements. This dissertation acknowledges the limitations in detecting extreme highs, in particular those accompanied by cloud cover as is common near the coasts.

At higher elevations and away from coastal cloudiness, it is more likely that the summertime energy balance is dominated by R_{net}. In particular, during clear, sunny conditions, the downwelling solar radiation may be especially strong, and warm air advection may further enhance extreme temperatures. These very conditions are the attributed causes of the current

high-elevation record, −7.0 °C, at Automated Weather Station D-80, set on 28 December 1989 [Skansi et al., 2017]. Similar events should be detectable using the condensed LST databases.

A rapid search of the condensed databases revealed several candidate events (see Appendix B for example queries). Importantly, the temperatures of the inland Antarctic Plateau tend to be extremely cold, even at the height of summer; some of the candidate warm extremes were cloud contaminated, and the cloud top LSTs were actually warmer than any cloud-free surface observations. While the presence of warm cloud contamination complicated the search for extreme high events, it also suggests that influxes of warm, moist air may be setting record temperatures in the continental interior as well, not only the coasts. For this search, then, visible imagery was manually consulted to verify clear conditions. Candidate warm LST events were discarded if the imagery was unavailable or indicated clouds.

The remaining anomalously warm events often shared a unique characteristic: In many cases, the warmest extremes occurred not on the ice sheet but on north-aspect mountain faces, most with low-albedo bare rock surfaces. The case study described in the following paragraphs details one such event.

Exposed rock slopes near Mount Cavaney (2,820 m) in the Deep Freeze Range, just inland from the Ross Sea, reached a land surface temperature of −5.2 °C on 04 January 2014. Visible imagery confirms that the warmest point was cloud-free at the time of observation. The extreme temperatures occurred during an extended period of clear weather, and were followed by a large breakup of nearby fast ice (Figure 4.33). Melt pools or bare ice are also visible on the nearby Nansen Ice Sheet (not shown).

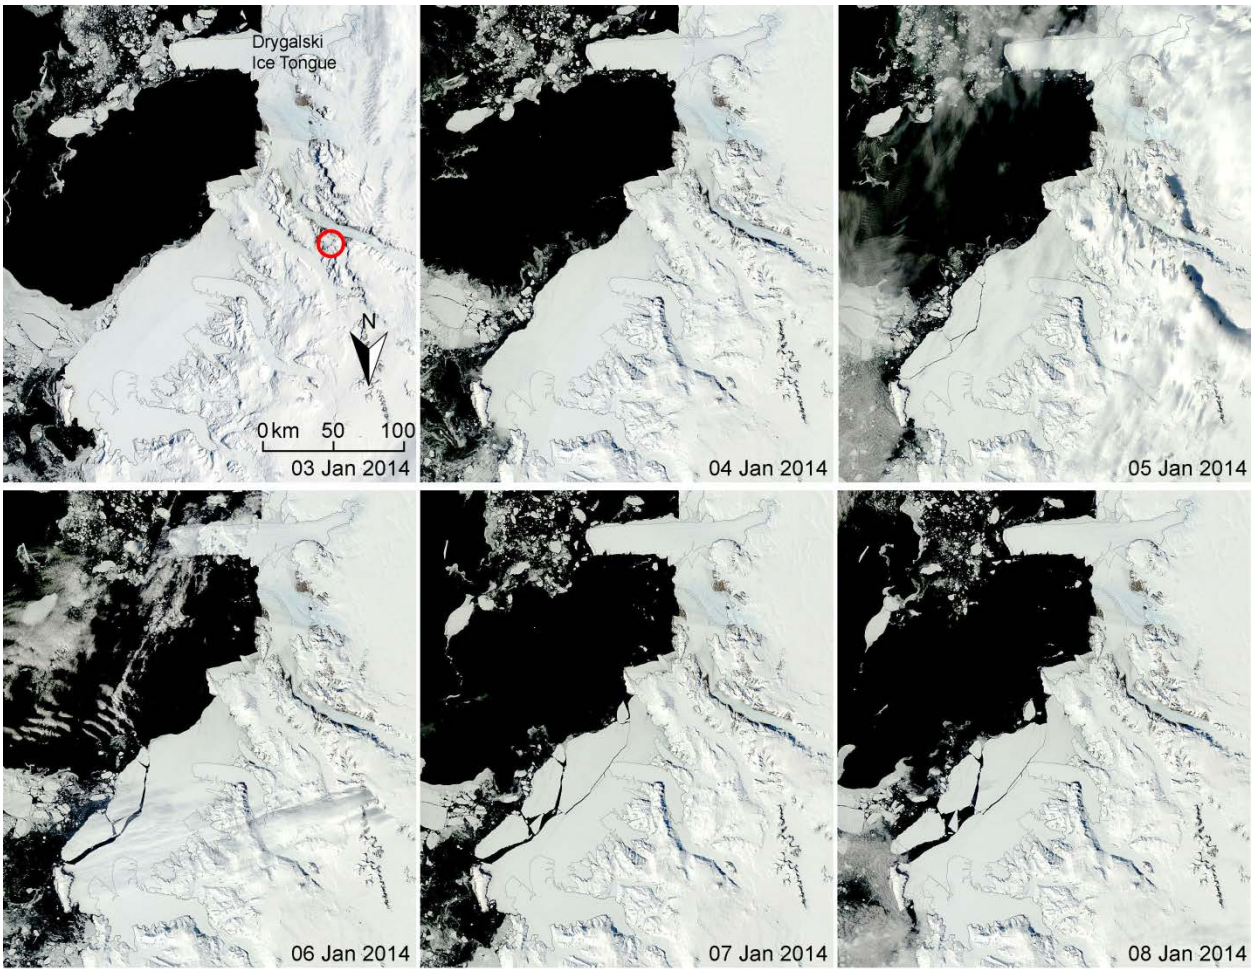

Figure 4.33. Visible images of the Ross Sea and Deep Freeze Mountain Range, including the Drygalski Ice Tongue. The location of the warm extreme on 04 January 2014 is in the red circle (first image, upper left panel). A section of fast ice approximately 150 km long broke free in subsequent days. [MODIS/Terra visible imagery courtesy NASA]

LST observations indicate a sharp rise in temperatures beginning on January 1st, culminating in the warm extreme on the 3rd and 4th (Figure 4.34). Reanalysis data is too spatially coarse to discriminate such small (approximately 1 to 5 km^2) areas, but suggests a daily temperature profile similar to the LST observations. Regional wind speeds ranged to 10 m/s or more on the days preceding the event, but diminished on the 3rd to ~2 m/s (Figure 4.35). Sea level corrected atmospheric pressure during the first week in January averaged around 990 mb,

with no indication of strong storm fronts moving through the area. Thus, light winds prevailed through the extreme high LST event.

The site location, −73.98° south 163.01° east, has an elevation of 2843 m and is located among other mountain faces that also appear anomalously warm on a frequent basis. Similar to the exposed, north-facing slopes along the Peninsula (see Section 4.2.1) which show occasional near-record temperatures for low elevations, high elevation slopes may also exceed the existing >2500 m temperature records in Antarctica. If the near-surface air temperatures are close to the observed LSTs, this region, when instrumented, could potentially set new Antarctic warm temperature records for elevations above 2500 m.

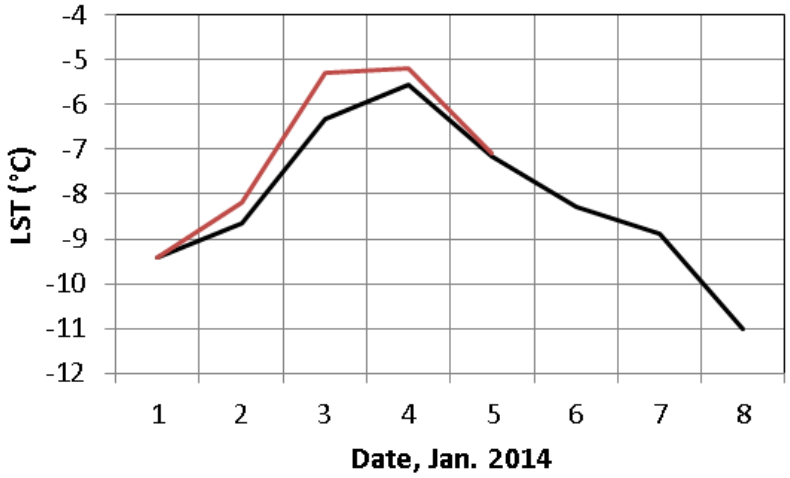

Figure 4.34. Daily LST observations at the site of the warm temperature event (red line) and the average of the adjoining 3x3 grid cells (black line). Cloud cover obscured observations prior to January 1st and partially after the 5th.

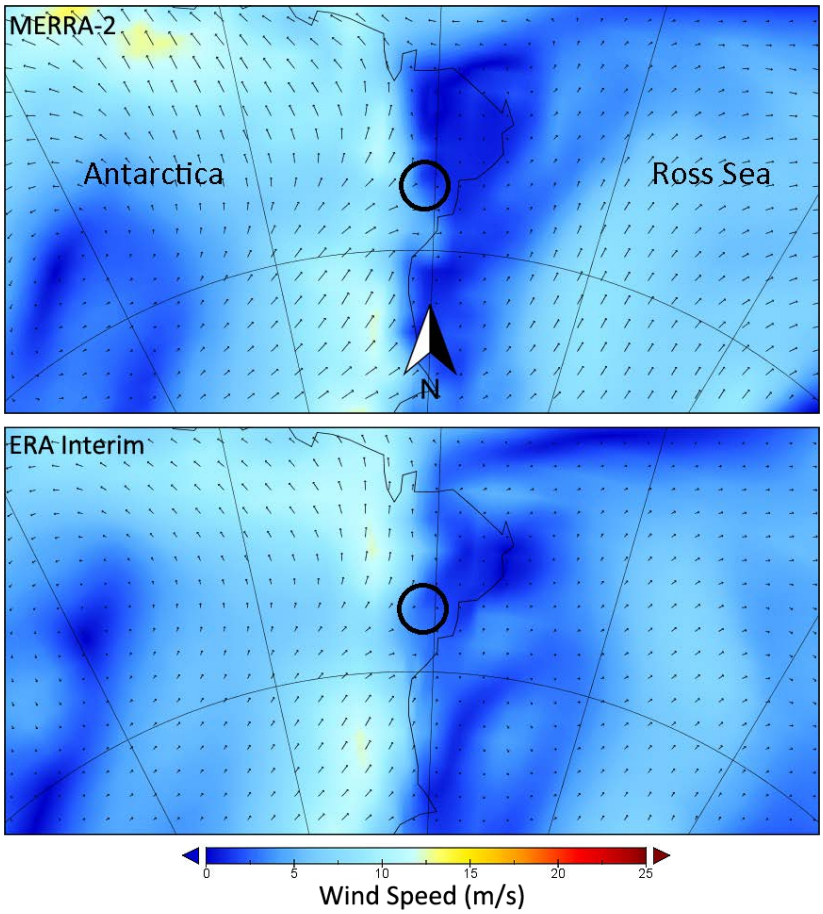

Figure 4.35. 2-meter wind vectors along the Ross Sea Coastline for 3 January 2014 from MERR-2 reanalysis data (top) and ERA Interim (bottom), and the location of the anomalous warm event (black circles).

4.5.4 Cloud-Free Observations Discontinuity

The total population of LST observations and clean, cloud-free samples are byproducts of the condensate database generation process. The ratio of cloud-free samples to total observations provides a rough measure of cloudiness, and is a unique tool for data exploration and quality assurance. In Figure 4.36, an unexplained rise in cloud-free observations is clearly visible starting in 2011. There is also an apparent change in the seasonal pattern of cloud-free data, with the most conspicuous changes to the cycle occurring during the austral summers.

The pattern of cloud-free pixels is regular, if seasonally dependent; a sudden, natural jump into a permanently different – but also regular – cloud regime would seem unlikely. This suggests there was a hardware change, processing change, or perhaps sudden sensor degradation in one of the MODIS channels used for cloud detection.

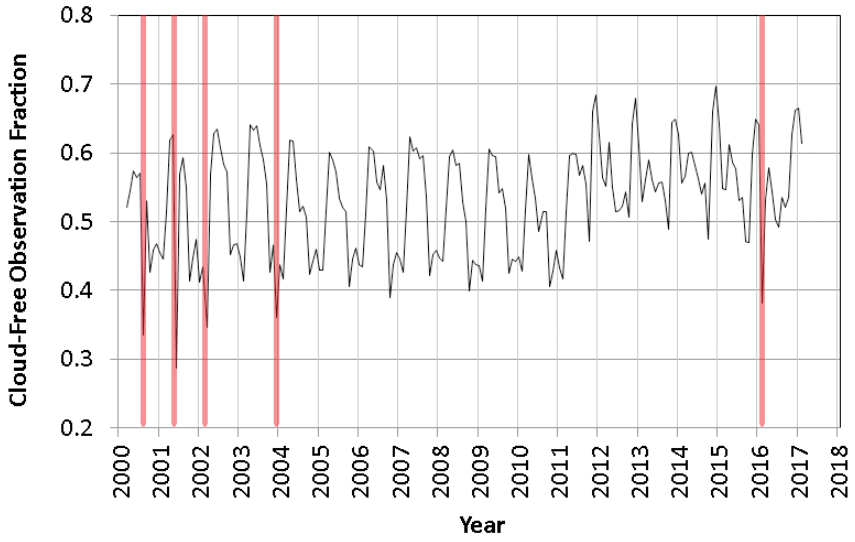

Figure 4.36. The fraction of cloud-free observations, by month, relative to the total population of observations (with or without clouds). Red lines show periods of substantial missing data due to satellite problems.

NASA's timeline of MODIS/Terra cloud processing shows no specific events during 2011 that may have caused a sudden change in observations (https://modis-atmosphere.gsfc.nasa.gov/data-issues/cloud). The MODIS C6 Product Generation Executive (PGE) Version History Summary shows no processing changes taking place that would account for the changes (https://modaps.modaps.eosdis.nasa.gov/services/production/pge_all.html). Likewise, there is no evidence of sensor health issues that would account for the change (https://mcst.gsfc.nasa.gov/sites/mcst.gsfc/files/file_attachments/Terra_C6_DQF_Table_07261

6_0.xlsx). Thus, the jump in the cloud-free fraction is either a natural phenomenon or, more likely, the result of a heretofore undocumented change in sensor performance or processing.

The changes seen in Figure 4.36 are a product of the LST cloud masking, and this provides a starting point for tracking down the cause. LST clear-sky confidence levels are reported as quality flag SDSs in each daily file. The quality flags are determined using the C6 cloud masking product, MOD35_L2 [Wan, 2013]. In the MOD35 product, cloud masking over polar snow and ice is performed using a series of algorithms that vary by day and night[4] [Strabala, 2004]. Specifically, the cloud masking algorithms use combinations of MODIS bands centered at wavelengths of 0.87, 0.935, and 1.38 μm during the daytime, 3.7 and 12 μm during nighttime, and 6.7, 11, 12, and 13.9 μm during both the daytime and nighttime. Since the cloud masking changes are apparent throughout the year, the most suspect channels are the latter four. These wavelengths correspond to MODIS bands 27, 31, 32, and 35. Of these, bands 31 and 32 are used in the split-window algorithm that generates the LSTs. If a sensor problem has occurred in one of these bands, it may be visible in the LST data.

A cursory check of the LST time series quality is possible using the condensate database statistics. The percentage of monthly LSTs that were found to be anomalous is shown in Figure 4.37. Visually, the percentages show no discernable discontinuities but are highly noisy. Linear regressions of the 5^{th} and 95^{th} percentiles' cloud-free sample populations have no statistically significant trend ($p = 0.167$ and $p = 0.781$, respectively, with $\alpha = 0.05$). However, the noisiness of data may be obscuring any trends. Similarly, plots of regional LST values over time have too

[4] It is unclear how the MODIS cloud masking decision tree differentiates day from night in polar regions with 24 hours of daylight or nighttime.

much natural variability to reveal minor temperature discontinuities or trends (not shown). Thus, this initial check on bands 31 and 32 is inconclusive.

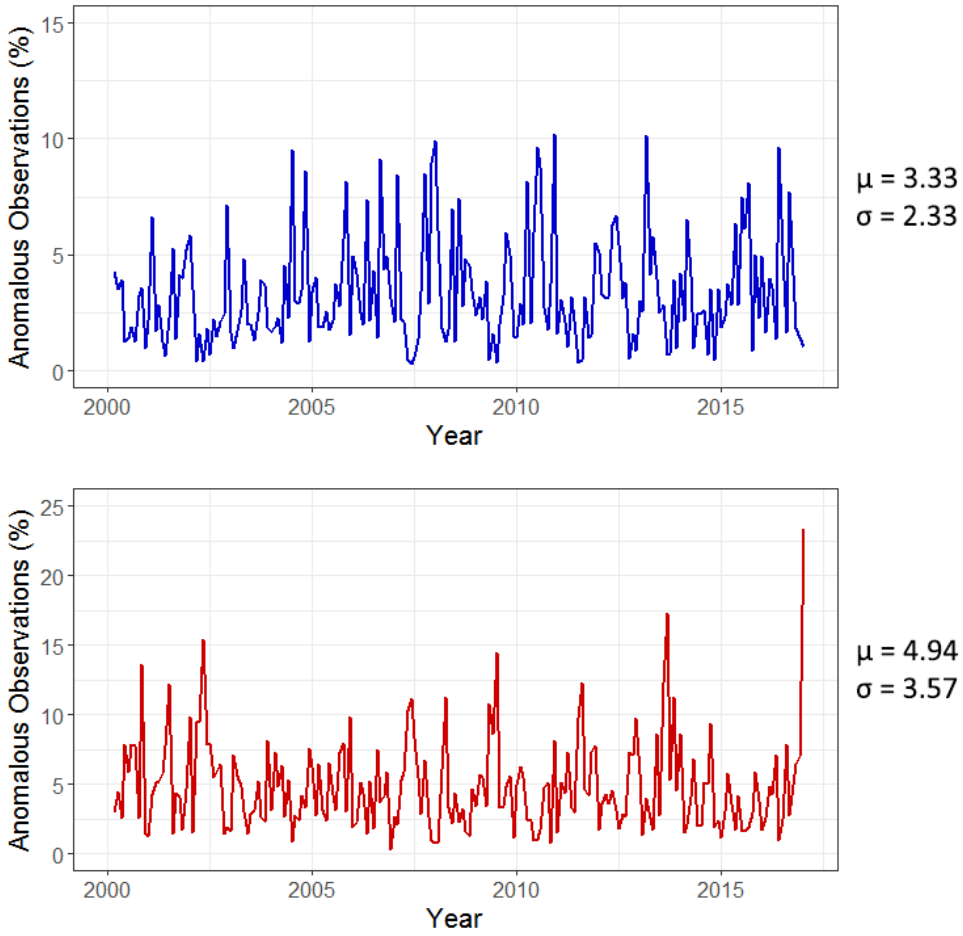

Figure 4.37. The percentage of cloud-free LST observations each month that were determined to be anomalous, showing the 5th percentile (top graph) and 95th percentile (bottom graph). Also shown (right) are the mean monthly percentages (μ) and standard deviations of the monthly percentage (σ), for a sample size of N=204 months.

Another approach to determine the cause of the 2011 jump in clear-sky observations is a comparison between MODIS/Terra and MODIS/Aqua data. If they both show the same effect, then the cause is either natural or a product of NASA's cloud masking processes. As shown in Figure 4.38, both Terra and Aqua exhibit similar discontinuities in 2011, thus Terra sensor degradation can be ruled out.

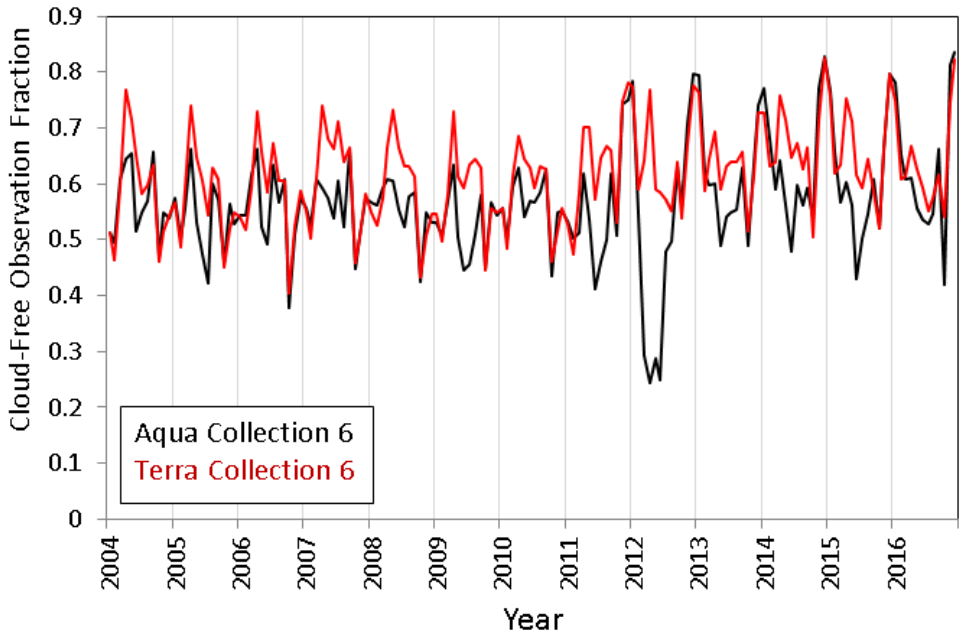

Figure 4.38. The fraction of Terra cloud-free LST observations versus Aqua, both using Collection 6. The data shown is for a single tile, h=18 v=16 (East Antarctica), for the years 2004 through 2016. Both Terra and Aqua show similar changes in cloud detection starting in 2011.

The next logical step is to determine if the 2011 changes in cloud detections are a natural phenomenon or a data processing problem. Figure 4.39 compares Terra Collection 6 with Terra Collection 5. As shown, the 2011 increase did not occur in Collection 5. A similar comparison using MODIS/Aqua (not shown) is the same. For comparison, a study using a tile in the Amazon Rainforest showed very few differences between Collections 5 and 6, and no sign of the 2011 discontinuity (not shown). Another test using tiles in the Sahara Desert showed a slight discontinuity in 2011, suggesting that the problem is related to cloud detection over bright or snow-covered surfaces. These results point to the NASA's cloud masking algorithm in the Product Generation Executive as the most likely source of the cloud detection discontinuity.

The exact cause of the problem or location within the code is out of the scope of this dissertation, however the results found here will be forwarded to NASA.

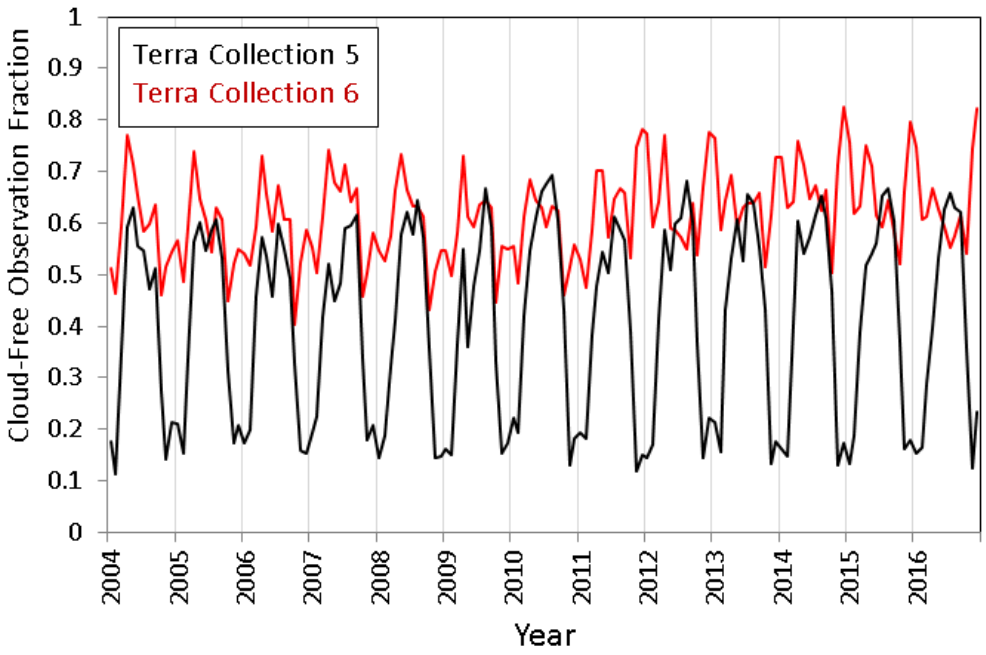

Figure 4.39. The fraction of Terra Collection 6 cloud-free LST observations versus Terra Collection 5. The data shown is for a single tile, h=18 v=16 (East Antarctica), for the years 2004 through 2016. Collection 6 shows an increase in clear-sky detections starting in 2011 whereas Collection 5 does not. The difference in variability between the two collections also raises questions, which remain unresolved.

5 Conclusion

This dissertation investigated the land surface temperature extremes of Antarctica using a new set of data processing methods. The methods were developed as a part of the *Condensate Database Project*, which sought to condense massive satellite datasets into easily-handled databases of anomalies. The goal of the project was to expedite data discovery and exploration, bypassing the typical, cumbersome process of downloading masses of data and then writing customized analysis code. In contrast to typical analysis techniques, the condensate databases provide a reusable method for repeated, varied data exploration. This dissertation acts as a proof-of-concept for the condensate methods. The techniques were remarkably successful, and especially effective at quality checking datasets.

To demonstrate the effectiveness and efficiency of the condensate databases, approximately 434 GB of MODIS/Terra land surface temperature (LST) data, 17 years of twice-daily observations, was condensed and entered into databases. This data spanned the entirety of the Antarctic continent at a resolution of 1 km. Condensation efficiency was approximately 3 to 4%, meaning that 96 to 97% of the source data was discarded by the condensation algorithm, leaving only anomalous data of high interest. The algorithm eliminated unnecessary data, such as off-the-map points and cloudy pixels, created statistical baselines of the original dataset, and then used the baselines to further reduce the dataset to observations considered anomalous. Anomalous LSTs were defined as those below or above the 5^{th} and 95^{th} percentiles. An additional convolution filter reduced noise by ensuring that candidate anomalies were adjacent to at least one other anomaly. The final output of the condensation process was 1) statistical baselines of the source LST dataset, and 2) a set of compact databases containing

anomalous LSTs. The baselines provided a “first look” at the distribution of LSTs around Antarctica plus a quick method of gauging data quality. The databases of anomalies were accessed for repeated and varying queries in support of data exploration and quality checking.

Queries against the anomaly databases were typically fast, returning images and results within a few minutes, often within seconds. Hardware capabilities, of course, were a limitation on the performance of the database queries, and too-broad queries (e.g., requesting all data for all times, across the entire continent) resulted in stalled searches. As such, the condensed databases are best used for specific questions, constrained by geographic boundaries, selected timespans, or a narrow range of sensor values. Within those constraints, the architecture and methods performed well.

Because the databases contained only anomalous data, the types of scientific questions that could be posed were limited. Nevertheless, as demonstrated by the examples in this dissertation, the possibilities are highly varied. Once an interesting event or data quality concern is found, the user may be obligated to return to the original, full dataset for further investigation. Thus, while the anomaly databases cannot answer all questions, they do provide a useful starting point for data exploration. Instead of an investigator downloading an entire dataset and building customized tools for analysis, the *Condensate Database Project* provides a set of tools to rapidly identify areas of special interest and return the most intriguing data. The system is especially adept at identifying quality problems.

5.1 Findings

5.1.1 Data Quality

Significant, systematic errors were exposed during ingestion of the MODIS/Terra LST dataset. Some MODIS tiles contain LST values that do not correspond to any valid terrestrial location; within the sinusoidal projection of the source MODIS tiles, the data points appear off the Earth (Section 2.1.2). Similar misplaced points were not found in MODIS Collection 5, implicating NASA's Collection 6 processing. Secondly, a ring of anomalously warm temperatures was found at approximately 87 degrees South latitude. The cause of the spurious temperatures is undetermined, but the characteristics strongly suggest a geometric error somewhere in MODIS LST processing.

Of greater concern, the cloud masking for Collection 6 appears to be inconsistent. Since cloud contamination may affect the final results of any study using LSTs, this issue has serious implications for the overall quality and usefulness of the dataset. This dissertation performed a deeper study of the extent and causes of the cloud masking irregularities, discussed in Section 4.5.4 and summarized below in Section 5.1.6.1.

5.1.2 Baseline Descriptive Statistics

The condensate algorithm's baseline statistical images revealed additional data quality issues. Although not necessarily errors, *per se*, LST dataset users should be aware of the potential for false conclusions due to the data's idiosyncrasies. The discovery of these problems emphasizes the value of the condensate methodologies.

First, NASA's LST processing uses a variable cloud masking confidence level based on elevation (Section 2.1.1). Below 2000 meters the processing relies on a 95% confidence of clear skies, while above 2000 meters the confidence level drops to 66%. The result is a discontinuity in LSTs at 2000 m, where higher elevations are distinctly colder due to cloud contamination (visible in Figure 4.4) but have a greater sample population (Figure 4.1). The lower threshold for clear skies also affects the absolute minima and maxima, where cloud contamination is clearly visible in the baseline images (e.g., Figure 4.10c).

Secondly, the baseline statistics also show that sample populations diminish in a circular pattern around the South Pole (e.g., Figure 4.2a). As discussed in Section 4.1.1., the reduced sample size is likely the result of oblique viewing angles. With an orbital inclination of 98.2°, the Terra satellite never flies directly over the pole, thus all observations of the South Pole are off-nadir. This geometry leads to increased atmospheric optical depth and greater cloud detection uncertainty. The lower sample size and oblique viewing angles may also be causing an additional data quality issue: Linear regressions of LSTs near the South Pole show a similar circular pattern, and suggest a cooling trend that runs counter to the warming shown elsewhere on the continent (Figure 4.8). However, the time series data at any given grid cell is noisy, and the correlation coefficients for the linear regressions are low. A single unusually cold or warm season, occurring naturally or through cloud contamination, strongly affects the outcome of the regressions. Effectively, it may appear as though a trend is occurring when in fact the statistical significance is far too low (Section 4.5.1).

Byproducts of the condensation process include the population of all observations at any grid cell, the number of cloud-free samples within that population, and the number of

those samples that met the criteria to be considered anomalous. While exploring the time series of yearly clear observations, an unexpected change in the ratio of cloud-free samples was discovered (Figure 4.16). A detailed look at the ratio revealed a startling increase beginning sometime in mid-2011, accompanied by a change in the seasonal pattern (Figure 4.36). The cause of the increase is still being investigated, although it seems likely that sensor degradation or perhaps processing changes are responsible. Because the MODIS data is supposedly a consistent, long-term Earth observation dataset, any unexplained changes occurring during the time series are cause for concern.

The means and standard deviations of LSTs (Figure 4.5) show that East Antarctica has both the lowest mean temperature on the continent, and, in general, the greatest overall variance in temperatures. Regions of the Transantarctic Mountains, however, exhibit a much greater LST variance during the winter season than any other location (Figure 4.7). Absolute maximum and minimum LSTs (Figure 4.10 and Figure 4.11) show wide-spread cloud contamination despite thorough dataset cleaning. Noting that the most extreme LSTs are likely to be cloud-contaminated, this result cautions against assuming the validity of isolated anomalies without supporting information, such as reanalysis data or visible imagery.

5.1.3 Spatial Patterns of Extremes

Queries for high and low LST anomalies were rapid, confirming known areas of extreme air temperatures and suggesting new locations for study.

The availability (or lack) of LST anomaly samples provided clues to the energy exchange causing extreme land surface temperatures. For instance, warm air advection plays a significant role in creating high temperature extremes. Near the coasts, the warmest events are often

accompanied by considerable cloud cover, hindering satellite observations of LSTs (see Section 4.2.1). The current record holder for Antarctic low-elevation high temperature is Base Esperanza, along the Antarctic Peninsula. When the Esperanza record was set, no LST observations were available due to cloud cover, implying that longwave energy exchange and sensible heat transfer played a greater role than shortwave insolation (Equation (I). While advection also plays a role in high elevation (>2500 m) warm extremes, insolation may have a larger effect. Thus, more high LST extremes may be observable in the plateau regions. In all locations, however, terrain aspect angle and ground coverage is especially important; the queries revealed that bare, north-facing slopes of exposed rock show frequent high LST extremes (Figure 4.14). These areas are not instrumented, leaving open the question of whether record extreme high temperatures are occurring there quite often but simply remain unrecorded.

The lowest extreme temperatures occur along the East Antarctic Ridge (see Figure 4.13). This information prompted a case study to find the area of absolute coldest temperatures and provided causality (Section 4.5.2). The results are summarized below, in Section 5.1.6.

5.1.4 Temporal Patterns of Extremes

A series of condensate database queries provided a look at the temporal trends in LSTs from 2001 to 2016. The queries explored the yearly means, medians, and standard deviations of the continent-wide LST extremes, as well as the cloud-free sample populations. On both a yearly and seasonal basis, and without subsetting to specific regions, the variance of LSTs is large; no statistically significant trends in the means or medians were found (Section 4.3 and Table 1). For the sample populations, however, a jump of approximately 7% in clear-sky

samples appeared in 2011 (Figure 4.16). The change in cloud fraction was further explored as a case study (Section 4.5.4). While the cause of the change in cloud detection remains undetermined, the pattern change is even more distinct when viewed on a monthly basis (Figure 4.36).

The initial development of the condensation algorithm relied on LST means and standard deviations for determining anomalies. Queries against the resulting anomaly databases, however, revealed that the annual and seasonal distributions of LSTs exhibit strong skewness. The skewness of the distributions caused the mean/standard-deviation algorithmic technique to bypass entirely many extreme LST values (see Section 4.3.1). Based on this evaluation, the Condensate Database Project changed the anomaly determination algorithm from using means and standard deviations to percentiles.

5.1.5 Teleconnections

Past studies have examined the teleconnections between Antarctic air temperatures and climate indices, such as the Southern Annular Mode (SAM) and Southern Oscillation Index (SOI). As a demonstration of the ease with which similar studies could be done using the condensed anomaly databases, this dissertation queried for high and low LST anomalies and then evaluated the anomaly populations in relation to the SAM, SOI, and sea ice extent. Using cross correlation functions, the current study found strong no-lag correlations (without attempting to pinpoint causality) between the SAM and the sample population of extreme LSTs (Figure 4.19). This result is in agreement with the findings of Kwok and Comiso [2002].

In contrast, no correlations were found between the SOI and 5th percentile (low) LST anomalies. However, a statistically significant positive correlation was found between the SOI

and 95th percentile (warm) anomalies. The correlation shows a 5-6 month lead time, i.e., changes to the SOI lead LST anomaly populations. The conflicting correlations, less distinct than with the SAM, are also in line with past studies [Turner, 2004; Smith and Stearns, 1993]

No correlations were found between sea ice extent and LST anomalies.

Importantly, the time required to perform the queries and retrieve the anomaly data was minimal. This demonstration show that, when using the condensate database methods, the bulk of data exploration effort is now in designing the experiments, creating publication-quality images, and interpretation.

5.1.6 Case Studies

Prompted by the results from the condensate database queries, the case studies (Section 4.5) performed expanded scientific investigations of the LST data. The case studies augmented the LST extremes and sample population information with auxiliary data from reanalyses, ground observations, and visible satellite imagery, as well as the uncondensed LST dataset.

An analysis of the LST trends around the South Pole showed no statistical significance, despite the general appearance of a cooling trend in the baseline linear regressions (Section 4.5.1). While individual grid cells in this region appear to be cooling, the interannual variability is very high, and one unusually cold winter is responsible for much of the apparent trend. Observations of surface air temperature from the South Pole Meteorology Department show an opposite trend, but again without statistical significance.

A look at an extreme cold LST event at high elevation (Section 4.5.2) confirmed previous work that evaluated the conditions leading up to such events [Turner et al., 2009; Campbell et

al., 2016]. The coldest land surface temperatures in Antarctica (during the timespan of the MODIS dataset) are shown to occur in shallow valleys downslope from the East Antarctic Ridge and, in particular, Dome Argus (Figure 4.27). In the example event, clear skies over an extended period, associated with light winds and a high pressure system over East Antarctica, permitted upwelling longwave radiation to cool the surface. Lack of warm air advection kept sensible heat transfer to a minimum, and cold air, flowing downhill from higher elevations, pooled in topographic low points.

A different case study examined extreme warm LSTs at high elevations and attributed causality to a somewhat different set of conditions (Section 4.5.3). While low advection rates also played a role, insolation, surface type, and aspect angle were critical. During calm, clear summer days, bare (low albedo) north-facing slopes in the Transantarctic Mountains showed the highest surface temperatures. This dissertation speculates that if some of these locations were instrumented they would likely set new air temperature records for high elevation Antarctica. Finding extreme high LSTs at lower elevations proved more difficult because warm, advected air masses are often responsible for the high temperatures; this warm air is typically accompanied by cloud cover, precluding IR satellite observation of the surface.

5.1.6.1 Cloud Masking Irregularities

A final case study evaluated a change in MODIS cloud detections (Section 4.5.4). The total population of LST observations, with and without cloud contamination, is a byproduct of the data condensation process, as is the sample population of cloud-free LSTs. The ratio of cloud-free samples to total observations provides a rough measure of clear skies. The fraction of cloud-free LSTs shows an irreconcilable change starting in 2011 (Figure 4.36). The increase in

cloud-free samples is not specific to the Terra MODIS sensor since a similar effect occurs within Aqua data as well (Figure 4.38). When MODIS LST collections are compared against each other the discontinuity in cloud detections is not present, demonstrating that the effect is not a natural event but most likely a defect in the Collection 6 cloud detection processing (Figure 4.39). Further testing suggests that the problem is centered around cloud detections over bright or snow-covered surfaces. Collectively, the evidence suggests that a systematic problem has occurred in NASA's Product Generation Executive or other code related to cloud detections. Since an inconsistency in cloud detections will influence cloud contamination (and thus the bias in LSTs), this revelation calls into question the accuracy of any time-series statistics generated using MODIS Collection 6 LSTs.

5.2 Applications

The Condensate Database Project and its associated methodologies demonstrate an alternative, highly efficient process for scientific data exploration and quality assurance. The methods are especially useful when applied to massive datasets. In addition to these features, the base concept of condensing a data stream has commercial, public safety, and military applications that require high-speed decision making but are constrained by low-bandwidth capabilities.

In the commercial sector, the ability to distill data to its essential core has potential value in the space, transportation, and finance industries. For example, the advent of small "CubeSat" satellites is revolutionizing the space industry, yet these miniature satellites often have limited downlink capabilities due to power and bandwidth constraints. In satellite-based data collection scenarios where real-time or near-real-time information is needed, a

constellation of CubeSats may utilize condensed data techniques – broadcasting only the data that exceeds a baseline threshold – to provide immediate, time-critical information. Storm front locations, for instance, may be directly broadcast from satellites to commercial aircraft or common carriers, enhancing avoidance of severe weather. The distillation of real-time geo-tagged social media trends may also alert law enforcement and public safety personnel to emergency situations. Within the finance industry, transaction speed is critical; any system that can rapidly discriminate between unneeded and crucial bits of information may provide a wealth of benefits.

On the military side, battlespace awareness relies on rapid, real-time intelligence to detect threats and targets. In a cluttered, dynamic operational theater, a condensed flow of information from source to soldier may save lives and promote successful mission accomplishments.

The ability to assess a dataset quickly for possible quality issues is perhaps the most immediately useful outcome. Earth Science datasets, especially the ever-growing collection of satellite observations, are often examined for quality and accuracy. However, the QC techniques and thoroughness are inconsistent. As demonstrated in this dissertation, even reputedly high-quality data, the subject of extensive evolution, testing and supporting literature, may have hidden internal flaws. The *Condensate Database Project* offers a starting point for a rapid, standardized methodology of verifying data quality. This capability comes at a time when the veracity of observational data is under constant scrutiny. More than ever, the scientific community has a responsibility to use techniques such as these to investigate and ensure data quality.

5.3 Future Work

The condensate databases provide a tool for rapid data exploration and a new perspective on data quality. As the *Condensate Database Project* progressed, the utility was further enhanced with the addition of DEM information, sample population data, and expanded descriptive statistics. While the condensate databases are clearly not designed to answer all questions or address all types of scientific analyses, specific topics of research are expedited.

The output from queries against the anomaly databases were images and CSV files, but not data from the full, uncondensed dataset. Thus, to complete most scientific analyses, the project required one additional tool: a separate application to subset and report the original, uncondensed source data for an area of interest. While it was not initially part of the project, a program with a command-line user interface was created to fulfill this need. The tool accepts date ranges, bounding boxes of grid cells and MODIS tile numbers as inputs, and delivers the requested time series of LST data, uncondensed, as an output. The data output format is compatible with ancillary analysis applications such as Excel, MATLAB, and R. Future developments to the Condensate Database Project should formalize and refine a similar application to subset the uncondensed data. The application should also have the option of applying the algorithmic cleaning filters to avoid reporting unneeded data.

One important shortcoming of the anomaly databases is the inability to reject queries that would generate too many results. For example, it is possible – indeed, likely – that at some point a user will submit a query that effectively asks for all anomalies for all time and space. While the anomaly databases represent a substantial condensation of the original dataset, this type of query may still attempt to report billions of data entries. In its current form, using

PostgreSQL databases on a single server, the database software will consume all available resources in an attempt to fulfill the query. Eventually a memory fault occurs. While this does not physically harm the server, and generally does not adversely affect the integrity of the databases, it does cause the server to appear hung for many minutes, perhaps hours, until it resets itself. During this time, other applications running on the server (including user command line response) may stall. Attempts to interrupt such queries take database administrator-level privileges, and occasionally cause the database software to hang in an unresolved state. The latter outcome necessitates a complete restart of the server processes. This “infinite query” problem remains unresolved, and currently users must remain vigilant about posing overly-broad queries. In a production setting, either the database software or the user interface should be intelligent enough to intervene when such queries are requested.

Of the scientific investigations undertaken in the Case Studies section, perhaps the most critical is determining the underlying cause of the cloud fraction inconsistency (Section 4.5.4). The results found in this dissertation will be forwarded to NASA for further review.

References

Ackerman, S. A., Holz, R. E., Frey, R., Eloranta, E. W., Maddux, B. C., & McGill, M. (2008). Cloud detection with MODIS. Part II: Validation. *Journal of Atmospheric and Oceanic Technology, 25*(7), 1073-1086. doi:10.1175/2007JTECHA1053.1

Ackerman, S. A., Strabala, K. I., Menzel, W. P., Frey, R. A., Moeller, C. C., & Gumley, L. E. (1998). Discriminating clear sky from clouds with MODIS. *Journal of Geophysical Research-Atmospheres, 103*(D24), 32141-32157. doi:10.1029/1998JD200032

AMRC. (2017). *South Pole Surface Observations, 1988-2016*. Retrieved from: ftp://amrc.ssec.wisc.edu/pub/southpole/surface_observations/

Bamber, J. L., Gomez-Dans, J. L., & Griggs, J. A. (2009). A new 1 km digital elevation model of the Antarctic derived from combined satellite radar and laser data - Part 1: Data and methods. *Cryosphere, 3*(1), 101-111. doi:10.5194/tc-3-101-2009

Baum, B. A., Kratz, D. P., Yang, P., Ou, S. C., Hu, Y. X., Soulen, P. F., & Tsay, S. C. (2000). Remote sensing of cloud properties using MODIS airborne simulator imagery during SUCCESS 1. Data and models. *Journal of Geophysical Research-Atmospheres, 105*(D9), 11767-11780. doi:10.1029/1999JD901089

Bosilovich, M. G., Lucchesi, R., & Suarez, M. (2016). MERRA-2: File Specification. In *Office Note No. 9 (Version 1.1)* (pp. 73): NASA Global Modeling and Assimilation Office (GMAO).

Bracegirdle, T. J., Connolley, W. M., & Turner, J. (2008). Antarctic climate change over the twenty first century. *Journal of Geophysical Research-Atmospheres, 113*(D3). doi:10.1029/2007JD008933

Brodzik, M. J., Billingsley, B., Haran, T., Raup, B., & Savoie, M. H. (2012). EASE-Grid 2.0: Incremental but Significant Improvements for Earth-Gridded Data Sets. *Isprs International Journal of Geo-Information, 1*(1), 32-45. doi:10.3390/ijgi1010032

Bromwich, D. H., Nicolas, J. P., Monaghan, A. J., Lazzara, M. A., Keller, L. M., Weidner, G. A., & Wilson, A. B. (2013). Central West Antarctica among the most rapidly warming regions on Earth. *Nature Geoscience, 6*(2), 139-145. doi:10.1038/NGEO1671

Campbell, G., Pope, A., Lazzara, M. A., & Scambos, T. A. (2013). *The Coldest Place On Earth: −90° C and Below in East Antarctica from Landsat-8 and other Thermal Sensors*. Paper presented at the American Geophysical Union, Fall Meeting 2013, San Francisco, CA.

Chapman, W. L., & Walsh, J. E. (2007). A synthesis of Antarctic temperatures. *Journal of Climate, 20*(16), 4096-4117. doi:10.1175/JCLI4236.1

Coles, S., Bawa, J., Trenner, L., & Dorazio, P. (2001). *An Introduction to Statistical Modeling of Extreme Values* (Vol. 208). London: Springer.

Comiso, J. C. (2000). Variability and trends in Antarctic surface temperatures from in situ and satellite infrared measurements. *Journal of Climate, 13*(10), 1674-1696. doi:10.1175/1520-0442(2000)013<1674:VATIAS>2.0.CO;2

Comiso, J. (2003). Warming trends in the Arctic from clear sky satellite observations. *Journal of Climate, 16*(21), 3498-3510.

Dash, P., Gottsche, F. M., Olesen, F. S., & Fischer, H. (2002). Land surface temperature and emissivity estimation from passive sensor data: theory and practice-current trends. *International Journal of Remote Sensing, 23*(13), 2563-2594. doi:10.1080/01431160110115041

Dee, D. P., et al. (2011). The ERA-Interim reanalysis: configuration and performance of the data assimilation system. *Quarterly Journal of the Royal Meteorological Society, 137*(656), 553-597. doi:10.1002/qj.828

Ding, Q. H., Steig, E. J., Battisti, D. S., & Kuttel, M. (2011). Winter warming in West Antarctica caused by central tropical Pacific warming. *Nature Geoscience, 4*(6), 398-403. doi:10.1038/NGEO1129

Dozier, J., & Warren, S. G. (1982). Effect of Viewing Angle on the Infrared Brightness Temperature of Snow. *Water Resources Research, 18*(5), 1424-1434. doi:10.1029/WR018i005p01424

Drewry, D. J., Jordan, S. R., & Jankowski, E. (1982). Measured properties of the Antarctic ice sheet: surface configuration, ice thickness, volume and bedrock characteristics. *Annals of Glaciology, 3*(1).

Easterling, D. R., Kunkel, K. E., Wehner, M. F., & Sun, L. (2016). Detection and Attribution of Climate Extremes in the Observed Record. *Weather and Climate Extremes, 11*, 17-27.

Eisenman, I., Meier, W. N., & Norris, J. R. (2014). A spurious jump in the satellite record: has Antarctic sea ice expansion been overestimated? *Cryosphere, 8*(4), 1289-1296. doi:10.5194/tc-8-1289-2014

Fetterer, F., Knowles, K., Meier, W., & Savoie, M. (2017). *Sea Ice Index*.

Gallaher, D. W., & Grant, G. E. (2012). Data Rods: High Speed, Time-Series Analysis of Massive Cryospheric Data Sets Using Pure Object Databases. *2012 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, 5297-5300. doi:10.1109/IGARSS.2012.6352413

GMAO. (2015). *MERRA-2: tavg1_2d_slv_Nx: 2d, 1-Hourly, Time-Averaged, Single-Level, Assimilation, Single-Level Diagnostics*.

Gong, D. Y., & Wang, S. W. (1999). Definition of Antarctic Oscillation Index. *Geophysical Research Letters, 26*(4), 459-462. doi:10.1029/1999GL900003

Grant, G. E. (2012). *An Analysis of Greenland Ice Sheet Variability Using 25 Years of AVHRR Polar Pathfinder Data.* (M.A., Geography), University of Colorado at Boulder, Boulder, CO. Retrieved from http://0-search.proquest.com.libraries.colorado.edu/docview/1112071640?accountid=14503 (1519557)

Grant, G. E., & Gallaher, D. W. (2015). A Novel Technique for Time-Centric Analysis of Massive Remotely-Sensed Datasets. *Remote Sensing, 7*(4), 3986-4001. doi:10.3390/rs70403986

Hall, D. K., Box, J. E., Casey, K. A., Hook, S. J., Shuman, C. A., & Steffen, K. (2008). Comparison of satellite-derived and in-situ observations of ice and snow surface temperatures over Greenland. *Remote Sensing of Environment, 112*(10), 3739-3749. doi:10.1016/j.rse.2008.05.007

Hall, D. K., Comiso, J. C., DiGirolamo, N. E., Shuman, C. A., Key, J. R., & Koenig, L. S. (2012). A Satellite-Derived Climate-Quality Data Record of the Clear-Sky Surface Temperature of the Greenland Ice Sheet. *Journal of Climate, 25*(14), 4785-4798. doi:10.1175/JCLI-D-11-00365.1

Haran, T. M., Campbell, G. G., Scambos, T. A., & Pope, A. (2016). *Coldest Place on Earth: New MODIS and Landsat 8 Thermal Data and Detailed Time Series of Cold Events.* Paper presented at the *AGU Fall Meeting*.

Hartmann, D. L., et al. (2013). Observations: Atmosphere and Surface. In *In Climate Change 2013 the Physical Science Basis: Working Group I Contribution to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*: Cambridge University Press.

Hori, M., Aoki, T., Tanikawa, T., Motoyoshi, H., Hachikubo, A., Sugiura, K., and Takahashi, F. (2006). In-situ measured spectral directional emissivity of snow and ice in the 8-14 mu m atmospheric window. *Remote Sensing of Environment, 100*(4), 486-502. doi:10.1016/j.rse.2005.11.001

Key, J. R., Collins, J. B., Fowler, C., & Stone, R. S. (1997). High-latitude surface temperature estimates from thermal satellite data. *Remote Sensing of Environment, 61*(2), 302-309. doi:10.1016/S0034-4257(97)89497-7

King, J. C., & Anderson, P. S. (1994). Heat and Water-Vapor Fluxes and Scalar Roughness Lengths over an Antarctic Ice Shelf. *Boundary-Layer Meteorology, 69*(1-2), 101-121. doi:10.1007/BF00713297

Koenig, L. S., & Hall, D. K. (2010). Comparison of satellite, thermochron and air temperatures at Summit, Greenland, during the winter of 2008/09. *Journal of Glaciology, 56*(198), 735-741.

Krause, P. F., & Flood, K. L. (1997). *Weather and Climate Extremes*. Alexandria, VA, USA: Army Topographic Engineering Center.

Kwok, R., & Comiso, J. C. (2002). Spatial patterns of variability in Antarctic surface temperature: Connections to the Southern Hemisphere Annular Mode and the Southern Oscillation. *Geophysical Research Letters, 29*(14). doi:10.1029/2002GL015415

Lewis, S. C., & King, A. D. (2017). Evolution of mean, variance and extremes in 21st century temperatures. *Weather and Climate Extremes*(15), 1-10.

Li, S. (2017). Data Reduction Techniques for Scientific Visualization and Data Analysis. *STAR, 36*(*3*).

Liu, Y. H., Ackerman, S. A., Maddux, B. C., Key, J. R., & Frey, R. A. (2010). Errors in Cloud Detection over the Arctic Using a Satellite Imager and Implications for Observing Feedback Mechanisms. *Journal of Climate, 23*(7), 1894-1907. doi:10.1175/2009JCLI3386.1

Liu, Y. H., Key, J. R., Frey, R. A., Ackerman, S. A., & Menzel, W. P. (2004). Nighttime polar cloud detection with MODIS. *Remote Sensing of Environment, 92*(2), 181-194. doi:10.1016/j.rse.2004.06.004

Marshall, G. (2016). The Climate Data Guide: Marshall Southern Annular Mode (SAM) Index (Station-based). Retrieved from https://climatedataguide.ucar.edu/climate-data/marshall-southern-annular-mode-sam-index-station-based

Marshall, G. (2007). Half-century seasonal relationships between the Southern Annular Mode and Antarctic temperatures. *International Journal of Climatology, 27*(3), 373-383. doi:10.1002/joc.1407

Marshall, G. J., Orr, A., van Lipzig, N. P., & King, J. C. (2006). The impact of a changing Southern Hemisphere Annular Mode on Antarctic Peninsula summer temperatures. *Journal of Climate*, 19(20), 5388-5404.

Marshall, G. J., & King, J. C. (1998). Southern hemisphere circulation anomalies associated with extreme Antarctic Peninsula winter temperatures. *Geophysical Research Letters, 25*(13), 2437-2440. doi:10.1029/98GL01651

McCormick, R. A. (1958). An estimate of the minimum possible surface temperature at the South Pole. *Monthly Weather Review, 86*(1), 1-5.

Meyer, H., Katurji, M., Appelhans, T., Muller, M. U., Nauss, T., Roudier, P., & Zawar-Reza, P. (2016). Mapping Daily Air Temperature for Antarctica Based on MODIS LST. *Remote Sensing, 8*(9). doi:10.3390/rs8090732

Miliaresis, G. C. (2014). Spatiotemporal patterns of land surface temperature of Antarctica from MODIS monthly LST (MYD11C3) data. *Journal of Spatial Science, 59*(1), 157-166.

Monaghan, A. J., Bromwich, D. H., Chapman, W., & Comiso, J. C. (2008). Recent variability and trends of Antarctic near-surface temperature. *Journal of Geophysical Research-Atmospheres, 113*(D4). doi:10.1029/2007JD009094

NASA (2016). MODIS Specifications. Retrieved from https://modis.gsfc.nasa.gov/about/specifications.php

NASA LP DAAC (2017) Land Surface Temperature and Emissivity Daily L3 Global 1 km Grid SIN. Version 6. NASA EOSDIS Land Processes DAAC, USGS Earth Resources Observation and Science (EROS) Center, Sioux Falls, South Dakota (https://lpdaac.usgs.gov), accessed March 1, 2017, at https://search.earthdata.nasa.gov/search?q=MOD11A1+V006.

NCEI. (2017). Southern Oscillation Index (SOI). https://www.ncdc.noaa.gov/teleconnections/enso/indicators/soi/

Nghiem, S. V., Hall, D. K., Mote, T. L., Tedesco, M., Albert, M. R., Keegan, K., Neumann, G. (2012). The extreme melt across the Greenland ice sheet in 2012. *Geophysical Research Letters, 39*. doi:10.1029/2012GL053611

Peck, L. S., Webb, K. E., & Bailey, D. M. (2004). Extreme sensitivity of biological function to temperature in Antarctic marine species. *Functional Ecology, 18*(5), 625-630. doi:10.1111/j.0269-8463.2004.00903.x

Petty, G. W. (2006). *A first course in atmospheric radiation* (2nd ed.). Madison, Wis.: Sundog Pub.

Price, J. C. (1984). Land Surface-Temperature Measurements from the Split Window Channels of the NOAA-7 Advanced Very High-Resolution Radiometer. *Journal of Geophysical Research-Atmospheres, 89*(ND5), 7231-7237. doi:10.1029/JD089iD05p07231

Quayle, W. C., Peck, L. S., Peat, H., Ellis-Evans, J. C., & Harrigan, P. R. (2002). Extreme responses to climate change in Antarctic lakes. *Science, 295*(5555), 645-645. doi:10.1126/science.1064074

Scambos, T. A., Bohlander, J. A., Shuman, C. A., & Skvarca, P. (2004). Glacier acceleration and thinning after ice shelf collapse in the Larsen B embayment, Antarctica. *Geophysical Research Letters, 31*(18). doi:10.1029/2004GL020670

Schneider, D. P., Deser, C., & Okumura, Y. (2012). An assessment and interpretation of the observed warming of West Antarctica in the austral spring. *Climate Dynamics, 38*(1-2), 323-347. doi:10.1007/s00382-010-0985-x

Schneider, D. P., Steig, E. J., & Comiso, J. C. (2004). Recent climate variability in Antarctica from satellite-derived temperature data. *Journal of Climate, 17*(7), 1569-1583. doi:10.1175/1520-0442(2004)017<1569:RCVIAF>2.0.CO;2

Schneider, D. P., Steig, E. J., van Ommen, T. D., Dixon, D. A., Mayewski, P. A., Jones, J. M., & Bitz, C. M. (2006). Antarctic temperatures over the past two centuries from ice cores. *Geophysical Research Letters, 33*(16). doi:10.1029/2006GL027057

Serreze, M. C., & Barry, R. G. (2014). *The Arctic Climate System*: Cambridge University Press.

Shliakhov, V. I. (1958). "0 Minimal'nykh temperaturakh v Antarktide" (On Minimum Temperatures in Antarctica). *Meteorologiia i Gidrologiia, 4*(1058), 5-7.

Shuman, C. A., Comiso, J. C., Winther, J. G., & Solberg, R. (2002). In situ and satellite surface temperature records in Antarctica. *Annals of Glaciology, Vol 34, 2002, 34*, 113-120. doi:10.3189/172756402781818003

Skansi, M., et al. (2017.). Evaluating highest-temperature extremes in the Antarctic. *EOS*, 01 March 2017.

Smith, S. R., & Stearns, C. R. (1993). Antarctic Pressure and Temperature Anomalies Surrounding the Minimum in the Southern Oscillation Index. *Journal of Geophysical Research-Atmospheres, 98*(D7), 13071-13083. doi:10.1029/92JD02157

Sobrino, J. A., Kharraz, J. E., & Li, Z. L. (2003). Surface temperature and water vapour retrieval from MODIS data. *International Journal of Remote Sensing, 24*(24), 5161-5182. doi:10.1080/0143116031000102502

Steig, E. J., Schneider, D. P., Rutherford, S. D., Mann, M. E., Comiso, J. C., & Shindell, D. T. (2009). Warming of the Antarctic ice-sheet surface since the 1957 International Geophysical Year. *Nature, 457*(7228), 459-462. doi:10.1038/nature07669

Stepanova, N. A. (1959). The lowest temperature in Antarctica. *Monthly Weather Review, 87*(4), 145-146.

Strabala, K. (2016). MODIS Cloud Mask User's Guide. Retrieved from http://modis-atmos.gsfc.nasa.gov/_docs/CMUSERSGUIDE.pdf

Su, Y., Agrawal, G., Woodring, J., Myers, K., Wendelberger, J., & Ahrens, J. (2013). *Taming massive distributed datasets: data sampling using bitmap indices.* Paper presented at the 22nd international symposium on High-performance parallel and distributed computing.

Thompson, D. W. J., & Solomon, S. (2002). Interpretation of recent Southern Hemisphere climate change. *Science, 296*(5569), 895-899. doi:10.1126/science.1069270

Turner, J. (2004). The El Nino-southern oscillation and Antarctica. *International Journal of Climatology, 24*(1), 1-31. doi:10.1002/joc.965

Turner, J., Anderson, P., Lachlan-Cope, T., Colwell, S., Phillips, T., Kirchgaessner, A., Orr, A. (2009). Record low surface air temperature at Vostok station, Antarctica. *Journal of Geophysical Research-Atmospheres, 114*. doi:10.1029/2009JD012104

Turner, J., Colwell, S. R., Marshall, G. J., Lachlan-Cope, T. A., Carleton, A. M., Jones, P. D., Iagovkina, S. (2004). The SCAR READER project: Toward a high-quality database of mean Antarctic meteorological observations. *Journal of Climate, 17*(14), 2890-2898. doi:10.1175/1520-0442(2004)017<2890:TSRPTA>2.0.CO;2

Turner, J., Lu, H., White, I., King, J. C., Phillips, T., Hosking, J. S., Deb, P. (2016). Absence of 21st century warming on Antarctic Peninsula consistent with natural variability. *Nature, 535*(7612), 411-+. doi:10.1038/nature18645

Urban, M., Eberle, J., Huttich, C., Schmullius, C., & Herold, M. (2013). Comparison of Satellite-Derived Land Surface Temperature and Air Temperature from Meteorological Stations on the Pan-Arctic Scale. *Remote Sensing, 5*(5), 2348-2367. doi:10.3390/rs5052348

Valor, E., & Caselles, V. (1996). Mapping land surface emissivity from NDVI: Application to European, African, and South American areas. *Remote Sensing of Environment, 57*(3), 167-184. doi:10.1016/0034-4257(96)00039-9

Vaughan, D. G., Marshall, G. J., Connolley, W. M., Parkinson, C., Mulvaney, R., Hodgson, D. A., et al. (2003). Recent rapid regional climate warming on the Antarctic Peninsula. *Climatic Change, 60*(3), 243-274. doi:10.1023/A:1026021217991

Wan, Z. M. (1999). MODIS land-surface temperature algorithm theoretical basis document (LST ATBD). In (pp. 75). Santa Barbara, CA, USA: Institute for Computational Earth System Science.

Wan, Z. M. (2008). New refinements and validation of the MODIS Land-Surface Temperature/Emissivity products. *Remote Sensing of Environment, 112*(1), 59-74. doi:10.1016/j.rse.2006.06.026

Wan, Z. M. (2013). Collection-6 MODIS Land Surface Temperature Products Users' Guide. In (pp. 33). University of California, Santa Barbara CA: Earth Research Institute.

Wan, Z. M. (2014). New refinements and validation of the collection-6 MODIS land-surface temperature/emissivity product. *Remote Sensing of Environment, 140*, 36-45. doi:10.1016/j.rse.2013.08.027

Wan, Z. M., & Dozier, J. (1996). A generalized split-window algorithm for retrieving land-surface temperature from space. *IEEE Transactions on Geoscience and Remote Sensing, 34*(4), 892-905. doi:10.1109/36.508406

Wan, Z. M., Zhang, Y., Zhang, Q., & Li, Z. L. (2002). Validation of the land-surface temperature products retrieved from Terra Moderate Resolution Imaging Spectroradiometer data. *Remote Sensing of Environment, 83*(1-2), 163-180. doi:10.1016/S0034-4257(02)00093-7

Wan, Z. M., Zhang, Y., Zhang, Q., & Li, Z. L. (2004). Quality assessment and validation of the MODIS global land surface temperature. *International Journal of Remote Sensing, 25*(1), 261-274. doi:10.1080/0143116031000116417

Wang, W., Liang, S., & Meyers, T. (2008). Validating MODIS land surface temperature products using long-term nighttime ground measurements. *Remote Sensing of Environment, 112*(3), 623-635. doi:10.1016/j.rse.2007.05.024

Warren, S. G. (1982). Optical Properties of Snow. *Reviews of Geophysics, 20*(1), 67-89. doi:10.1029/RG020i001p00067

Weatherly, J. W., Walsh, J. E., & Zwally, H. J. (1991). Antarctic Sea Ice Variations and Seasonal Air-Temperature Relationships. *Journal of Geophysical Research-Oceans, 96*(C8), 15119-15130. doi:10.1029/91JC01432

Westermann, S., Langer, M., & Boike, J. (2012). Systematic bias of average winter-time land surface temperatures inferred from MODIS at a site on Svalbard, Norway. *Remote Sensing of Environment, 118*, 162-167. doi:10.1016/j.rse.2011.10.025

Wexler, H. (1959). Note on Lowest Antarctic Temperature Estimated by Shliakhov. *Monthly Weather Review, 87*(4), 147.

WMO. (2008). Guide to Meteorological Instruments and Methods of Observation. In *WMO-No. 8* (Seventh ed.): World Meteorological Organization.

WMO. (2017). Esperanza Temp [Press release]. Retrieved from https://public.wmo.int/en/media/press-release/wmo-verifies-highest-temperatures-antarctic-region

Xiong, X. X., & Barnes, W. (2006). An overview of MODIS radiometric calibration and characterization. *Advances in Atmospheric Sciences, 23*(1), 69-79. doi:10.1007/s00376-006-0008-3

Appendix A. Database System Design and Performance

The baseline statistics and condensed datasets exist as plain-text CSV files, independent from any database structures. While the data could be searched and used in this simple format, a primary goal of the *Condensate Database Project* was to ingest the condensed data into a series of databases that would enable immediate exploration without the overhead or coding necessary for processing flat files. The following sections describe the considerations, architecture, and results of this effort.

A.1 Database Architecture

Multiple open-source database systems were evaluated and discarded before PostgreSQL was chosen. The other database systems were H2, Perst, MongoDB, and MySQL. Ultimately, PostgreSQL version 9.4.5 fulfilled the project's needs with the best performance.

Key considerations in the selection and design of the condensed databases included:

- Speed of data retrieval
- Industry standard query formatting
- Database size
- Ability to handle billions of entries
- Ease of administration
- Use of open-source database software

Note that *speed of ingestion* does not appear on this list; an underlying assumption is that front-end processing (that of creating the condensed databases in the first place) will be transparent to the user, so the speed at which they are created is not important so long as it can be accomplished in a reasonable amount of time. The generation of the CSV datasets was,

in fact, an intermediate step intended to accelerated database creation. Specifically, our final choice of database software, PostgreSQL, can ingest large CSV files much more rapidly than piecemeal insertion of each anomalous data point.

The condensation algorithm, when applied to MODIS LST data, typically condensed the source data points down to roughly 3-4% of the original population. However, the source data originated from flat image files, essentially matrices where any pixel's geographical location was immediately known from its row and column position in the matrix: When an individual pixel is extracted from an image, as the condensation process does, the location information is lost. Consequently, the original row and column location of each pixel must also be stored in the database (and CSV files), along with a table for translating the rows and columns into latitudes and longitudes. A similar limitation holds true for the observation time of any pixel, which is originally contained only in the image file's name. The database, then, must store references to the location and time for each pixel, plus tables of the locations and timestamps, requirements that "re-inflate" the size of the condensed data. As a result, although the source pixels may have been condensed down to just 3-4% of the original volume, the physical disk space occupied by the databases is often 40-60% of the original source CDR files.

Because physically large databases can become administratively unwieldy, keeping the databases compact was an important goal. The selection of data storage types (e.g., big integers versus small integers) was driven by this need. Likewise, some data values were encoded to save space, and references to location and timestamp table entries were used instead of storing the actual latitudes, longitudes, and times with every anomalous pixel. The data formats and table joins are shown in Figure 3.2. Table 2 through Table 6 provide additional

detail for each of the database tables and variables shown in Figure 3.2. Note that within PostgreSQL a “Small Integer”, or SMALLINT, is a 2-byte signed integer, an INT is a 4-byte signed integer, BIGINTs are 8-byte signed integers, FLOAT (or “REAL”) is a 4-byte floating point number, and DOUBLE PRECISION (or simply “DOUBLE”) is an 8-byte floating point number.

Table 2. Metadata table, describing the contents of a condensed database.

Variable	Format	Encoding	Description
Observations	BIGINT	none	The total population of dataset observations before cleaning but after the application of a land mask.
Samples	BIGINT	none	Population of Observations remaining after dataset cleaning.
Anomalies	BIGINT	none	Population of Samples that were determined to be anomalies.
Timestamps	SMALLINT	none	Number of unique timestamps where anomalies were found.
Locations	SMALLINT	none	Number of possible locations before land masking.
Valid Locs	SMALLINT	none	The number of valid locations remaining after application of the land mask.
Rows	SMALLINT	none	Number of rows in the source images.
Cols	SMALLINT	none	Number of columns in the source images.

Table 3. Locations table. Each pixel location in the source images is described by this table.

Variable	Format	Encoding	Description
ID	INT	none	A unique identification reference for each pixel location.
Row	SMALLINT	none	Source image row.
Col	SMALLINT	none	Source image column.
Lat	DOUBLE PRECISION	none	Latitude.
Lon	DOUBLE PRECISION	none	Longitude.
Elev	SMALLINT	none	Elevation in meters.
Htile	SMALLINT	none	Source tile horizontal number.
Vtile	SMALLINT	none	Source tile vertical number.

Table 4. Timestamps table. Each anomalous pixel entry refers to an entry in this table.

Variable	Format	Encoding	Description
ID	SMALLINT	none	A reference identification for each unique anomaly timestamp.
EpochDate	FLOAT	none	The number of days since January 1, 1970, as a decimal value.
Year	SMALLINT	none	Year of the observation.
Month	SMALLINT	none	Month of the observation.
Day	SMALLINT	none	Day of the month.
DOY	SMALLINT	none	Day of year.
Time	SMALLINT	none	Hour of the observation.

Table 5. Anomalies table. Each anomalous LST is one table entry with these values.

Variable	Format	Encoding	Description
Value	SMALLINT	*10	The LST value in degrees K, encoded by a factor of 10. Example: 2730 represents 273.0 K.
LocID	INT	none	Refers to a unique entry in the Locations table.
TimeID	SMALLINT	none	Refers to a unique entry in the Timestamps table.

Table 6. Stats table. The baseline statistics are generated against the cleaned dataset, before condensation, across the timespan of the database (i.e., interannual season, month, or year). The population variables are added after creation of the anomaly databases.

Variable	Format	Encoding	Description
Row	SMALLINT	none	Pixel row in the source image.
Col	SMALLINT	none	Pixel column in the source image.
LocID	INT	none	Location table entry identifier.
Mean	DOUBLE PRECISION	*10	Mean LST value.
SD	DOUBLE PRECISION	*10	Standard deviation of LSTs.
FifthPct	SMALLINT	*10	Fifth percentile LST.
Median	SMALLINT	*10	Median LST.
NinetyFifthPct	SMALLINT	*10	Ninety fifth percentile LST.
Min	SMALLINT	*10	Absolute minimum.
Max	SMALLINT	*10	Absolute maximum.
Slope	SMALLINT	*1000	Slope (Beta1) of a simple linear regression of the time series of samples.

Corr	SMALLINT	*1000	Correlation coefficient (R) for the linear regression.
Observations	SMALLINT	none	The number of observations at this location prior to dataset cleaning.
Samples	SMALLINT	none	The number of observations at this location after dataset cleaning.
Anomalies	SMALLINT	none	The number of anomalies found within the samples.

A.2 Database Administration

Testing demonstrated that PostgreSQL databases could handle one billion table entries or more, although performance becomes sluggish at those volumes. To improve performance, many individual anomaly databases were created instead of one massive, all-encompassing database. These smaller databases were subdivided by MODIS tile and temporal span (interannual months, seasons, and full years). With databases covering 4 seasons, 12 months, and one full year for each of the 23 MODIS tiles, the result was 391 total databases. Subsequent evaluation revealed that most queries were aimed at anomalies that exceeded either the upper or lower percentile, but not both. Further subdividing the databases into low anomalies ($<5^{th}$ percentile) and high anomalies ($>95^{th}$ percentile) improved performance but doubled the number of databases to 782. This number of databases is, of course, wholly dependent on the organization, span, and spatiotemporal density of the source data being condensed. Conceivably, an entire global high-resolution dataset could be condensed into many thousands of contiguous databases.

Administratively, this large number of databases is not difficult to manage, although some shell scripting became necessary for Database Management System (DBMS) commands needing execution across all of them. The subdivision of databases also improved (although did

not fully cure) another problem, that of the user inadvertently querying for too much information. The smaller dataset sizes did not prevent exceeding computing memory resources, but did reduce the amount of time between query and error. In general, when a memory error occurred the query would abort. On a few occasions, administrative action was required to restart the DBMS software. Another unresolved issue is that of PostgreSQL daemon lock-up, where the background database manager process would continue trying to execute a query long after resources had been exhausted and never stop. At these times, the Linux process prioritization levels caused every process on the computer to stall. Fortunately, this occurrence was rare, but it revealed a potential problem with the current configuration: The Linux OS tagged the DBMS process as unstoppable, even for system administrators, necessitating a reboot. At no time were the databases corrupted due to a misguided query, although DBMS restarts sometimes required several minutes to complete.

The large number of databases was transparent to users. The ancillary software written to assist with queries opened and closed databases as needed, without user knowledge or intervention. This functionality would be a requirement for any operational version of the condensate databases since manually querying one or more of 700+ databases would be extremely tedious. Importantly, it proved much easier to query across multiple databases if they were subdivided spatially rather than temporally. For example, attempts to take the 17-year span of MODIS anomaly data and subdivide it into databases containing, say, 5-year chunks, became confusing. Using temporal divisions, high performance queries needed parsing so the software could understand the temporal span and then decide which databases were required. The alternative was to open and query more databases than needed, drastically

slowing response times. If the number of database entries is growing too large, this dissertation recommends further subdivision by space rather than time, thus preserving the full time series of data within each database.

A.3 Data Ingestion

Data ingestion was divided into two parts: 1) baseline statistics calculations, and 2) condensing the dataset using the baselines as a reference. Dataset cleaning was performed on-the-fly in the same manner in both steps. No secondary "cleaned" dataset was created. The output from the baseline stage was a series of CSV files containing the statistical information. Originally, the output from the condensing stage was directly inserted into PostgreSQL databases. However, the entry-by-entry insertion proved too time consuming due to DBMS overhead and was abandoned. The final method for creating the databases was the creation of plain-text CSV files, each containing all the anomalies in an individual database, followed by discrete DBMS commands that fed whole CSV files into PostgreSQL. This last step, ingesting a CSV file into a PostgreSQL database, was highly efficient, typically taking no more than a few seconds.

The speed of data ingestion is, of course, platform, code, and dataset dependent. This dissertation used a series of Java executables, created and optimized using Eclipse on a 12-core SUSE Linux server, with 24GB of RAM and network drives housing the 17-year MODIS LST CDR. Benchmarking demonstrated that, on our system, performance could be enhanced by

multithreading (note: *not* multiprocessing) the baseline and condensation steps.[5] Baseline CSV generation using all 12 cores took approximately 2.5 days.

Condensate CSV file generation was accelerated when using up to 4 cores. Above 4 cores there reached a point of diminishing returns and potential resource overload. Using only two cores, condensing the dataset took approximately 4 days and did not risk memory errors.

Given more time, further performance enhancements may be expected with better memory management, coding optimization, and the possible use of a different programming language. Of particular note, Java does not currently support advanced parallel processing techniques such as MPI, OpenMP or OpenACC, which may offer substantial performance improvements. This dissertation recommends moving future versions of the *Condensate Database Project* to either Python or C/C++, expanding the multiprocessing options.

A.4 Data Retrieval

Data can be retrieved from any of the databases using standard SQL syntax. Multithreaded queries (querying multiple databases simultaneously) are effective, to a point, and accelerate results retrieval. Actual retrieval times are clouded by the overhead of storing the results as either CSV data or images, or both, which may be a substantial portion of the total query time. Appendix B shows some example queries and the time necessary for completion.

[5] Multithreading consists of executing multiple independent processes and then re-assembling the results within a single program. On a Linux server, a multithreaded query appears as if many different query programs are running at the same time, which in fact they are. Since there is overhead involved whenever a new process is started, this methodology can be somewhat slow. In contrast, multiprocessing (or "parallel processing") is executed from a single program with shared resources. Each query may be assigned its own CPU but have access to a shared memory block. With care, this approach can be a much more efficient use of computing resources and potentially accelerate query performance, assuming the system is not constrained by I/O bottlenecks.

In general, a single query against one database, retrieving only a small number of anomalies (<1000 pixels) may take just a few seconds. A large query involving the entire Antarctic continent and retrieving millions of anomalies may take 5 to 15 minutes or more. Queries that attempt to retrieve all pixels for all time may cause the system to fail.

Appendix B. Structured Query Language (SQL) Queries

The following sections show examples of queries used during the analysis portion of this dissertation. The timing metrics for some queries are inflated by approximately 50% due to the overhead of storing the results on disk as CSV files, and unique to the hardware configuration used.

B.1 Spatial Extremes Queries

The highest LST anomalies along the Antarctic Peninsula (Figure 4.12a) were determined using a simple query of the database containing anomalies for the Peninsula Region (tile coordinates h=15, v=15) during the austral summer season (DJF), and filtering for observations less than 2500 meters in elevation. The query includes a filter restricting anomalies to greater than or equal to 2382 (238.2 K, −35 °C). This single-threaded query returns 2,699,351 pixels and takes 30 seconds to complete:

```
SELECT * FROM ANOMALIES A, LOCATIONS L WHERE
A.VALUE>=2382 AND A.LOCID=L.ID AND L.ELEV<2500
```

A query for the highest continental LST anomalies at upper elevations (>2500 m) is similar. A variation on this query was used to identify hot spots in the Transantarctic Mountain Range (Figure 4.12b). To reduce the results set, a filter is set to exclude LSTs below 2682 (268.2 K, −5 °C). The query is posed to all databases containing the upper percentile anomalies during the austral summer season (23 total databases for continental coverage). A multithreaded query returns 600 anomalies and takes approximately 10 minutes:

```
SELECT * FROM ANOMALIES A, LOCATIONS L WHERE
A.VALUE>=2682 AND A.LOCID=L.ID and L.ELEV>2500
```

A query for high elevation extreme cold temperatures uses the austral winter seasonal databases (JJA). This query rapidly identifies the coldest LSTs in Antarctica where temperatures fall below −93 °C (180.2 K). A similar query was used to generate Figure 4.26. Multithreaded across 12 cores, it returns 71761 anomalies in 7 minutes:

```
SELECT * FROM ANOMALIES A, LOCATIONS L WHERE
A.VALUE<=1802 AND A.LOCID=L.ID AND L.ELEV>2500
```

B.2 Temporal Analysis Queries

The temporal analyses in this dissertation used a variety of monthly and yearly statistics. For instance, to create a graph of the monthly means, a query was executed the retrieved the statistics for a single month. A variation on this query was then repeated for each month in the timeline. The following example retrieves all statistical information available for a single month, for each grid cell (1,200,000 results) in less than 2 seconds:

```
SELECT * FROM STATS S, ANOMALIES A, TIMESTAMPS T
WHERE T.YEAR=2001 AND T.MONTH=1 AND A.LOCID=S.LOCID
AND A.TIMEID=T.ID
```

The results from each monthly query were later assembled in an Excel spreadsheet to produce the graphics.

Total observation populations and sample sizes were not originally part of the STATS table, only the total number of anomalies found. Analyses that relied on these values were initially conducted using custom code. However, the population statistics proved so valuable that subsequent versions of the condensed databases were expanded to include the full set of population numbers. The Observations, Samples, and Anomalies populations (see Table 6) were used in both the temporal analyses and the teleconnection portion of this dissertation.

B.3 Case Studies

The East Antarctic ridge (Figure 4.26) spans two MODIS tiles, h=18 v=17 and h=19 v=17. Rather than querying all tiles and seasons for the coldest temperatures, the study reduced the queries to the winter seasonal databases for these two tiles, e.g.:

```
SELECT * FROM ANOMALIES A, LOCATIONS L WHERE
A.VALUE<=1882 AND A.LOCID=L.ID AND L.ELEV>2500
```

A query for the high LSTs percentile focusing on the Transantarctic Mountains near the Drygalski Ice Tongue was also reduced to a single MODIS tile, h=22 v=16. This final query finds 105 extreme LST pixels in 35 seconds:

```
SELECT * FROM ANOMALIES A, LOCATIONS L WHERE
A.VALUE>=2682 AND A.LOCID=L.ID AND L.ELEV>2500
```

Other case studies, examining the long-term trends in South Pole surface temperatures and the frequency of cloudy pixels, primarily utilized the original source MODIS LST dataset. This approach is in keeping with the larger goals of the *Condensate Database Project*, where the condensed databases serve as a tool for data exploration and the starting point for more in-depth analyses.

Appendix C. Acronyms and Symbols

AVHRR Advanced Very High Resolution Radiometer, the 5-channel imager aboard NOAA satellites. Channels 4 and 5 measure in the TIR spectrum, at 11 μm and 12 μm.

AWS Automated Weather Station, an unmanned weather station that reports autonomously.

B Radiant intensity, $W/sr \cdot m^2$

B_λ Black body spectral radiant intensity, $W/sr \cdot m^2 \cdot \mu m$

c Speed of light, 3×10^8 m/s

CCF Cross Correlation Function, a statistical technique for comparing time series data.

CDR Climate Data Record, a time series of climate data, typically stored in daily files.

DAAC Distributed Active Archive Center.

DB Database.

DBMS Database Management System.

DJF December, January, and February, austral summer.

G Ground energy flux.

h Planck's constant, 6.626×10^{-34} J·s.

HDF Hierarchical Data Format, a self-describing file format.

H_{LE} Turbulent heat flux from latent energy (evaporation, sublimation).

H_S Turbulent sensible heat flux.

IR Infrared.

I_λ Spectral radiant intensity, $W/sr \cdot m^2 \cdot \mu m$

J Joules.

j Radiant flux from a surface, $W \cdot m^{-2}$

JJA June, July, and August, austral winter.

K Degrees Kelvin.

k	Boltzmann constant, 1.38×10^{-23} J/K
LP DAAC	Land Processes Distributed Active Archive Center.
LST	Land Surface Temperature, the temperature of the topmost layer of a land surface.
L↓	Downwelling longwave radiation.
L↑	Upwelling longwave radiation.
MAM	March, April, and May, austral fall.
MERRA	Modern-Era Retrospective analysis for Research and Applications, a climatology reanalysis dataset produced by NASA.
MODIS	Moderate Resolution Imaging Spectroradiometer, the 36 channel imaging instrument aboard NASA's Aqua and Terra satellites.
NASA	National Aeronautics and Space Administration.
NetCDF	Network Common Data Form, a self-describing file format.
NIR	Near-IR, the infrared spectrum from approximately 750 nm to 3000 nm.
nm	Nanometers, 10^{-9} meters.
NOAA	National Oceanic and Atmospheric Administration.
NSF	National Science Foundation.
NSIDC	National Snow and Ice Data Center.
OLS	Ordinary Least Squares.
PGE	Product Generation Executive, NASA's processing for MODIS product generation.
QC	Quality Control or Quality Checking.
R_{net}	Net surface radiation, the sum of the incoming and outgoing longwave and shortwave radiation.
R_{ss}	Total solar shortwave radiation, herein defined as wavelengths from 100 nm to 3000 nm.
s	Seconds.
SAM	Southern Annular Mode, a climate index in the Southern Hemisphere.

SAT	Surface Air Temperature, the air temperature at approximately 2 m elevation. In equations, identified as T_{air}.
SDS	Scientific Data Set, a subset of information often contained in a CDR file.
SOI	Southern Oscillation Index, a climate index.
SON	September, October, and November, austral spring.
SQL	Structured Query Language.
sr	Steradians, a measure of solid angle.
S↓	Downwelling shortwave radiation.
S↑	Upwelling shortwave radiation.
T	Temperature.
t	Time.
Tair	Near-surface air temperature, the temperature of the air approximately 2 m above the surface. See: SAT.
Tb	Brightness temperature, the sensible temperature T of a theoretical black-body, determined from measuring its thermal radiance at any wavelength.
TIR	Thermal Infrared, the spectrum of IR with wavelengths longer than approximately 8 µm (8000 nm).
$t_{LocalSolar}$	Local Solar Time, the time of day based on the Sun's zenith angle.
Ts	Surface temperature (see LST), typically in degrees Kelvin.
T_{UTC}	Coordinated Universal Time.
UTC	Universal Coordinated Time.
α	Albedo, or reflectivity of a surface, as a percentage or on a scale from 0 to 1.0.
ε	Emissivity, assumed to be broadband.
ε_λ	Spectral emissivity, the emissivity at a particular wavelength.
λ	Wavelength, in meters.
θ	Viewing angle, measured from 0° (zenith) to 90° (horizontal).

μm	Micrometers, 10^{-6} meters.
σ	Stefan-Boltzmann constant, 5.67×10^{-8} $W\cdot m^{-2}\cdot K^{-4}$
Φ	Radiant flux emitted from a surface, Watts or J/s.