

Methods for Background Temperature Estimation in the Context of Active Fire Detection

Bryan Allen Hally

**METHODS FOR BACKGROUND TEMPERATURE
ESTIMATION IN THE CONTEXT OF ACTIVE FIRE
DETECTION**

DISSERTATION

to obtain a double-badged Degree, namely:
the degree of doctor of philosophy at the University of Twente
and the Royal Melbourne Institute of Technology,
on the authority of the rector magnificus of the University of
Twente,
prof. dr. T. T. M. Palstra,
and the Council of the Royal Melbourne Institute of Technology,
on account of the decision of the Doctorate Board,
to be publicly defended at the University of Twente
on Thursday, September 19, 2019 at 14:45uur

by

Bryan Allen Hally
born on April 15, 1982
in Mornington, Vic, Australia

This dissertation is approved by:

Prof. dr. A. K. Skidmore (supervisor)

Prof. dr. S. D. Jones (supervisor)

A/Prof. dr. K. J. Reinke (co-supervisor)

Dr. L. O. Wallace (co-supervisor)

ITC dissertation number 363

ITC, P.O. Box 217, 7500 AE Enschede, The Netherlands

ISBN: 978-90-365-4834-2

DOI: 10.3990/1.978903648342

Printed by: ITC Printing Department, Enschede, The Netherlands

© Bryan Allen Hally, Enschede, The Netherlands

© Cover design by Adam Mattinson

All rights reserved. No part of this publication may be reproduced without the prior written permission of the author.



Graduation committee**Chair**

Prof. dr. A. Veldkamp University of Twente

Supervisors

Prof. dr. A. K. Skidmore University of Twente

Prof. dr. S. D. Jones RMIT University

Co-supervisors

A/Prof. dr. K. J. Reinke RMIT University

Dr. L. O. Wallace RMIT University

Members

Prof. dr. Z. Su Univeristy of Twente

Prof. dr. ing. W. Verhoef University of Twente

Prof. dr. K. Tansey University of Leicester

Prof. dr. P. Pilesjo Lund University

'Don't you want to know how we keep starting fires?'

Dick Valentine, *Danger! High Voltage*

Acknowledgements

Firstly I would like to give thanks to the patience and support provided by my supervisory panel, Prof. Simon Jones, A Prof. Karin Reinke and Dr. Luke Wallace from RMIT, and Prof. Andrew Skidmore from the University of Twente. Simon, thank you once again for firstly convincing me to embark upon this journey, and thank you for the numerous times you talked me off the ledge of giving up during the first year. Your ability to frame the project amongst the bigger picture and your conviction in my ability to get the work done have paid off, and I can only hope to be as successfully irreverent as you are most days. Karin, thanks for pulling me back to reality time and again. As much as I might make fun of some of our discussions, it's truly valuable to have someone asking questions and challenging assumptions, and also someone to sympathise with about deadlines. Luke, thanks for being my mental punching bag, provider of ideas and reassurances that I know all of this already, for listening and being happy to humour me up to the point where things no longer work. As much as it pains me at times, your confidence in your skills and the ability to back up that talk is something to be admired. Andrew, thank you for facilitating the dual-badged program and for making me feel welcome at ITC. Despite their infrequency, I really enjoyed our meetings and your insight into making foreign situations and opportunities work for you. I would also like to thank Lucas Holden for dragging me in for the initial chat that set me on this path, and also for his ongoing support. I would also like to thank Peter Hudson at Geoscience Australia for assistance with the Sentinel fire project data, Beth Ebert and Leon Majewski from the Australian Bureau of Meteorology for assistance with Himawari data, and Lyndsey Wright, Michael Rumsewicz and the rest of the staff at the Bushfire and Natural Hazards Co-operative Research Centre (BNHCRC). I would also like to thank Esther Hondebrink for the translation of the Summary chapter of this thesis into Dutch.

I would like to thank the BNHCRC, the NSW Rural Fire Service, the Victorian Department of Environment, Land, Water and Planning, and other partners for providing the funding that made the *Fire Surveillance and Hazard Mapping*

project possible, the RMIT School of Graduate Research for awarding me an Australian Postgraduate Award, and the BNHCRC for awarding me a top-up scholarship. I would also like to acknowledge the support of the Japanese Aerospace Exploration Agency, the Australian Bureau of Meteorology and the National Computing Infrastructure for providing access to Himawari data, and NASA and the NOAA for providing access to other remote sensing data used in this project.

I would also like to thank my fellow PhD students, post docs, and other research staff. From RMIT I would like to thank Chris Bellman, Mariela Soto-Berelov, Lola Suarez, Sam Hillman, Daisy San Martin Saldias, Chathura Wickramasinghe, Phil Wilkes, Margi McFadyen and Sam Hislop. From the ITC Department of Natural Resources I would like to thank Anahita Khosravipour, Jing Liu, Xi Zhu, Elnaz Neinavaz, Zhihui Wang, Trini Del Rio, Maria Buitrago-Acevedo, Yifang Shi and Festus Ihwagi, and everyone else who made me feel welcome at ITC, and I would also like to thank Linda van der Hout, Alexander Dijkshoorn, Nico Hendrickx, Yannick Donners, Roelofjan Velthuys, Yanna Kraakman, Henk Meijer, Nienke Nooren and Meike Nauta for welcoming me to the Bovenmaat and putting up with me during my five month stay in their house.

I would also like to thank my wider network of friends for their assistance and support during the last four years - they have been tumultuous times indeed. Thank you to my former housemates Adam Mattinson, Nicky Harris, Alistair Hunt, James Ahern, Morgan Tipper and Kylie Butler - your support throughout the Crumpet Age will be a cherished time of my life, and whilst they may rebuild the house, our memories from the Ritz shall hopefully never need such repair. I need to also thank Rob Fuller and Joel Graham for their post-fire assistance, and Tineke Fitzgerald and Stephen Impey for their ongoing support in the recovery.

Lastly, I would like to thank my parents, Jan and Rodney Hally. Whilst your unwavering trust in me to make the right decisions in life can make getting advice hard, I've never struggled to find a welcoming ear and full support on the occasions I do ask for it. Dad, even though you missed out on seeing me finish, I'm sure you'd be proud of who I am today, even though you were never any good at showing it. Mum, thanks for everything you do, you're a true champion and a super strong person, and despite the annoyances at times you continue to be in my corner for anything I take on in life.

Contents

Acknowledgements	i
Contents	iii
List of Figures	iv
List of Tables	viii
1 Introduction	1
1.1 General Introduction	2
1.2 Problem Statement	5
1.3 Research Questions	6
1.4 Thesis Structure	7
2 Estimating Fire Background Temperature at a Geostationary Scale	
— An Evaluation of Contextual Methods for AHI-8	9
2.1 Introduction	10
2.2 Method	13
2.3 Results	19
2.4 Discussion	32
2.5 Conclusions	35
2.6 Thesis Context	35
3 A Broad-Area Method for the Diurnal Characterisation of Upwelling	
Medium Wave Infrared Radiation	39
3.1 Introduction	40
3.2 Materials and Methods	43
3.3 Results	52
3.4 Discussion	57
3.5 Conclusions	60
3.6 Thesis Context	61
	iii

4	Advances in Active Fire Detection Using a Multi-Temporal Method for Next-Generation Geostationary Satellite Data	63
4.1	Introduction	64
4.2	Data	67
4.3	Method	69
4.4	Results	73
4.5	Discussion	77
4.6	Conclusions	83
4.7	Thesis Context	84
5	A New Spatio-Temporal Selection Algorithm for Estimating Up-welling Medium-Wave Radiation	85
5.1	Introduction	86
5.2	Spatio-Temporal Selection	87
5.3	Results	93
5.4	Discussion	100
5.5	Conclusion	103
5.6	Thesis Context	104
6	Synthesis	105
6.1	Summary of Results	106
6.2	The Background Behind Background	110
6.3	A New Direction for Real-Time Anomaly Detection	112
	References	116
	Summary	126
	Samenvatting	128

List of Figures

2.1	Examples of contextual temperature determination scenarios — (a) uniform contextual surroundings, with low spatial variance; (b) land cover change (yellow/green), with pixels of multiple land cover classes contributing to the estimate; (c) waterbodies (dark blue), which permanently obscure part of the contextual kernel; (d) cloud obscuration (hatched blue), which intermittently cause missing contextual data; and (e) smoke (grey), which provides directional partial obscuration of downwind pixels, and is less likely to be masked out of images than cloud.	11
2.2	(a) land area of the full disk covered by the AHI sensor; (b) 500×500 image tiles with sufficient land surface processed for the full disk analysis. The horizontal banding of the full disk image in (b) also corresponds to the areas selected for the cloud analysis presented in Table 2.2.	14
2.3	Case study areas selected for examination.	18
2.4	(a) Mean brightness temperature difference between contextual estimates and the central pixel for the ring of pixels at the edge of each window across the full disk for 0500 UTC B07 AHI-8 images. (b) Standard deviation of contextual estimates derived from each window edge by percentage of available pixels in the window edge.	22
2.5	Breakdown of temperature estimation pass rate on pixels that have no solution in their 5×5 window. The percentage of pixels covered by each bar this figure as a portion of all pixels examined is shown at the top of the figure. Each bar in the figure represents a minimum percentage level of valid contextual pixels for temperature calculation, and each coloured section represents the portion of pixels that are successful in deriving an estimate at each window size. The balance of exhausted pixels with no solution at each assessed percentage is also shown.	25

2.6	Mean difference between contextual estimates and the central pixel for the selected period for each area. (a) south-eastern Australia (sea); (b) north-western Australia (nwa); (c) Borneo (bor); and (d) central Thailand (thl).	27
2.7	Mean difference between contextual estimates and the central pixel for the selected period for each area. (a) eastern China (chn); (b) central Honshu (jpn); and (c) Siberia (sib).	28
2.8	Changes in spatial and statistical distribution of temperature estimates for the south-eastern Australia (sea) study area by window size. Window levels shown are (a) 5×5 window; (b) 7×7 window; (c) 9×9 window; and (d) 11×11 window.	30
2.9	Changes in spatial and statistical distribution of temperature estimates for the north-western Australia (nwa) study area by window size. Window levels shown are (a) 5×5 window; (b) 7×7 window; (c) 9×9 window; and (d) 11×11 window.	31
3.1	Time series diagram for a swath of latitude 25.75°S – 26°S from 135°E – 150°E longitude on 2015 day 319 (15 November 2015). Each square represents the median temperature of the 0.25° block at the image time on the y-axis. These blocks represent one minute of training data that can be fed into the brightness temperature aggregation process.	45
3.2	Figure 3.1 visualised relative to local solar time. Each of the swaths of block values extracted from each image are shown as a grey line. The four coloured lines depict the trajectory of individual block temperatures at 135°E , 140°E , 145°E and 150°E as the day passes.	46
3.3	An example of the training data fitting process on a swath for a 24-h period. Grey lines represent the raw median values for each swath from the time of each image, the blue data points are representative of the median brightness temperature of the training data at each local solar time, and the red trend line represents the filtered medians of the training data. (Due to the nature of the filtering process, the lack of data at each end of the data results in anomalous fitting and as such the swath sampling has been extended one hour either side of the 24-h period to minimise these errors.)	47
3.4	Examples of model fitting using the four training data derivations. Figure (a) shows a typical day with less than ten cloud instances, (b) shows a day with between 20–30 instances of cloud, and (c) is typical of a day with more than 70 identified cloud periods.	53

- 3.5 Availability of training data from the block and pixel based methods. (a,b) show the mean instances of training data available using the BAT method for October and November respectively; (c,d) show the training data available using the pixel method for the same months.(e,f) demonstrate the number of 24-h periods that could be utilised as training data for each block in October and November, and (g,h) show this same criteria using the pixel based method. . . . 56
- 4.1 Diagram showing the location of the selected study area in north Western Australia, along with the areas affected by fire between 1st - 28th of August 2016 as determined by the MODIS burned area product [53]. 69
- 4.2 An example illustrating the shifted temporal window described in Section 4.3.5 in operation. In this instance, a significant portion of the original fitting window, which is the time period to the right of the grey vertical line, is affected by active fire in the pixel, and this results in a higher fitted background temperature than expected during the period between 2016-218 09:00 and 2016-218 18:00, as shown by the line in dark green. By moving the start of the fitting earlier, based upon detection by the VIIRS active fire product in this case, the fitting process is less affected by the elevated temperatures associated with the ongoing fire. This produces a fitting such as that shown by the light green line, leading to fire-related anomalies being identified earlier, dependent upon the threshold set. 72
- 4.3 Associated fittings applied to a pixel at 15.5409°S, 129.2377°E, with a MODIS burned area product detection at 2016-219 05:20 UTC (shown by the red vertical line). The algorithm detection threshold set is 4 K. This figure shows ongoing fire activity in the AHI Band 7 brightness temperatures, shown here in blue, surrounding a single VIIRS active fire detection at 2016-218 16:40 UTC (in orange). With the temporal window based upon the time of the burned area product minus 23 hours, the first AHI detection at the 4 K threshold occurs 90 minutes after the VIIRS active fire detection (initial window detections are black circles). However, the shifted temporal frame based upon the time of this VIIRS detection produces a lower fit for background temperature during this night-time period, and the initial fire detection from AHI moves to 190 minutes before the VIIRS overpass. 77

4.4	Examples of temperature fitting and sources of perturbation. a) shows an example of anomalies in pixel brightness temperature caused by fire activity, b) is an example of negative temperature anomalies causing false detections, in this case cloud cover, and c) shows an example of false detection caused by improper fitting of the diurnal model. .	79
5.1	Locations of the case study areas selected for analysis in this paper, depicted on the AHI full disk.	88
5.2	Flowchart of the STS selection and estimation process.	91
5.3	Pixel training comparisons for selected pixels in the bor_l group. (Left) shows the spatial distribution of points selected during the training process relative to the 50-pixel radius selection area; and (Right) depicts the pixel trajectories over the image set examined for training, with the prediction target pixel value shown in red, the STS training pixel values shown in blue, and the surrounding context pixel values in green. Shown at time t is the distribution of values in the prediction image from both prediction methods, with their respective means shown as coloured crosses in comparison to the recorded brightness temperature shown as a red dot. Pixels are shown as labelled with prediction time t at 2016-067 04:20 UTC. . . .	94
5.4	A series of brightness temperature images and related estimations for the thl_j region. From left to right, the AHI B07 brightness temperature at the prediction time, the STS prediction image of the area, the contextual estimation of the area, and the differences between the AHI image and STS estimates, and the AHI image and context, are shown. The differences shown highlight positions where the recorded image value is higher than the estimation (red) and vice versa (blue). Prediction times are shown next to each figure.	98
5.5	Examples of common error in contextual brightness temperature estimation and the results using STS in similar conditions.	101

List of Tables

2.1	Specifications for the timeframes, area of the AHI disk and UTC times for analysis of each of the case study areas.	18
2.2	Average and standard deviation of cloud coverage for the AHI land areas covered in the study. The figures are an aggregate of 36 images recorded at 0500 UTC as mentioned in Section 2.2.1, broken into horizontal slices of the AHI disk as shown in Fig 2.2.	20
2.3	Number and percentage of pixels that are lacking sufficient adjacent pixels to provide contextual estimation at various window sizes and percentages across the AHI disk. A total of 4,663,165 AHI land pixels were evaluated.	20
2.4	Mean and standard deviation of the contextual estimate differences from central brightness temperature (AHI Band 7) for all available pixels in the 36 day set of full disk images at 0500 UTC. A total of 76,023,810 pixels were examined over the 36 images used in the study.	21
2.5	Mean and standard deviation of brightness temperature differences between the central pixels and the contextual surrounds at each window level per percentage level. Numbers shown in the 5×5 window row report statistics for pixels that would be added to the 1.00 pixels if the valid context percentage shown was used to accept contextual estimates. The percentage of total pixels with estimates available at the 5×5 window for each valid context percentage is also shown. The rows for each subsequent window size describe the number of temperature estimations that would be added from failures at the previous window size by expanding the examined window, and the subsequent means and variances of pixels included from these window sizes. A total of 76,023,810 pixels were examined over the 36 images used in the study.	24

2.6	Mean and standard deviation of brightness temperature differences between the central pixels and the contextual surrounds at each window level per percentage level, or where number of context pixels reaches 10. The 5×5 window statistics show the global rates for pixels which have equal or greater contextual pixels than the minimum for estimation. The rows for each window size describe the number of calculated values that would be added by expanding to each window size, and the subsequent means and variances of pixels included from these window sizes.	36
2.7	Mean and standard deviation of mean brightness temperature differences of each case study area for each 31 day period. Pixel values were averaged over the 31 day period for each site, and global means and standard deviations of these averages are reported.	37
2.8	Mean and standard deviation of brightness temperature differences between the central pixels and the contextual surrounds at the specified percentage levels for the 5×5 window in each case study area. Each column reports the statistics of accepting the available pixels above the denoted percentage level. Pixels with full contextual coverage are reported in the 1.00 column.	37
3.1	Comparison of MTSAT-2 and AHI-8 sensors for fire detection using a MWIR ($\sim 4\mu\text{m}$) channel.	43
3.2	Comparison of fitting techniques to brightness temperatures recorded by the AHI sensor using root mean square error after eliminating incidences of Clear Sky Probability (CSP) of less than one from the evaluation.	52
3.3	Comparison of time taken to provide a temperature fitting using the pixel-based training technique and the BAT fitting process.	55
4.1	Raw anomaly rates for the threshold algorithm for the selected temperature thresholds, and the proportion of these anomalies which have an associated disturbance detected by the MODIS burned area product, from a total selection of 93,906 cloud-free diurnal fittings. .	74

4.2	Detection results of the thresholding algorithm on 150 fire incidents in each detection grouping per temperature threshold. Detections occur where at least one brightness temperature measurement exceeds the fitted brightness temperature by the selected threshold. Synchronous fire detections are classified as where an anomaly detected by one or both of the active fire products has at least one corresponding detection from the threshold algorithm within twenty minutes of the LEO detection.	75
4.3	Time of detection of fires using the threshold algorithm in comparison to times of first detection using the two LEO active fire products at each temperature threshold. Times shown are the average time of detection prior to LEO active fire detection, with numbers shown for both the diurnal temporal window commencing 23 h prior to burned area detection, and for the shifted temporal window commencing 22 h before initial active fire detection. Times shown are in hours and minutes.	76
5.1	Specifications for the timeframes, area of the AHI disk and UTC time offsets for each of the case study areas examined.	90
5.2	Accuracy of estimation techniques against brightness temperature values from the assessed images by case study area. $\Delta\sigma$ is the percentage change between the standard deviations of the context and STS estimation methods.	95
5.3	Breakdown of the availability of temperature values using the two estimation methods against total image pixels present. n BT Obs gives the number of cloud free image pixels out of the total possible shown in Total Pixel Obs.	99

Introduction

1

1.1 General Introduction

Fire is one of the most important ecological drivers on the planet. Whilst its initial impacts upon the environment are seemingly destructive, fire also plays an important role in the functioning and regeneration of ecosystems, along with providing a catalyst for changes to ecosystem successional processes. Approximately 348 Mha of the earth's surface is affected by fire every year [24], with fire activity predicted to increase due to the effects of a changing climate [10, 47, 14]. Fires can occur in almost all of the varied land cover types throughout the world, with the exception of desert and arctic areas, with the spread and intensity of fire driven by fuel volume, fuel moisture and weather conditions. Whilst lightning is generally the cause of fires in more remote parts of the landscape, the majority of fires are started by humans. In the United States for instance, it is estimated that 84% of all wildfires are started by humans, with strong correlation between fire location and human activity [1, 44]. The economic impact of wildfire is estimated to be between USD 70 billion and USD 348 billion annually in the United States alone [81], and in Australia the impact of individual fire events on residences, agriculture, forestry and the environment can reach beyond AUD 1 billion alone [79]. Fire also contributes a significant amount to global carbon emissions into the atmosphere, with [87] estimating total wildfire emissions worldwide at 3.53 Pg annually, which equates to almost 40% of of total carbon emissions [70].

With these impacts in mind, finding methods to efficiently measure and monitor wildfire is of vital importance. Information about fire in the environment can assist in mitigation planning and asset management for fire authorities, assists in the minimisation of fire impacts on human lives and communities, and of course can help to address problems such as carbon accounting and ecosystem impacts of fire.

Early detection of potential wildfires is also vital to assist with the strategic and effective management of wildfire situations. Obtaining this information in an accurate and timely manner *in-situ* can be a difficult task. Historically, fire authorities have placed equipment and manpower in elevated areas for fire spotting purposes, with fire reporting to this day remaining reliant upon eyewitness accounts from people close to the fire activity. In populous areas this type of fire detection can be effective at generating location information about a fire, and may generally estimate the time of ignition well, but lacks the ability to generate quantitative information about fire size, intensity, growth rates and general behaviour. However, fires may also go undetected in remote, underpopulated areas, or where resources do not exist to capture all fire activity comprehensively.

Remote sensing has become a vital source of information about wildfire activity where traditional coverage is lacking. It also has become a vital source of information with regard to fire intensity, burned area, and for estimates of carbon emission. Satellite remote sensing of fire is possible due to the radiative emission of wildfire and the transmittance of important wavelengths of electromagnetic radiation through the earth's atmosphere. Depending upon the type of fuel and weather conditions, wildfire generally burns at between 700–1000 K, which provides a peak in blackbody radiative output at around 4 μm in the electromagnetic spectrum. This peak in radiative energy is several orders of magnitude above the radiative emission of land at a typical background temperature of 300 K. This wavelength in the electromagnetic spectrum happens to coincide with a "window" or peak in the atmospheric transmittance of the earth's atmosphere, which means that radiation emitted at this wavelength travels relatively unimpeded through the atmosphere into space. These two factors contribute to the situation where the radiation at the peak fire wavelength is not only easy to detect, but that fire activity stands out in high contrast to the background temperature. This means that fires do not have to be large relative to the ground sampling size of the sensor in order to be detectable as anomalous energy sources. This lucky break has led to the proliferation of methods for fire detection from various satellite remote sensors.

The possibility of fire detection and monitoring was first explored by [11], whose work used a bispectral method to attempt to identify hot sub-pixel targets within a uniform background. This was the first work utilising remote sensing to attempt to identify characteristics of anomalous temperatures, including the portion of pixel affected by the anomalous activity, and the temperature of the anomaly source. This method was utilised for a number of different fire detection algorithms for a number of different sensors [48, 63, 57, 42], and is still the main driver of the WF-ABBA fire algorithm used upon GOES satellites over the Americas [45]. A number of flaws were identified in the use of this bi-spectral relationship, and these are highlighted in [22], which included the sensitivity of the method to error when dealing with small fires. In the mid-1990s, a new method of anomaly detection was proposed by [40], which involved use of the contextual area surrounding the potential anomalous pixel as the source of the estimate of the temperature of the target pixel. This method quickly caught on as the preferred method of determining fire background temperature, and became the basis for a number of fire detection algorithms [15, 23, 43, 9, 19, 61, 76], and is the main driver of fire background temperature for the commonly used MODIS [26] and VIIRS [75] active fire detection products.

The detection of fire from satellite remote sensor systems involves trade-

offs between spatial, radiometric and spectral resolutions, and varying accuracy - all dependent upon the type of sensor in use for the task. Early remote sensing sensor systems (e.g. AVHRR [16]) lack the spectral band-pass widths found on more modern sensors (e.g. VIIRS [92]), and often have poor radiometric resolution and low saturation temperatures. This hampers their usefulness for accurate description of fire activity. The bulk of effort from the scientific community in this field in the 1990s and 2000s focused upon the sensors in Low Earth Orbit (LEO), such as AVHRR, MODIS, VIIRS, BIRD, Firebird, and others. Sensors in these polar orbits are set to pass over the observation areas at specific times of the day (generally late morning or early afternoon) to maximise the effectiveness of the visible bands for image capture. With the lower orbits of LEO sensors, generally between 450-900 km, the spatial resolution available is very high, enabling small areas of fire activity to be identified in detail. Despite this, coverage of the temporal activity of fire is limited due to their sun-synchronicity, with typically between 3 - 5 images available daily from a constellation of LEO sensors such as MODIS. In the mid 2000s, advances in the sensors being placed in geostationary orbits led to increased activity in the detection of fire from sensors at this orbit. Whilst geostationary (GEO) satellites orbit the earth at 35.786 km, which leads to diminished spatial resolution due to the distance to the earth, these GEO sensors are fixed in their viewpoint of the earth, and provide continuous coverage dependent upon the temporal resolution of the sensor. This fixed view provides the ability to monitor change over time, and provides the ability to approach fire detection using a different framework.

Whilst many efforts such as the WF-ABBA [45], GOES [95] and MSG-SEVIRI related fire detection algorithms [28] focused upon extension of the single image contextual algorithms into the GEO sensor space; innovative ways of using the temporal stream of data supplied by GEO sensors started to appear in the mid 2000s. The work of [27] first examined the modelling of the diurnal temperature cycle for land surface temperature estimation, by way of a prescriptive model using data from the METEOSAT sensor. Use of the diurnal cycle for potential fire anomaly isolation was proposed in [84], which used a Kalman filter for temperature modelling purposes. A number of different multi-temporal techniques for determining fire activity have spawned from this work [86, 83, 85, 65, 13] and research continues to focus upon refinement of background temperature modelling in the diurnal temporal space.

This decade has seen opportunities grow in the GEO space for more accurate fire detection, with new sensors providing improvements in their temporal, spatial and radiometric resolutions. The previous generation of GEO satellites launched in the mid-late 2000s, which includes MSG-SEVIRI, MT-SAT2, and

GOES 13-15, provided between 15–30 min temporal coverage of their respective full disk areas, with a typical spatial resolution of 4 km in the thermal bands used for fire detection. These sensors (with the exception of SEVIRI) are now making way for a new generation of GEO satellites launched in the mid-late 2010s. The launch of AHI-8 [60] in 2014 by JAXA over the Asia-Pacific, followed by the launches of GOES-16 and 17 in 2016 and 2018 [72] has led to significant improvements in sensing capability. These new sensors are capable of capturing full disk images with 10 min recapture time, and the spatial resolution of all infrared bands on these sensors is 2 km. These sensors are also capable of short term mesoscale captures, with areas of up to 1000×1000 km capable of being captured at one minute intervals. This gives an unprecedented view into short-term changes on the earth's surface, at a spatial resolution far more suited to the isolation of fire activity. This drive from new sensors and data streams also heightens the necessity for more research in this area, to most effectively utilise the new information about the earth's systems coming from these sources.

1.2 Problem Statement

Early and reliable active fire detection is of great importance to land managers, to assist in risk assessment, mitigation strategies and minimisation of harm to both people and assets [80]. With the timeliness of fire detection in mind, and with shortcomings in current LEO fire detection products due to temporal coverage [61], a need for strategies and systems that apply predominately to geostationary sensor imagery has been identified. Whilst geostationary sensors, especially those from the current generation of new satellites (AHI-8, GOES-R), exhibit high temporal refresh rates, their spatial resolution is far more coarse than the sensors that the current standard fire detection products are based upon. The effect of systematic errors that occur in the use of brightness temperature estimation from situational context is acknowledged but poorly understood, and the effect of coarser spatial resolution may exacerbate these errors further.

Whilst the common remotely sensed fire products are generally based upon techniques that are applicable to discretely captured events, due to being based on information taken from sensors that move relative to the earth's surface, geostationary images allow for the continuous capture of information from a fixed location in orbit over time. Leverage of the temporal domain for fire detection has been successful from geostationary images using various methods [56, 65, 84, 66], but these have focused upon deriving estimates based upon

the data from individual pixels, sacrificing the use of relevant spatial-based correlations. An opportunity lies in the development of techniques that utilise both spatial and temporal relationships to drive more robust and accurate estimation than those driven solely from either of these domains.

Given the problems intrinsic to fire detection, and more broadly anomaly detection, techniques for estimation must be robust in the face of anomalies in any training process attached to them. Clouds, smoke and fires have large roles to play in the necessity for secondary false alarm processes in current fire detection methods. The focus of any estimation technique developed should be in mitigating the influence of these factors on the training, modelling and subsequent estimation process, with a view to minimisation of these secondary tests that may introduce or exacerbate error in resultant anomaly attribution.

1.3 Research Questions

In order to address the gaps identified in section 1.2, four research questions are identified and outlined below:

Question 1. What is the effect of systematic and structural errors caused by the use of spatial contextual estimation in common fire detection techniques?

Whilst contextual estimation is an accepted form of temperature estimation for fire detection, there have been no studies previously that demonstrate the expected errors in such estimates, or the specific landforms and conditions that may propagate these errors. This question aims to provide a comprehensive breakdown of the expected error in such estimations based upon application to imagery from the AHI-8 geostationary sensor, and examine the causes over a number of case study areas.

Question 2. How can we use the common diurnal variation of upwelling radiation to estimate brightness temperature in a robust fashion?

Areas at similar latitudes will receive similar solar radiation budgets — use of this assumption can allow us to create models of the expected temperature based upon the standardised form of the diurnal cycle for a specific latitude. This question aims to derive a method of temperature estimation that takes advantage of the high temporal frequency of the imagery from AHI-8 to create a time-corrected idealised model of diurnal variation in brightness temperature. This will exploit non-cloud affected areas and apply this information to more obscured regions. This technique is called the Broad Area Training method (or BAT).

Question 3. How effective is the new Broad Area Training method at identifying fire-related brightness temperature anomalies in comparison to other

fire detection methods?

This question addresses the ability of the brightness temperature estimation method outlined in the response to Question 2 to identify anomalous pixels in a near-real-time fashion. The method is tested over a range of anomaly temperature thresholds, and comparisons are made to fire products derived from low earth orbit imagery to compare timeliness and potential for omission error.

Question 4. How can we use similarities in image characteristics over time to improve temperature estimation over a single-image contextual approach?

Given the weakness of contextual estimation results as a function of available adjacent land pixels, an opportunity lies in leveraging similarities in image values over both time and a wider area to provide an improved set of candidate pixels for temperature estimation. This question outlines a new method of spatio-temporal sampling of candidate training pixels for use as brightness temperature estimators in subsequent imagery. The work identifies criteria for training pixel selection, and compares estimates of brightness temperature back to those derived from contextual estimation.

1.4 Thesis Structure

Presented in this thesis are four research chapters that address the questions asked in 1.3. Chapter 2 tests the currently accepted method of background temperature estimation over imagery from a geostationary sensor to determine the errors associated with this method's use, along with isolation of the potential causes of such errors. Chapter 3 of the thesis introduces the Broad Area Training method of brightness temperature estimation, which provides temperature fitting of candidate pixels based upon the idealised diurnal cycle of pixels at similar latitudes. Chapter 4 examines an application of the method outlined in Chapter 3 for the purpose of isolating potential temperature anomalies, plus provides a comparison to commonly used polar earth orbiting sensor based fire products to determine rates and times of detection. Chapter 5 introduces the use of the spatio-temporal selection method for background temperature estimation, and provides example images and comparisons to contextual estimates for a number of case study areas. Finally, the thesis concludes with a synthesis chapter which collates the research presented in the thesis and discusses the potential of the methods presented not only for fire detection purposes but for other types of environmental monitoring.

Estimating Fire Background Temperature at a Geostationary Scale — An Evaluation of Contextual Methods for AHI-8

2

2.1 Introduction

Satellite remote sensing has become a vital tool in the arsenal of land managers, not only for the initial detection of active fire, but as part of inputs for modelling and planning purposes. Timely and accurate fire information from remote sensing enables preparation and planning for mitigation activities, along with providing vital information about fire behaviour and characteristics [78]. Increasing importance is being placed upon active fire products to calculate metrics such as fire radiative power and burn severity [68], in order to obtain an understanding of how the environment burns, and also to provide input for environmental modelling and quantifying outputs such as carbon emissions from fire.

Active fire detection from remote sensing relies on elevated levels of radiation in the infra-red wavelengths caused by the blackbody radiation emitted from fire [68]. The typical energy emitted by fire at medium wave infra-red (3 – 4 μm) wavelengths can be several orders of magnitude higher than regular radiation levels, which are primarily made up of thermal emission from the surface and solar reflection [67, 29]. This disparity in energy levels allow fires that are much smaller than the pixel area to be detected, as the extra energy from a fire will overwhelm the background level of radiation [78]. This propensity of fire to overwhelm the background signal presents a problem for fire detection purposes as well. The ability to determine whether a pixel is fire-affected is dependent upon knowing what the pixel should look like in the absence of fire [23]. Accurate knowledge of the differential between fire signal and background allows fire to be detected, and enables the calculation of common fire-related metrics such as fire radiative power (FRP) [66].

Without the ability to directly measure the background temperature of a pixel in the event of fire, fire algorithms have largely utilised the land area surrounding a target pixel to facilitate estimation of the background temperature, a method known as contextual estimation [21, 19, 9, 73, 46, 66, 88]. For pixel brightness temperatures in the medium wave infrared, spatial autocorrelation is primarily driven by latitude, with adjacent pixels receiving similar amounts of solar radiation, along with climatic conditions, which homogenise land cover over localised regions. This was highlighted in [66], who stated the assumption of neighbouring pixels having the same surface background characteristics was implicit in the fire algorithm developed in that work. This work [66] also stated that *"... the extent to which this is true depends of surface spatial*

This chapter was published in a peer-reviewed journal as: Hally, B., Wallace, L., Reinke, K., Jones, S., Engel, C., & Skidmore, A. (2018). Estimating Fire Background Temperature at a Geostationary Scale — An Evaluation of Contextual Methods for AHI-8. *Remote Sensing*, 10(9). <https://doi.org/10.3390/rs10091368>

homogeneity and the sensor spatial resolution.” There has been no thorough examination of how surface homogeneity affects the accuracy of fire detection algorithms, despite this assumption being prevalent in active fire algorithms and products. Contextual measurements are also influenced by obscuration due to cloud or smoke, which may lead to decreased infra-red radiation in pixels adjacent to a target pixel [58]. Additionally, adjacency to water bodies may eliminate some pixels from being used in contextual calculations, with islands and coastal regions particularly susceptible to errors caused by reduced land surface availability. Examples of how these scenarios may influence the calculation of background temperature may be seen in Figure 2.1.



Figure 2.1: Examples of contextual temperature determination scenarios — (a) uniform contextual surroundings, with low spatial variance; (b) land cover change (yellow/green), with pixels of multiple land cover classes contributing to the estimate; (c) waterbodies (dark blue), which permanently obscure part of the contextual kernel; (d) cloud obscuration (hatched blue), which intermittently cause missing contextual data; and (e) smoke (grey), which provides directional partial obscuration of downwind pixels, and is less likely to be masked out of images than cloud.

Land surface temperature is a well covered topic in remote sensing [50, 17, 69, 93, 52, 96], but most techniques focus upon use of thermal infrared (8 – 12 μm), which lacks a solar reflection component. This has led to an integration of land surface temperature techniques encompassing a combination of medium-wave and thermal infrared bands for fire detection purposes [26, 90, 9, 66, 62], due to the differential response between these two wavelengths to emitted energy from fire. Such methods rely on accurate knowledge of the sensor response to temperature in both infrared bands and their relation to one another, and often rely on arbitrary statistical thresholds to relate the two bands for detection purposes, and studies such as [22] have highlighted issues with the use of bispectral methods of fire detection. Algorithms exclusively using medium-wave infrared for background temperature detection have generally used this approach for calculation of metrics such as FRP, which is less reliant on highly accurate temperature information to achieve satisfactory results [95, 65, 91].

The successful launch of the AHI-8 sensor in 2015 has expanded the availability of geostationary satellite image data for the Asia-Pacific, both in the spatial and temporal resolution domains [39]. The increased spatial resolution

of the sensor, which achieves $2\text{ km} \times 2\text{ km}$ resolution in the medium wave and thermal infrared bands, and the increased temporal coverage of the sensor, which records an as-yet unparalleled 10 min refresh rate for geostationary full disk images, provide opportunities to image and analyse the sensor's coverage area in far greater detail than previously [30]. The fire detection and examination capabilities of the sensor have already been demonstrated in multiple studies [59, 94, 88, 31]. These studies use a mix of contextual and multi-temporal techniques to detect and monitor fire activity, but as yet there has been no definitive fire algorithm for all conditions adopted for use with this sensor.

Fire detection algorithms perform a number of tests to not only isolate elevated sources of radiation, but to also eliminate false positive detections. Tests are usually made to mask cloud, which can trigger some detections through elevated reflectivity in the medium-wave infra-red, for masking excess solar reflectivity in the form of sun glint, and to flag areas of water, which will bias infra-red measurements downwards. Once these sources of error are eliminated from evaluation, decisions are then made about the suitability of pixels surrounding a potential fire for fire background temperature calculation. For instance, the MODIS MxD14 product [26] uses values initially from a 3×3 (3 km) pixel window surrounding the target pixel (without the leading and trailing pixels in the cross-swath direction due to pixel smearing) to determine this temperature. The algorithm then tests how many suitable contextual pixels are available for evaluation, with a successful set of target pixels isolated for temperature calculation when the number of valid contextual pixels reaches at least 25% of the total, with a minimum of eight contextual pixels used for calculation. If the algorithm cannot find sufficient pixels at the first window (in this case, only six pixels are available and eight are required), the window expands to 5×5 pixels, and the tests are repeated. If the test fails again, the cycle repeats expanding the window to the maximum size of 21×21 , at which point the tests conclude with no result.

This technique of the expanding window is not exclusive to use for MODIS. The VIIRS VNP14 product [75] has a background temperature calculation based upon a starting window of 11×11 ($\sim 4\text{ km}$ in length), a success rate based on 25% of valid contextual pixels available for calculation and a 10 pixel minimum, and a maximum window range of 31×31 ($\sim 10\text{ km}$ in length). The Fire Identification, Mapping and Monitoring Algorithm (FIMMA) for use on AVHRR sensors [49] started with a 5×5 window, ended at the 41×41 pixel level, and used 35% of total contextual pixels available with a minimum number of eight pixels used. Work involving fire detection using Landsat-8 [76] involved evaluation of a fixed 61×61 pixel window for background temperature calculation, with no limits placed upon the number of pixels used. Geostationary

satellite algorithms apply these contextual tests as well - the MSG-SEVIRI sensor fire algorithm [66] starts at a 5×5 window (15 km due to the sensor spatial resolution), with a maximum window size of 15×15 (45 km) evaluated before calculation failure. The pixels inside each window are tested against cloud, sun glint and anomalous differences between medium wave and thermal infra-red, and only if at least 65% valid context pixels are available will an estimation take place. This work on SEVIRI has also been extended for use on the GOES sensors [93], with similar parameters used for contextual pixel utilisation.

These expanding window methods for evaluating temperature from pixel context are applied to sensors with different spatial and radiometric characteristics, so they should differ slightly in application based upon each sensor. Despite this, apart from a rough relationship of spatial scaling between some of the products, there is no general consensus as to the ideal dimensions for contextual window evaluation, and indeed no optimal value for minimum percentage of valid contextual pixels to use for deriving an accurate background temperature.

The objectives of this work are to examine common methods of deriving land surface temperature from a target's surroundings in the context of fire detection. To achieve this, the enhanced temporal and spatial capabilities of the AHI-8 sensor are exploited in a large-area study. This paper presents the effects of variation of examined window sizes and valid contextual pixel percentages on background temperature. This work also highlights the challenges faced in using contextual estimation effectively, with in depth examinations of a number of case study areas to determine the effectiveness of contextual temperature calculation.

2.2 Method

2.2.1 Data

This study utilises images from the Advanced Himawari Imager-8 (AHI-8), a geostationary sensor located at 140.7° E longitude [60], data from which was obtained from the Japan Meteorological Agency (JMA) via the Australian Bureau of Meteorology (ABOM). This geostationary sensor provides coverage over the Asia-Pacific region over 16 bands, with an image captured every 10 min. Images were obtained from the 3.9 μm medium-wave infrared band (AHI-8 Band 7) data, which is available in Australia from the National Computing Infrastructure (NCI). Dates were randomly selected for 36 days of the year 2016, with a distribution of three per calendar month in order to provide a representative sample of times in the results. The Julian dates selected were days 6, 10, 20,

2. Estimating Fire Background Temperature at a Geostationary Scale — An Evaluation of Contextual Methods for AHI-8

35, 36, 41, 71, 72, 82, 97, 101, 103, 133, 144, 149, 153, 164, 173, 184, 188, 200, 222, 230, 236, 253, 257, 274, 279, 286, 290, 314, 322, 323, 343, 353 and 355 of 2016. A single image was examined at each of these days for the full disk examination, which was taken at 0500 UTC. This time was selected for full disk processing to maximise the amount of the land surface in daylight, along with examination of much of the disk at, or near, peak daily temperatures. This timing also coincides with the afternoon overpass of the VIIRS sensor for much of the land areas of the disk. This study utilises a cloud mask algorithm used in a study of AHI fire detection by [94], which was adapted from use on the GOES-11 and GOES-12 geostationary sensors from [95]. This mask is calculated using AHI Bands 3, 7 and 13, along with solar zenith information at each image time, from products supplied by ABOM.

To enable efficient processing of full disk images the size of those captured by AHI, each full disk image was divided into component arrays of 500 x 500 pixels in size. The number of land pixels in each of these component arrays was then counted, and arrays containing less than 100 land pixels were discarded from analysis. Along with these omitted areas, arrays comprising solely of land constituting the continent of Antarctica were also discarded. Once these tiles were identified, selections from each image with a 12 pixel buffer (for expanding window analysis purposes) were made of each tile and processing was performed. The areas with sufficient land for analysis are shown in Figure 2.2.

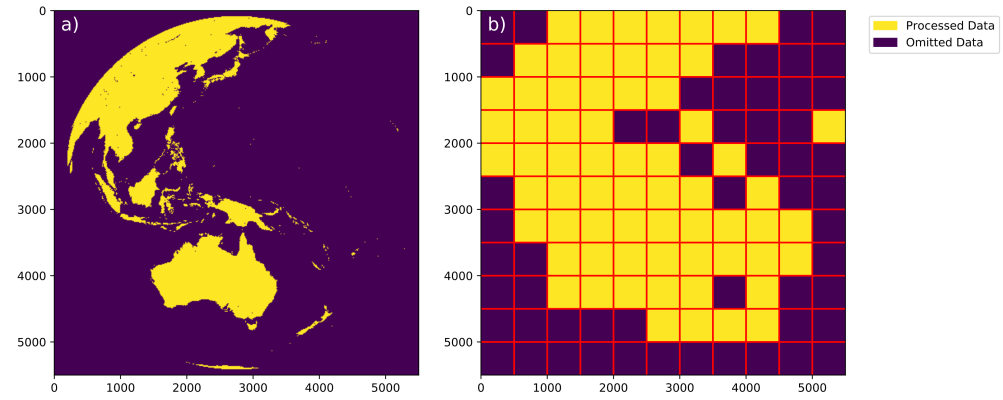


Figure 2.2: (a) land area of the full disk covered by the AHI sensor; (b) 500×500 image tiles with sufficient land surface processed for the full disk analysis. The horizontal banding of the full disk image in (b) also corresponds to the areas selected for the cloud analysis presented in Table 2.2.

As the focus of this study is determination of brightness temperature of land pixels, a land/sea mask supplied as part of the AHI ancillary data was applied to imagery to mask non-land pixels. Pixels close to the edge of the

full disk are stretched over a large area of land surface, and also suffer from refraction due to the longer transmission period through the atmosphere. Pixels that have a sensor zenith angle greater than 80° were masked from further analysis using the AHI sensor ancillary product provided by ABOM.

2.2.2 AHI Disk Characterisation

Cloud is a major source of occlusion when measuring brightness temperature values. In order to obtain an understanding of the role cloud cover plays in an AHI full disk image, and by extension the distribution of clear sky pixels for analysis, the AHI image was broken into sub-images of 500 rows, for the first 5000 rows of the 5500×5500 image. The number of land pixels available in each of these sub-images was tallied, and the cloud coverage from the cloud mask was recorded for each full disk image. This breakdown of the AHI full disk into sub-images can be seen in the horizontal banding depicted in Figure 2.2b.

The land area covered by AHI can be quite discontinuous, especially in the equatorial regions where many islands are present. These islands and coastal areas will have permanent gaps in their contextual coverage area due to the land forms surrounding them. In order to gain an understanding of the magnitude of these standing anomalies, an analysis of the land mask was conducted. Pixels were selected by the number of contextual pixels available for estimation during a cloud-free period, and categorised into percentage classes (75 %, 65 %, 55 %, 45 %, 35 %, 25 %, 15 %). Pixels that had less than the required percentage of pixels available on the land mask were flagged, and counts of these unusable pixels were tabled.

To investigate the effectiveness of contextual estimation at a full disk level, the mean of all available contextual pixels was taken for each window size for each cloud-free pixel in the 36 images selected for study. The difference between each of these contextual estimates and the benchmark central pixel was calculated, and mean and standard deviations of these differences were aggregated for analysis. These values were further broken down by the exact percentage of contextual pixels available at each window level, in order to understand how percentage of valid pixels affects the ultimate calculation of contextual temperature.

The size of the land area covered by individual pixels in a geostationary image increases as the sensor zenith angle increases. To determine whether this expansion of pixel area has an effect on contextual temperature calculations, all pixels from the dataset with contextual estimates were then divided into classes based upon their sensor zenith angle (eight classes spanning 10° from

0 - 80°), and statistics were aggregated for each of these classes.

2.2.3 *Expanding the Window*

As noted in the introduction, there have been many approaches taken to determine a suitable window size for contextual calculation, and no general consensus has been reached for ideal parameters, apart from a rough $10\text{ km} \times 10\text{ km}$ maximum window size for the LEO sensor algorithms. For a geostationary sensor like AHI, we are limited as to the spatial bounds of the minimum window size we can select, as the sensor resolution prevents us from resolving at better than two kilometres in the infra-red bands. A minimum sampling window of 5×5 has been set around each pixel, which corresponds to $10\text{ km} \times 10\text{ km}$ at sensor nadir. A number of window sizes were examined, with values selected in two pixel increments up to a maximum window size of 25×25 pixels. Each of these windows had a count of valid pixels, and the mean and standard deviation of differences between the contextual mean and the central pixel value recorded for each pixel for each image.

A common feature of contextual algorithms is the use of a threshold of valid pixels as a portion of total examination window as a limiting factor for estimation validity. If the target pixel has at least the number of valid context pixels set by this threshold, the target's contextual pixel values are used to calculate a temperature estimate, otherwise the target is ignored. There is no consensus upon which to base a definitive decision about valid context percentage choice - the most commonly used success criterion is 25 % or an arbitrary number of pixels, as used by both MODIS and VIIRS in their respective fire products. This study has chosen to examine the use of seven percentage thresholds of contextual pixel availability, ranging from 75 % to 15 % in 10 % increments. A pixel is deemed to have sufficient contextual data to make a calculation when the number of valid contextual pixels is equal to or greater than the selected percentage over the window being examined. For example, at the 5×5 window size, nine or more valid pixels need to be available for a temperature to be calculated at the 35 % threshold. At some thresholds, land pixels with proximity to oceans and lakes may have insufficient land available to calculate a temperature.

Another commonly utilised feature of contextual algorithms is the expanding window. When insufficient data is available at an inner window size, the window of examination grows outwards until it obtains sufficient data to make a temperature determination. For a true evaluation of the effects of the expanding window on contextual estimation, it is important to know not only how often this window expansion occurs, but the effect the expanding window has upon calculated contextual estimations. For the expanding window sec-

tion of this study, the portion of data with full contextual coverage at the 5×5 window was analysed separately from pixels with at least one contextual pixel obscured. From the remaining pixels for each of the valid context percentages, pixels with sufficient context available at the 5×5 were identified, and statistics calculated over these pixels. For the remaining pixels with no solution at the 5×5 window at each valid context percentage, the window of examination was expanded to 7×7 . At this point, the counts of valid context pixels were totalled for the current window and all previous windows. If the new number of contextual pixels was sufficient for the valid context percentage to be met, a contextual estimate was calculated over all contextual pixels available, and these statistics were recorded for reporting at the specified window size. After this, the examination window was expanded, and the process was repeated. Once the window of examination reached 25×25 , some pixels were unable to find a solution based upon the selected percentage of valid contextual pixels. Counts of these failed pixels were also recorded.

Also, some expanding window methods will in addition use an absolute threshold for the number of valid contextual pixels required for temperature estimation. Once the number of contextual pixels available satisfies this threshold of valid pixels, a contextual estimate will be made based upon the available pixels regardless of the valid context percentage set. The work presented in this paper also examined the effects of using an absolute threshold of valid pixels of 10, similar to the VIIRS VNP14 product. For this, the 5×5 window was firstly analysed, and as 10 pixels was the cutoff for validity for the 45% valid pixel class at 5×5 , no higher valid contextual pixel percentages were examined. If a target pixel had either the required percentage of contextual pixels available, or sufficient contextual pixels to reach the absolute cutoff, the target pixel had a context temperature estimate calculated and recorded. Where this requirement was not met, the window was expanded to the next window size. If a target pixel did not reach either the valid contextual percentage or the absolute threshold of contextual pixels by the 25×25 window, the target pixel was recorded as a failure and tallied.

2.2.4 Case Study Evaluation

A series of case study areas have also been evaluated in a more in-depth fashion, due to their land surface variation or their fire-prone nature. These areas include part of south-eastern Australia, part of north-western Australia, a section of Kalimantan's east coast, part of central Thailand, part of eastern China, the central part of Honshu in Japan, and part of Siberia east of Lake Baikal. Each of these areas consists of a section of the AHI image measuring 200×200

2. Estimating Fire Background Temperature at a Geostationary Scale — An Evaluation of Contextual Methods for AHI-8




Figure 2.3: Case study areas selected for examination.

pixels in size, with a small buffer to provide data for pixels at the edge of the selected window. These study areas are highlighted in Figure 2.3.

Table 2.1: Specifications for the timeframes, area of the AHI disk and UTC times for analysis of each of the case study areas.

Case study area	Start Date	End Date	AHI image area	Time (UTC)	Local time @ centroid
sea	2016-03-30	2016-04-29	[4400, 4600, 3050, 3250]	3:50	13:49
nwa	2016-10-23	2016-11-22	[3600, 3800, 2000, 2200]	5:00	13:32
bor	2016-02-14	2016-03-15	[2600, 2800, 1400, 1600]	5:40	13:22
thl	2016-02-28	2016-03-29	[1800, 2000, 800, 1000]	6:30	13:15
chn	2016-08-27	2016-09-26	[1000, 1200, 1600, 1800]	5:10	12:56
jpn	2016-05-03	2016-06-02	[900, 1100, 2500, 2700]	3:50	12:59
sib	2016-05-10	2016-06-09	[200, 400, 2000, 2200]	5:00	12:43

In order to provide a more representative understanding of how each of these landscapes behaves during fire-prone periods, a selection of images for each case study area was made based upon the prevalence of fire over 2016. The monthly VIIRS fire product (VNP14IMGML) [74] was subsampled for each of the study areas, and a rolling window of 30 days was applied to the sum

total of fires from each area over the course of the year. The point of time exhibiting maximum fire activity from this was then used as the central day in a 31 day window for in-depth analysis. The image time selected for each case study area was also derived from the time of fires detected during the day time period in each case study area. The selection criteria for each case study area is detailed in Table 2.1.

The counts of valid context pixels, and the difference of the context pixel mean from the central pixel were obtained for each window size, for each image, for each of the case study areas used for analysis. A visual examination of the causes of contextual estimate variation was also conducted based upon the spatial distribution of the mean temperature differences calculated, over window sizes from 5×5 pixels to 11×11 pixels, for each site.

2.3 Results

2.3.1 AHI Full Disk Characterisation

Cloud is a major impediment to any surface temperature estimation, and the area covered by the AHI disk is no exception. At the 0500 UTC time point, on average 55.6% of assessable land surfaces on the AHI disk are covered by cloud, with cloud coverage over land surfaces ranging from 45% to 73% over the images analysed. Cloud cover is most common over the northerly quarter of the disk, with areas north of AHI image row 1500 experiencing 68 – 74% cloud cover over the period examined. A full breakdown of cloud cover statistics can be found in Table 2.2. These areas of cloud cover, as determined by the cloud mask product, were removed from the context analysis, and form the bulk of the missing data in the window examinations.

Table 2.3 supplies a breakdown of pixels that are in permanent deficit of sufficient contextual pixels for temperature estimation at each valid context percentage at each window size. A requirement of at least 75% of contextual pixel availability is quite restrictive given the landforms present, and at least 2.2% of all land pixels cannot obtain this number of adjacent contextual pixels in the 5×5 window. The numbers in this table are adjusted for all window levels preceding — an assessment of a 7×7 window for instance takes into account pixels at the 5×5 window at the same time to determine whether an estimation is possible over all of the context pixels available to the target. These target pixels suffer permanent obscuration, and these locations can be flagged as problematic for contextual calculation for all periods.

Table 2.4 shows the global mean and standard deviation for all target pixels available for assessment at each window level individually. This assessment is

2. Estimating Fire Background Temperature at a Geostationary Scale — An Evaluation of Contextual Methods for AHI-8

Table 2.2: Average and standard deviation of cloud coverage for the AHI land areas covered in the study. The figures are an aggregate of 36 images recorded at 0500 UTC as mentioned in Section 2.2.1, broken into horizontal slices of the AHI disk as shown in Fig 2.2.

AHI Image rows	# of land pixels	Mean % cloud	SD % cloud
0 - 500	526506	74.1	15.7
500 - 1000	714119	69.1	15.0
1000 - 1500	663172	68.1	13.7
1500 - 2000	420460	49.2	23.0
2000 - 2500	184404	54.2	19.3
2500 - 3000	366370	62.7	10.4
3000 - 3500	248687	55.3	12.4
3500 - 4000	643030	28.6	14.0
4000 - 4500	793030	37.3	16.7
4500 - 5000	103387	58.1	19.4

Table 2.3: Number and percentage of pixels that are lacking sufficient adjacent pixels to provide contextual estimation at various window sizes and percentages across the AHI disk. A total of 4,663,165 AHI land pixels were evaluated.

Window size	Percentage of context pixels required for assessment						
	>75%	>65%	>55%	>45%	>35%	>25%	>15%
5 × 5	103801	74712	46141	18523	10918	4840	2389
	2.23%	1.60%	0.99%	0.40%	0.23%	0.10%	0.05%
7 × 7	136747	97771	54351	25771	13842	7322	3873
	2.93%	2.10%	1.17%	0.55%	0.30%	0.16%	0.08%
9 × 9	165592	110470	61786	31008	17290	9436	4544
	3.55%	2.37%	1.32%	0.66%	0.37%	0.20%	0.10%
11 × 11	192298	129744	73595	37000	21033	11510	5563
	4.12%	2.78%	1.58%	0.79%	0.45%	0.25%	0.12%
13 × 13	217235	150574	86662	43558	24681	13651	6794
	4.66%	3.23%	1.86%	0.93%	0.53%	0.29%	0.15%
15 × 15	240738	165472	97107	49446	28451	15689	7549
	5.16%	3.55%	2.08%	1.06%	0.61%	0.34%	0.16%
17 × 17	263862	182197	106023	55620	31895	17482	8466
	5.66%	3.91%	2.27%	1.19%	0.68%	0.37%	0.18%
19 × 19	286131	195443	114230	60973	35605	19496	9159
	6.14%	4.19%	2.45%	1.31%	0.76%	0.42%	0.20%
21 × 21	307516	210405	122986	66290	38851	21809	10196
	6.59%	4.51%	2.64%	1.42%	0.83%	0.47%	0.22%
23 × 23	328452	226933	132790	71657	42888	24078	11199
	7.04%	4.87%	2.85%	1.54%	0.92%	0.52%	0.24%
25 × 25	348645	240456	142150	75910	46572	25839	12100
	7.48%	5.16%	3.05%	1.63%	1.00%	0.55%	0.26%

conducted where there is at least one contextual pixel available at the denoted window size for comparison. As can be seen there is a global tendency to

overestimate temperature from the available contextual pixels, and there is little change in central tendency once the window of examination grows beyond 11×11 . The variation of the temperature estimation rises with the increased distance of assessed pixels from the centre, although the distance from the central pixel becomes less of an influence on variation once the window of examination grows beyond 11×11 . Global statistics such as these hide some of the more interesting trends in the data, and Figure 2.4 shows the breakdown of mean and standard deviation by contextual pixel availability at each window.

Table 2.4: Mean and standard deviation of the contextual estimate differences from central brightness temperature (AHI Band 7) for all available pixels in the 36 day set of full disk images at 0500 UTC. A total of 76,023,810 pixels were examined over the 36 images used in the study.

window size	5×5	7×7	9×9	11×11	13×13	15×15
mean (K)	0.037	0.031	0.029	0.027	0.025	0.024
std (K)	1.522	2.039	2.200	2.320	2.415	2.494
count	76023810	75858159	75871580	75880469	75888096	75893762
window size	17×17	19×19	21×21	23×23	25×25	
mean (K)	0.023	0.023	0.023	0.024	0.024	
std (K)	2.562	2.622	2.677	2.726	2.771	
count	75895983	75899037	75899238	75898553	75898041	

Figure 2.4a shows the mean value of the temperature difference as a function of the valid context percentage available at the outer edge of each window, apart from at the 5×5 window, where analysis includes all pixels inside this window. When all pixels are available for analysis at a particular window edge, the distance of the examined pixels from the central pixel has no influence upon the resulting temperature estimate, and the difference between estimates calculated using pixels from each window edge stays similar down to 75% of available pixels. At this point, having fewer pixels available in the 5×5 window of pixels causes a growth in temperature overestimation, which reaches a maximum when half of adjacent pixels are unavailable.

Figure 2.4b shows the standard deviation of the temperature difference as a function of the percentage of contextual pixels available, similar to Figure 2.4a. For all window sizes the standard deviation suffers a large increase once only one value is obscured in a window, with this effect most marked at the larger window sizes. Variation peaks in a similar fashion to the mean at around half of all contextual pixels available, with most window sizes seeing a levelling out of variation until only a handful of contextual pixels remain for estimation. The relative indifference to distance from the central pixel for the larger window sizes is due to the way pixels here are selected for analysis. The outer edge of the specified window is assessed, which is square in shape, and the pixels at

2. Estimating Fire Background Temperature at a Geostationary Scale — An Evaluation of Contextual Methods for AHI-8




Figure 2.4: (a) Mean brightness temperature difference between contextual estimates and the central pixel for the ring of pixels at the edge of each window across the full disk for 0500 UTC B07 AHI-8 images. (b) Standard deviation of contextual estimates derived from each window edge by percentage of available pixels in the window edge.

each outer edge exhibit a far greater range of distances from the central pixel as one moves further out, which would smooth out any purely distance-based variation.

The investigation into the effect of sensor zenith angle on temperature estimation found no marked influence. Mean values in the 5×5 window for temperature differences ranged from 0.07 K in the $0^\circ - 10^\circ$ view angle region, down to 0.025 K near the edge of the disk between $70^\circ - 80^\circ$ zenith angle over the images analysed. The largest errors were present in the two regions closest to nadir ($0^\circ - 10^\circ$ and $10^\circ - 20^\circ$), but the land surface area in these regions is much smaller than further out from the sensor nadir. There are no trends present due to sensor zenith angle in the standard deviation of contextual estimation either, apart from a slight drop in values close to nadir and at the $70^\circ - 80^\circ$ zenith angle.

2.3.2 Expanding Window Analysis

Table 2.5 demonstrates the breakdown of estimated pixel values when utilising an expanded window algorithm. Firstly, the rate reported in the 1.00 column represents the characteristics of pixels that have all contextual pixels available at the 5×5 window. These pixels, which make up 53.88% of all cloud-free pixels analysed, are generally underestimated by contextual methods, albeit only by 0.03 K, and display low variance. The other columns in the 5×5 row report statistics on the pixels that are added at each of the contextual percentage availabilities specified. For example, if a process accepted estimates with 45% or

more available contextual pixels, an extra 40.28% of all target pixels would be available for evaluation, in addition to the 53.88% from the full context (1.00) pixels. The additional pixels accepted at each valid context percentage have the means and standard deviations shown. For the remaining pixels without a solution, the examined contextual window is expanded through the window values shown, with statistics reported for pixels that achieve the valid context percentage at each window size. After the process is exhausted at the 25×25 window, the remaining pixels without a solution for each percentage are tallied in the total failures row at the bottom of the table.

The tendency of a target pixel's contextual surrounds to slightly underestimate temperature in optimal conditions, as seen in Figure 2.4, is also seen here in the 5×5 section of Table 2.5. As the threshold for valid contextual pixels is lowered, the mean temperature of all estimates rises and the variation in these estimates increase. Of course these trade-offs in temperature accuracy come with an increased level of coverage — accepting 65% contextual pixel availability allows 85.5% of all target pixels to be estimated with a neutral mean and relatively low variance. Conversely, accepting pixels at 15% contextual availability would allow for the calculation of temperature estimates over 99.2% of all target pixels, but with both higher mean and higher variance overall. Once the window of contextual pixels is expanded though, the accuracies coming from the contextual estimate deteriorates. In general the pixel's context tends to overestimate temperatures by an increasing amount, with mean temperature differences ranging between 0.47-2.11 K, and the standard deviation of results increases by around 50% by just moving from a 5×5 window to a 7×7 window of examination.

A further examination of calculation rates using the expanded window sizes is shown in Figure 2.5. For the portion of pixels that have no solution at the 5×5 window for each percentage, this figure shows the rate that target pixels subsequently obtain sufficient valid contextual pixels for calculation at each window size. The portion of target pixels that does not achieve sufficient contextual pixel counts for evaluation after expansion to the 25×25 window is shown in grey. As seen in Table 2.5, the higher contextual limitations have larger portions of the total data set suffer from insufficient data for estimation. Changing the acceptance percentage does not however affect the proportion of pixels that subsequently obtain sufficient contextual pixels for estimation at larger window sizes. This figure shows that no expanding window threshold will return values for more than 60.3% of the remaining pixels that fail to be calculated at the 5×5 window size, with the 75% threshold yielding less than 20% of extra pixels at larger windows. Of the pixels that do manage to obtain solutions, on average at least 69.5% of those occur at the 11×11 or

2. Estimating Fire Background Temperature at a Geostationary Scale — An Evaluation of Contextual Methods for AHI-8

Table 2.5: Mean and standard deviation of brightness temperature differences between the central pixels and the contextual surrounds at each window level per percentage level. Numbers shown in the 5×5 window row report statistics for pixels that would be added to the 1.00 pixels if the valid context percentage shown was used to accept contextual estimates. The percentage of total pixels with estimates available at the 5×5 window for each valid context percentage is also shown. The rows for each subsequent window size describe the number of temperature estimations that would be added from failures at the previous window size by expanding the examined window, and the subsequent means and variances of pixels included from these window sizes. A total of 76,023,810 pixels were examined over the 36 images used in the study.

window		1.00	Valid portion of total context pixels						
			0.99 - 0.75	0.99 - 0.65	0.99 - 0.55	0.99 - 0.45	0.99 - 0.35	0.99 - 0.25	0.99 - 0.15
5×5	mean (K)	-0.029	-0.083	-0.036	0.006	0.063	0.086	0.100	0.111
	std (K)	1.090	1.603	1.699	1.759	1.839	1.867	1.885	1.898
	count	40958274	18106490	24041100	27144999	30622724	32199899	33351687	34480202
	% avail	53.88%	23.82%	31.62%	35.71%	40.28%	42.36%	43.87%	45.35%
Total 5x5 success			77.69%	85.50%	89.58%	94.16%	96.23%	97.75%	99.23%
7×7	mean (K)	-	0.474	0.772	0.925	1.059	1.029	0.995	0.940
	std (K)	-	2.314	2.538	2.667	2.734	2.711	2.705	2.768
	count	N/A	1651297	948803	1096828	557910	562891	407382	160575
	% avail	N/A	2.17%	1.25%	1.44%	0.73%	0.74%	0.54%	0.21%
9×9	mean (K)	-	0.704	1.007	1.143	1.293	1.270	1.193	1.143
	std (K)	-	2.592	2.778	2.874	2.932	2.908	2.914	2.999
	count	N/A	502591	369619	289271	182700	127933	134785	53782
	% avail	N/A	0.66%	0.49%	0.38%	0.24%	0.17%	0.18%	0.07%
11×11	mean (K)	-	0.889	1.193	1.341	1.498	1.476	1.381	1.310
	std (K)	-	2.757	2.940	3.050	3.075	3.086	3.054	3.197
	count	N/A	320616	262912	221789	155173	118434	87380	36791
	% avail	N/A	0.42%	0.35%	0.29%	0.20%	0.16%	0.11%	0.05%
13×13	mean (K)	-	1.024	1.321	1.491	1.615	1.611	1.521	1.471
	std (K)	-	2.860	3.055	3.161	3.200	3.228	3.221	3.348
	count	N/A	228249	199477	177211	130197	102158	63145	27398
	% avail	N/A	0.30%	0.26%	0.23%	0.17%	0.13%	0.08%	0.04%
15×15	mean (K)	-	1.137	1.445	1.597	1.739	1.726	1.600	1.551
	std (K)	-	2.982	3.165	3.252	3.273	3.286	3.325	3.410
	count	N/A	174901	158520	121066	93067	63103	48553	21233
	% avail	N/A	0.23%	0.21%	0.16%	0.12%	0.08%	0.06%	0.03%
17×17	mean (K)	-	1.224	1.585	1.702	1.830	1.804	1.765	1.626
	std (K)	-	3.032	3.283	3.333	3.371	3.436	3.437	3.449
	count	N/A	139247	108539	105588	70645	58638	38539	14115
	% avail	N/A	0.18%	0.14%	0.14%	0.09%	0.08%	0.05%	0.02%
19×19	mean (K)	-	1.328	1.694	1.818	1.953	1.875	1.834	1.702
	std (K)	-	3.177	3.358	3.414	3.445	3.450	3.507	3.610
	count	N/A	113322	93057	79027	54876	46985	31733	12024
	% avail	N/A	0.15%	0.12%	0.10%	0.07%	0.06%	0.04%	0.02%
21×21	mean (K)	-	1.416	1.747	1.867	2.046	2.020	1.885	1.805
	std (K)	-	3.265	3.380	3.471	3.556	3.573	3.595	3.866
	count	N/A	94179	81879	71265	51677	33939	27491	10239
	% avail	N/A	0.12%	0.11%	0.09%	0.07%	0.04%	0.04%	0.01%
23×23	mean (K)	-	1.422	1.817	1.951	2.043	2.040	1.948	1.911
	std (K)	-	3.288	3.502	3.572	3.591	3.657	3.646	3.883
	count	N/A	80631	73046	63430	48480	36557	23016	9168
	% avail	N/A	0.11%	0.10%	0.08%	0.06%	0.05%	0.03%	0.01%
25×25	mean (K)	-	1.547	1.877	2.025	2.079	2.110	1.988	2.024
	std (K)	-	3.342	3.548	3.549	3.575	3.661	3.556	3.886
	count	N/A	70008	64301	51988	40127	27803	20150	8127
	% avail	N/A	0.09%	0.08%	0.07%	0.05%	0.04%	0.03%	0.01%
Total failures			13584005	8664283	5643074	3057960	1687196	831675	231882
			17.87%	11.40%	7.42%	4.02%	2.22%	1.09%	0.31%

window or lower, and 83.4% occur at window sizes at or smaller than 15×15 . This rate of return for the expanding window method, coupled with the vari-




Figure 2.5: Breakdown of temperature estimation pass rate on pixels that have no solution in their 5×5 window. The percentage of pixels covered by each bar this figure as a portion of all pixels examined is shown at the top of the figure. Each bar in the figure represents a minimum percentage level of valid contextual pixels for temperature calculation, and each coloured section represents the portion of pixels that are successful in deriving an estimate at each window size. The balance of exhausted pixels with no solution at each assessed percentage is also shown.

ability of results coming from estimations made at the larger window sizes, calls into question the overall effectiveness of using such a method, especially considering the computationally intensive nature of using pixels from a wider area.

Often in the case of some of the LEO fire products, an absolute cutoff threshold is used in order to calculate temperatures where a certain number of pixels are available for the calculation, regardless of their distance from the central pixel. A table demonstrating the effect of using a valid pixel threshold of 10 or more pixels is shown in Table 2.6. This table does not show valid percentages above 45%, as pixels that are only valid at these higher percentages trigger the absolute pixel threshold at the 5×5 window. The 10 pixel threshold homogenises the 45%, 35% and 25% classes to an extent, with very similar means and standard deviations emerging from each window size. Setting an absolute threshold of valid pixels does increase the total number of pixels that obtain temperature estimates, but even so there is still a number of pixels that a so-

lution is not possible for even at the lowest percentages. In comparison to the figures presented in Table 2.5, the estimated means at the higher window percentages using the absolute threshold are reduced, and the variation of temperature estimates smooths out once the window expands beyond 9×9 . This is due to more pixels in the original analysis expanding the window further than what was required to provide reasonably accurate temperature estimation. The major improvement from using an absolute pixel threshold is in the total percentage of pixels that are assessable, with the first two window sizes able to provide estimates in $\geq 98\%$ of cases in all percentage classes.

2.3.3 Case Study Areas

Figures 2.6 and 2.7 show the spatial distribution in the mean of the temperature differences at the 5×5 window for each of the case study areas, along with a histogram of the counts of these temperature differences per area. Each of the case study areas display a unique distribution. South-east Australia (Figure 2.6a), Thailand (Figure 2.6d) and Japan (Figure 2.7b) show marked linear features which line up with boundaries of land use areas. South-eastern Australia area has the most variation in the west where forested areas open to grazing and croplands, whilst Thailand and Japan have greatest variation in line with changes in relief. The Japan case study area has the most variation at the tree line high on Honshu's central range. The effect of coastline pixels is most evident in the Borneo area (Figure 2.6c), with the influence of swamp and mangrove along the coastline leading to an underestimation of temperatures in adjacent pixels. Urban areas are also a source of underestimation, most prevalent in the central China study area (Figure 2.7a) where cities in the north west of the area display a heat island effect. This effect is also seen to a lesser extent in the south-east Australia and Japan study areas. The Siberian (Figure 2.7c) area displayed relative uniformity outside of the central latitudes, where unmelted snow from mountain ranges caused commission errors in the cloud mask used, which led to large estimation errors on these interfaces. North-western Australia (Figure 2.6b) is characterised by high local variability, and high contrast between vegetated and bare earth areas coupled with the lack of surface moisture increases this local variability (shown in greater detail in Figure 2.9). All distributions of temperature differences are relatively uniform in nature, with the Japan, Siberia and Thai areas displaying longer tails than other areas.

Table 2.7 depicts the global mean and standard deviations of the case study areas compared to the outer edge of pixels at various window sizes. The general trend of overestimation of pixel temperatures when looking at the global

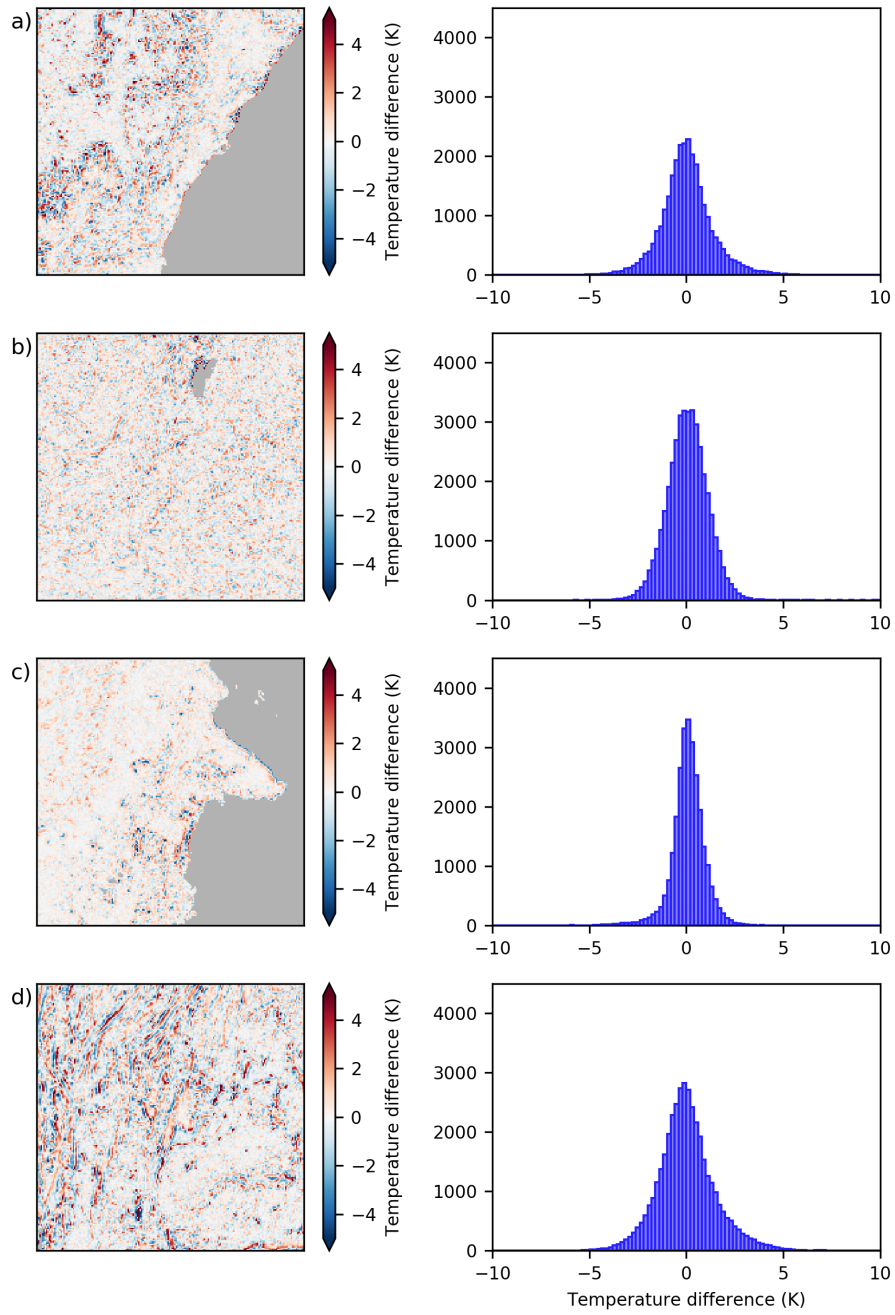


Figure 2.6: Mean difference between contextual estimates and the central pixel for the selected period for each area. (a) south-eastern Australia (sea); (b) north-western Australia (nwa); (c) Borneo (bor); and (d) central Thailand (thl).

2. Estimating Fire Background Temperature at a Geostationary Scale — An Evaluation of Contextual Methods for AHI-8

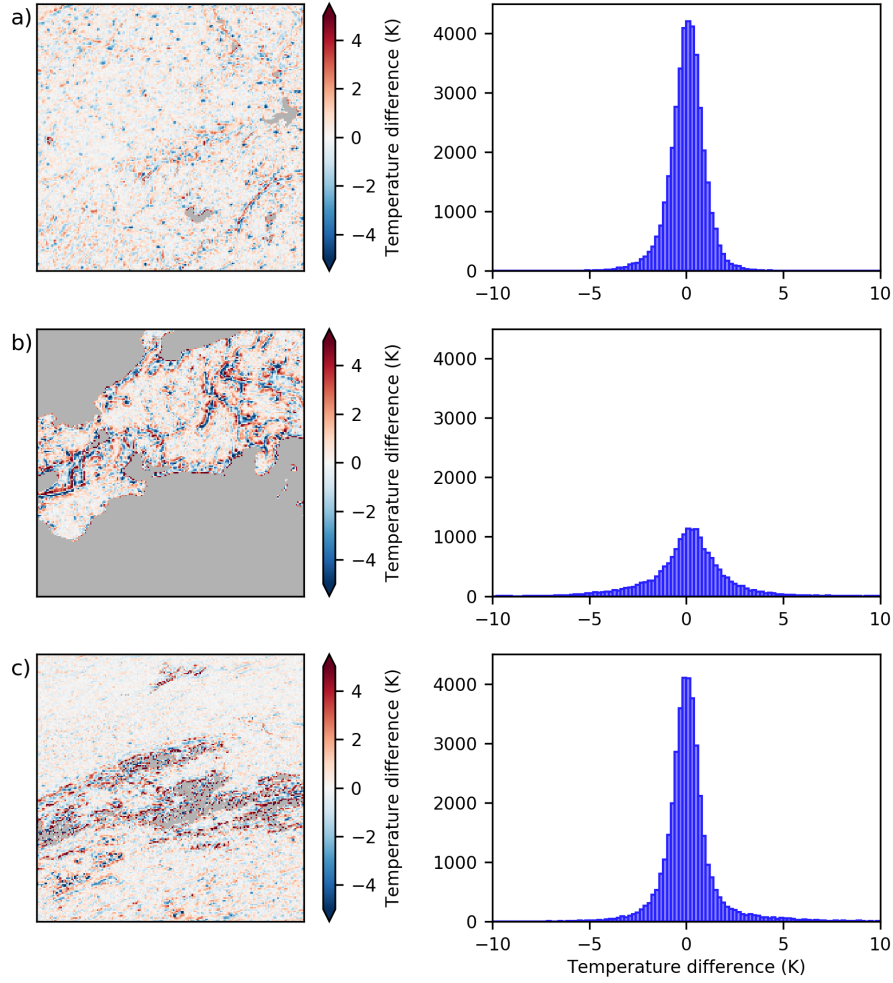


Figure 2.7: Mean difference between contextual estimates and the central pixel for the selected period for each area. **(a)** eastern China (chn); **(b)** central Honshu (jpn); and **(c)** Siberia (sib).

statistics is shown here, but the change in mean values is different from area to area. Stability in the mean temperatures here is a function of the amount of clear sky present during the times examined — Thailand for instance has a comparatively small number of pixels affected by cloud during the examined period, whereas Japan and Siberia are heavily cloud affected during their examined time periods. North-western Australia shows marked improvement in temperature recovery when looking at the more distant window edges, which is seemingly due to poor performance at the 5×5 window size. All areas have a notable gain in the temperature variance as the pixels examined become more

distant from the central pixel.

Table 2.8 reports statistics for each of the case study areas broken down by valid contextual pixel percentage. As can be seen in all areas, pixels with all contextual pixels available for calculation tend to underestimate the target temperature. An increasing tendency to overestimate temperature as the amount of contextual pixels available reduces is present at all sites. The stability of temperature estimation from a pixel with no contextual obscuration is also much better than from areas that are partially obscured. Some of the case study areas display a much larger variance once contextual pixels become partially obscured - the north-western Australia area is the median for variance during full availability, but is the worst performer once the contextual area is even slightly obscured. The trend of greater overestimation as obscuration of contextual pixels increases is caused by the target pixel temperature dropping due to cloud shadows causing lower solar reflectivity, in comparison to clearer and brighter valid pixels in the surroundings. The expected deterioration of accuracy for each of the percentage windows is seen clearly, with standard deviations increasing as more obscured estimations are accepted. The south-east Australia, Thailand and China areas display less variation than other areas as the percentage of valid contextual availability decreases. With regard to numbers of target pixel estimates available at each contextual percentage, these examples display a slight inflection in their trend around 45%, with numbers of estimates available increasing in greater quantities below this percentage and at lesser quantities above. Total recovery rates by percentage can be calculated by adding the percentage availability to the obscuration-free contextual (1.00) values.

Moving further away from the central pixel has the most marked effect on temperature variation, and this effect can be seen in Figure 2.8. This figure depicts the changes in the spatial and statistical distribution of contextual temperatures over the south-eastern Australian study area, for window sizes between 5×5 pixels and 11×11 pixels. Expanding the window of examination for pixel estimation exacerbates the edge effects seen in the east and south-eastern portions of this area, with much larger areas of high variation on the boundaries seen previously. The greater window size also highlights the larger variations at the urban interfaces of Sydney and the Illawarra region, and shows a general overestimation of temperatures along the coastline. The distributions of temperatures remain normal, but are flattened considerably compared to values from the most adjacent pixels.

2. Estimating Fire Background Temperature at a Geostationary Scale — An Evaluation of Contextual Methods for AHI-8

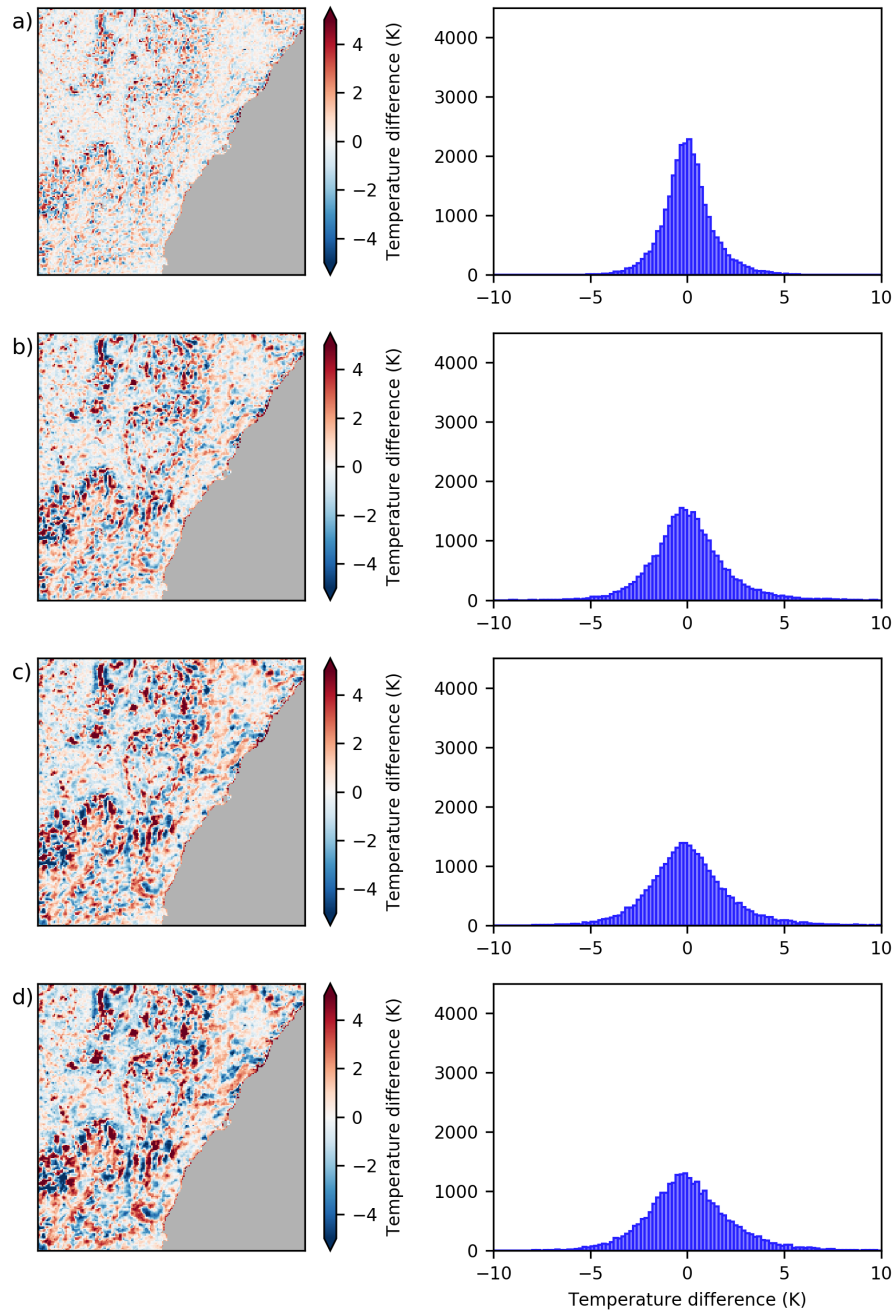


Figure 2.8: Changes in spatial and statistical distribution of temperature estimates for the south-eastern Australia (sea) study area by window size. Window levels shown are (a) 5×5 window; (b) 7×7 window; (c) 9×9 window; and (d) 11×11 window.

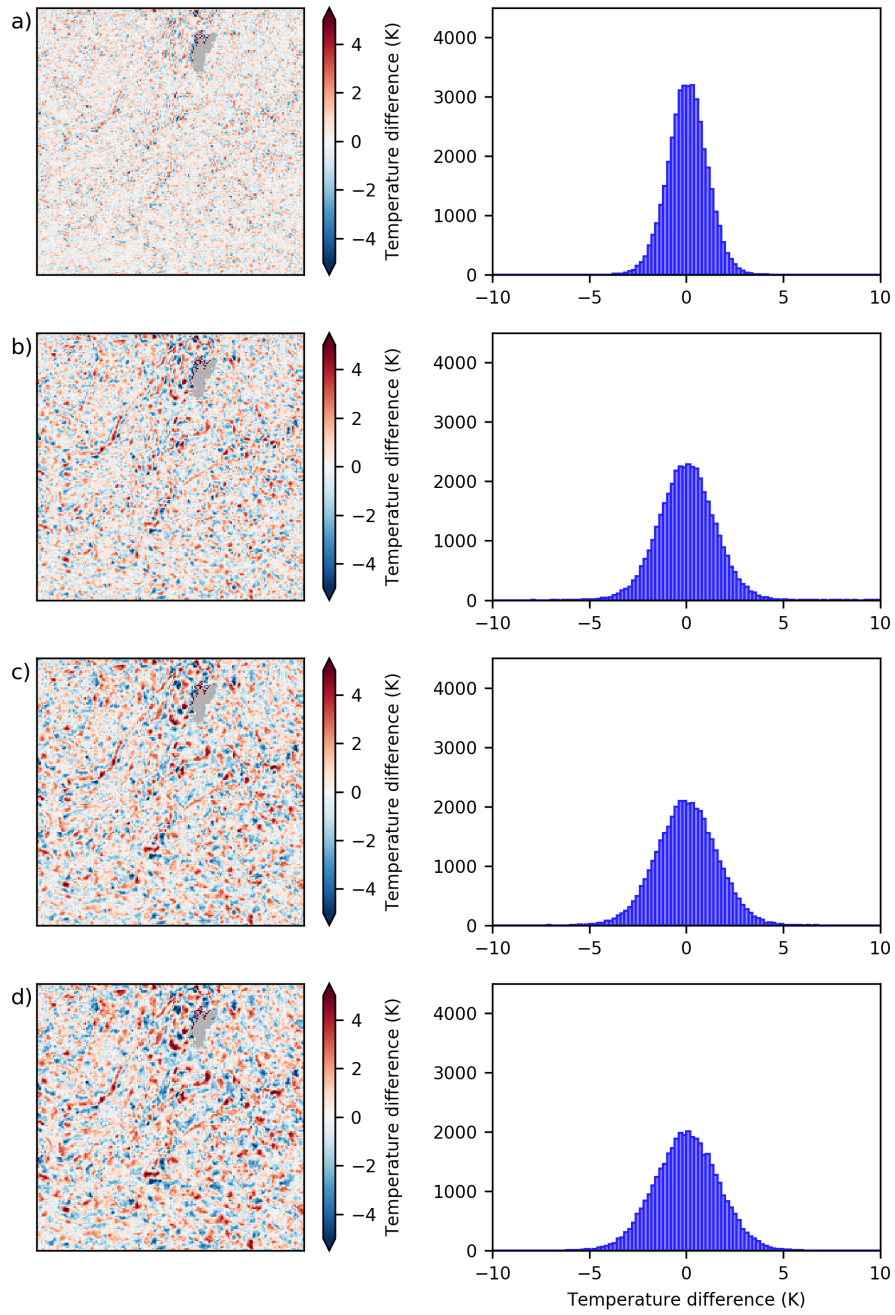


Figure 2.9: Changes in spatial and statistical distribution of temperature estimates for the north-western Australia (nwa) study area by window size. Window levels shown are (a) 5×5 window; (b) 7×7 window; (c) 9×9 window; and (d) 11×11 window.

2.4 Discussion

Whilst the numbers presented in Section 2.3.1 are specific to the AHI disk coverage area, the same factors that restrict calculation of background temperature should be common to any part of the globe where fire detection and attribution occurs. Cloud coverage is a major inhibiting factor in any satellite fire detection setup, and areas that display even moderate occlusion of the contextual surroundings tend to present less than ideal estimations of temperature. From the range of values of contextual availability shown in Figure 2.4a, there seems to be a break between results derived from pixels with at least 65% contextual availability and results from pixels with less contextual values available. Usage of estimates from target pixels with at least 65% available contextual information minimises the bias in the mean calculation of background temperature, especially at the larger window sizes, whilst also limiting the variation of the resultant estimations. The results presented in both Table 2.4 and Figure 2.4 also demonstrate the relative stability of temperatures derived from window sizes larger than 13×13 , or in AHI scale once pixels are at least 12 km from the pixel being estimated. If an increase in variance of calculated estimates of 60% over values derived at the 5×5 is acceptable for a specific purpose, then there is seemingly no reason not to set the initial area of examination for contextual temperature as large as practicable, but if this temperature variance is more of a concern, then using pixels from outside even the 11×11 window of pixels becomes problematic.

The effects at play when calculating contextual estimates as shown in Figure 2.4 bear further examination. The relative differences between the mean and variation seen at the higher window sizes reduces as the pixels examined increase in distance from the target, an effect noted in Section 2.3.1 being due to variations in the window edge radius. Examination of the effect of using pixels with similar distances to the target, in a circular ring, would most likely bear this out, though implementation of such a distance-based window of examination would become less trivial as sensor zenith angle increases. The pattern of mean difference as a function of valid pixels is worth mentioning as well, especially with regard to overestimation of the target temperature when valid contextual pixels approach 50%. This effect is likely due to shadowing of the target pixel and consequent reduction in solar reflectivity, with the target pixel most likely being immediately adjacent to the obscuration affecting the surrounding pixels. This effect is lessened in the rings of pixels situated further from the target pixel, as the source of obscuration at the outer edge of the window is less likely to be present closer in to the target pixel. This overestimation is not particularly large in magnitude, and is less likely to affect fire

detection for instance, but such information may assist in the adjustment of temperature-controlled metrics calculated from these estimates.

The results also cast the use of expanding windows for contextual temperature examination in a poor light, particularly for those sensors with larger spatial resolutions. The vast majority of all pixel calculations are achieved at the 5×5 window, with the recovery of data from using an expanding window ranging from 20% to 54% of all remaining target pixels. If we are to use the 65% window as an example, 85% of data is contributed from the 5×5 window, extra estimates from using the expanding window are just over 4%, and the majority of those extra estimates occur at or below the 11×11 window. There are also compromises involved in using the estimates, with a general positive bias and much higher variation in values at even the 7×7 level. Depending on the purpose of using these estimates, using the data coming from the combined windows could be detrimental to overall reporting accuracy. When evaluating how a background temperature method should be implemented, care needs to be taken to ensure that any need for comprehensive coverage, whether it be achieved by either using a smaller percentage of valid contextual pixels, by using larger window sizes, or both, does not inhibit the accuracy of the overall product.

With regard to the case study areas selected for analysis, the reasons for major variances in contextually determined temperature are as diverse as the case study sites selected. Phenomena affecting contextual estimation range from highly ephemeral conditions, such as fire and flooding, to seasonally changing influences such as snow and vegetation cover, to semi-permanent influences like urban-rural interfaces and land cover change, and on to permanent conditions such as relief, tree lines and coastlines. Each of these influencing factors need to be treated in a different way dependent upon the expected temporal duration of phenomena. Whilst setting global thresholds is satisfactory for more holistic measures such as carbon emissions and global FRP [73], in order to obtain more accurate estimates of pixel contrast, for metrics which require more accurate estimates of pixel temperature, use of a contextual method may require application of a-priori information. Conversely, a method that takes local variation into account by using such information needs to take into account the changes caused by more short-term influences mentioned here. This adds complexity to any system that uses fire background temperature in a rapid fashion, such as in active fire response.

Whilst this study demonstrates the effectiveness of contextual estimation when conditions are amenable, the deterioration of temperature estimation fidelity, and in some cases total loss of recovery, leads to investigation of other methods that may be able to bridge the gap in temperature retrieval. Investiga-

tion should be encouraged into the leveraging information from the temporal domain when looking at this problem. Methods such as those used in [65, 83] & [31] look at the diurnal temporal domain for temperature estimation, which is more suited to geostationary sensors such as AHI and GOES. This does not preclude the use of temporal information for LEO products though. An approach to integration of temporal modelling of background temperature could look at adjustment of measurements by images from previous time periods, with adjustments made for factors such as time of image capture. Looking at many different time points would provide redundancy against ephemeral conditions such as cloud, but looking too far back in time can lead to information not being representative of the current state of the landscape. A mix of ephemeral, seasonal and annual adjustments should be examined for their effectiveness in correcting estimated values for LEO-based products.

With regard to the direct applicability of these results to products and values from other sensors, caution should be exercised. The pixel sizes examined here from the AHI-8 sensor are much larger than their equivalents from images taken by low earth orbiting sensors. The rapid changes in landforms and land cover types seen in the case study areas may be smoothed or exacerbated by using smaller pixels, and the overall granularity of spatial homogeneity at varying scales should be taken into account when making comparisons across products and sensor scales. Sensor dependent effects such as sensor point spread function have also not been examined here, although these effects are mostly seen when dealing with high temperature anomalies in the MWIR band, which the vast majority of target pixels in this study do not encounter. The orbit of the sensor used in this study also grants the opportunity to examine targets at the same local time over many images, and application of methods used for analysis of LEO sensor information in a similar fashion would need to take into account variations in the time of image capture for longitudinal analysis purposes.

This study has assessed the overall ability to estimate background temperature from spatial context using AHI. In this study, temperature estimates from pixels with all context pixels available show a standard deviation of 1.09 K when examined across the full disk. In comparison, the global standard deviations for the case study areas were higher, ranging from 1.12 K in Siberia to 2.06 K in Japan. Whilst accuracy of background temperature is less emphasised for metrics such as FRP, information obtained from this study could be used in an adjustment of these metrics as calculated from AHI. Knowledge about the expected variation of medium-wave infrared radiation estimation may also play a role in development of new fire detection techniques, which use the expected variation of MWIR radiation in an area to identify anomalous values as a first-

pass filter. Providing simpler and more concise algorithms for fire detection reduces the data volumes and processing overhead required, leading to more rapid production and application of results.

2.5 Conclusions

An analysis of the effectiveness of contextual calculation of pixel background temperature has been conducted for a 36 image set from Band 7 from the AHI-8 sensor. Results show that estimates made from unobscured context pixels are very accurate, with a slight negative bias and low variation of temperature differences. Accuracy of the contextual method deteriorates with decreasing contextual pixel availability, with 65% a good balancing point between increased bias and variation of calculated values, and the overall availability of contextual data for estimation. Using a growing window for increasing the pixel availability by leveraging a larger window size decreases the accuracy of estimation results, with much larger values of bias and variation in resultant temperatures. Care needs to be taken with expanding window methods in order to balance comprehensive coverage of image data against the accuracy required from use of the results. A wide range of influences cause variation in temperature estimation, with each of the case study areas examined providing both unique problems for contextual estimation, and placing emphasis on the need for knowing the conditions specific to an area in order to provide highly accurate temperature estimation. Comprehensive coverage of all land areas is not achievable using contextual estimation, and in most cases is not desirable due to the deterioration of results as estimates use less optimal data. Alternative methods for temperature estimation need to be explored in order to overcome the limitations of contextual-based algorithms presented here, particularly when used with large resolution sensors such as AHI-8.

2.6 Thesis Context

This chapter outlines the error analysis of the commonly used contextual brightness temperature estimation method used by most of the common satellite-based fire detection products. The work highlights the need to explore other methods for brightness temperature estimation, especially in areas with high land cover variability. The next chapter provides the first exploratory examination of a method that discards examining the pixel context in favour of providing a robust diurnal fitting technique for temperature estimation.

2. Estimating Fire Background Temperature at a Geostationary Scale — An Evaluation of Contextual Methods for AHI-8

Table 2.6: Mean and standard deviation of brightness temperature differences between the central pixels and the contextual surrounds at each window level per percentage level, or where number of context pixels reaches 10. The 5×5 window statistics show the global rates for pixels which have equal or greater contextual pixels than the minimum for estimation. The rows for each window size describe the number of calculated values that would be added by expanding to each window size, and the subsequent means and variances of pixels included from these window sizes.

window		1.00	Valid portion of total context pixels			
			0.99 - 0.45	0.99 - 0.35	0.99 - 0.25	0.99 - 0.15
5×5	mean (K)	-0.029	0.076	0.086	0.100	0.111
	std (K)	1.090	1.856	1.867	1.885	1.898
	count	40958274	31473186	32199899	33351687	34480202
	% avail	53.88%	41.40%	42.36%	43.87%	45.35%
Total 5x5 success			95.27%	96.23%	97.75%	99.23%
7×7	mean (K)	-	0.709	0.746	0.874	0.940
	std (K)	-	2.550	2.568	2.642	2.768
	count	N/A	2456495	1734495	664734	160575
	% avail	N/A	3.23%	2.28%	0.87%	0.21%
9×9	mean (K)	-	0.623	0.628	0.703	0.996
	std (K)	-	2.639	2.640	2.673	2.928
	count	N/A	591757	589044	531807	97775
	% avail	N/A	0.78%	0.77%	0.70%	0.13%
11×11	mean (K)	-	0.544	0.548	0.588	0.854
	std (K)	-	2.723	2.723	2.745	2.963
	count	N/A	225018	224240	212723	119473
	% avail	N/A	0.30%	0.29%	0.28%	0.16%
13×13	mean (K)	-	0.485	0.487	0.518	0.701
	std (K)	-	2.789	2.792	2.816	2.971
	count	N/A	108023	107653	103138	66691
	% avail	N/A	0.14%	0.14%	0.14%	0.09%
15×15	mean (K)	-	0.448	0.451	0.481	0.637
	std (K)	-	2.828	2.831	2.852	3.018
	count	N/A	60176	59952	57566	39017
	% avail	N/A	0.08%	0.08%	0.08%	0.05%
17×17	mean (K)	-	0.413	0.414	0.435	0.584
	std (K)	-	2.844	2.845	2.869	3.019
	count	N/A	37688	37596	36118	24821
	% avail	N/A	0.05%	0.05%	0.05%	0.03%
19×19	mean (K)	-	0.401	0.403	0.434	0.562
	std (K)	-	2.864	2.867	2.897	3.057
	count	N/A	25000	24899	23883	16827
	% avail	N/A	0.03%	0.03%	0.03%	0.02%
21×21	mean (K)	-	0.439	0.441	0.464	0.607
	std (K)	-	2.996	3.000	3.031	3.226
	count	N/A	17483	17419	16712	12002
	% avail	N/A	0.02%	0.02%	0.02%	0.02%
23×23	mean (K)	-	0.316	0.318	0.341	0.428
	std (K)	-	2.913	2.919	2.943	3.092
	count	N/A	12125	12068	11667	8478
	% avail	N/A	0.02%	0.02%	0.02%	0.01%
25×25	mean (K)	-	0.304	0.306	0.324	0.415
	std (K)	-	2.869	2.874	2.897	2.998
	count	N/A	8910	8867	8596	6289
	% avail	N/A	0.01%	0.01%	0.01%	0.01%
Total failures			49675 0.07%	49404 0.06%	46905 0.06%	33386 0.04%

Table 2.7: Mean and standard deviation of mean brightness temperature differences of each case study area for each 31 day period. Pixel values were averaged over the 31 day period for each site, and global means and standard deviations of these averages are reported.

Window edge	5 × 5		7 × 7		9 × 9		11 × 11	
Case study area	\bar{x} Mean	\bar{x} SD	\bar{x} Mean	\bar{x} SD	\bar{x} Mean	\bar{x} SD	\bar{x} Mean	\bar{x} SD
sea	0.031	1.312	0.051	1.891	0.063	2.090	0.079	2.229
nwa	0.059	1.031	0.024	1.440	0.022	1.570	0.021	1.658
bor	0.089	0.856	0.089	1.231	0.097	1.360	0.101	1.454
thl	0.022	1.481	0.021	2.202	0.023	2.469	0.024	2.673
chn	0.023	0.942	0.024	1.348	0.020	1.494	0.014	1.605
jpn	0.092	1.928	0.140	2.862	0.162	3.259	0.178	3.553
sib	0.112	1.370	0.134	1.810	0.144	1.939	0.152	2.026

Table 2.8: Mean and standard deviation of brightness temperature differences between the central pixels and the contextual surrounds at the specified percentage levels for the 5 × 5 window in each case study area. Each column reports the statistics of accepting the available pixels above the denoted percentage level. Pixels with full contextual coverage are reported in the 1.00 column.

		Valid portion of total context pixels								all
		1.00	0.99 - 0.75	0.99 - 0.65	0.99 - 0.55	0.99 - 0.45	0.99 - 0.35	0.99 - 0.25	0.99 - 0.15	
sea	mean (K)	-0.021	-0.042	-0.013	0.010	0.043	0.056	0.063	0.068	0.022
	std (K)	1.670	1.832	1.862	1.880	1.921	1.931	1.935	1.941	1.804
	count	279250	152220	210297	243132	284703	308590	330486	363534	688739
	% avail	40.5%	22.1%	30.5%	35.3%	41.3%	44.8%	48.0%	52.8%	
nwa	mean (K)	-0.051	-0.134	-0.035	0.034	0.127	0.171	0.198	0.218	0.061
	std (K)	1.377	2.576	2.769	2.856	2.953	2.992	3.014	3.029	2.216
	count	548125	258015	339147	382837	438487	470677	499270	538353	1129978
	% avail	48.5%	22.8%	30.0%	33.9%	38.8%	41.7%	44.2%	47.6%	
bor	mean (K)	-0.106	-0.096	-0.038	0.003	0.051	0.071	0.086	0.096	0.061
	std (K)	1.121	1.472	1.585	1.651	1.719	1.746	1.764	1.777	1.681
	count	90734	250567	343181	392122	451781	485259	515552	559085	702114
	% avail	12.9%	35.7%	48.9%	55.8%	64.3%	69.1%	73.4%	79.6%	
thl	mean (K)	-0.033	0.000	0.047	0.079	0.109	0.118	0.122	0.125	0.016
	std (K)	1.679	1.874	1.920	1.941	1.961	1.965	1.967	1.970	1.776
	count	683361	224582	281720	310807	346989	367880	386865	415359	1134791
	% avail	60.2%	19.8%	24.8%	27.4%	30.6%	32.4%	34.1%	36.6%	
chn	mean (K)	-0.032	-0.041	0.006	0.039	0.079	0.092	0.100	0.104	0.021
	std (K)	1.159	1.310	1.345	1.370	1.407	1.418	1.424	1.428	1.272
	count	428453	176040	232020	262412	301287	324324	346985	384005	868807
	% avail	49.3%	20.3%	26.7%	30.2%	34.7%	37.3%	39.9%	44.2%	
jpn	mean (K)	-0.019	-0.151	-0.134	-0.056	0.079	0.102	0.116	0.125	0.046
	std (K)	2.061	2.246	2.269	2.332	2.460	2.479	2.486	2.490	2.265
	count	120759	54546	74758	86968	103879	114110	124201	141136	288787
	% avail	41.8%	18.9%	25.9%	30.1%	36.0%	39.5%	43.0%	48.9%	
sib	mean (K)	-0.057	-0.073	-0.017	0.020	0.066	0.080	0.088	0.092	0.037
	std (K)	1.120	1.746	1.814	1.859	1.947	1.969	1.980	1.996	1.745
	count	86220	66918	97011	117111	149287	173672	202360	260949	478458
	% avail	18.0%	14.0%	20.3%	24.5%	31.2%	36.3%	42.3%	54.5%	

A Broad-Area Method for the Diurnal Characterisation of Upwelling Medium Wave Infrared Radiation

3

3.1 Introduction

Active fire in the landscape is a major catalyst for environmental change, potentially resulting in large socio-economic impacts, including the high costs and risks associated with mitigation efforts and the disruptive evacuation of communities[4]. Fire authorities and land managers are constantly seeking new techniques for the early detection of fire to assist in the timely informing and evacuation of the public from at-risk areas, the planning and prioritisation of asset management strategies, and feasibility assessment of possible suppression efforts. This requirement for active fire detection in near real-time has seen the adoption of remote sensing from satellite sensors as an objective means to quantify and characterise the location, spread and intensity of fire events to support these important decisions [68]. The information derived from this imagery can also be used in conjunction with other data to provide models of an event, leading to more accurate understanding of the potential impacts of an event before they occur.

Remote sensing for fire detection and attribution has predominantly focused on imagery from low earth orbiting (LEO) sensors, which have significant advantages with regard to spatial resolution, and therefore to the minimum size of fire that can be detected. The trade-off with sensors of this type is that their orbital parameters preclude rapidly repeated observations of a single location, and without a significant investment in capital to provide for more missions, the ability to provide real-time observations of fire from these sensors will be hampered by extensive revisit times. The necessity for rapid fire detection sees the focus of fire detection shift to imagery obtained from geostationary sensors, which provide an increased revisit rate at the cost of a loss of fidelity in the spatial and radiometric realms [7]. Despite this, the launch of new sensors such as the Japanese Meteorological Agency's Advanced Himawari Imager (AHI) and the NOAA's Advanced Baseline Imager (ABI) provide an enhanced opportunity to examine fire ignitions and evolution due to improved spatial, radiometric and temporal resolutions compared to their geostationary predecessors.

One of the physical limitations of some techniques used for the remote sensing of fire is determination of the background temperature of a pixel. Having an accurate measure of this temperature is vital in order to be able to classify a target pixel as containing a fire in the first place, along with being able to accurately estimate the area of the pixel containing fire and the intensity or ra-

This chapter was published in a peer-reviewed journal as: Hally, B., Wallace, L., Reinke, K., & Jones, S. (2017). A Broad-Area Method for the Diurnal Characterisation of Upwelling Medium Wave Infrared Radiation. *Remote Sensing*, 9(2), 167. <https://doi.org/10.3390/rs9020167>

diative output of the fire [68]. Background temperature tends to be a difficult value to determine accurately because of the obscuring effects of the fire's output, which outweighs the background signal from a pixel in the medium wave infrared. Early efforts to correct for this behaviour used a bi-spectral approach [11], which used the response of thermal infrared bands in the same area to develop an estimate of fire characteristics. Thermal infrared bands also display sensitivity to fire outputs but to a much lesser extent and are generally used for false alarm detection, especially for marginal detections from the medium wave infrared caused by solar reflection [34]. The difference between signal response in these two bands is the basis for most current geostationary fire detection algorithms and similarly with the analysis of LEO sensor data [68]. Issues with these algorithms start when looking at fires of smaller extents. A study by [22] highlighted issues with fire retrievals using the bi-spectral method, especially with regard to smaller fires and background temperature characterisation. The study found that misattribution of the background temperature by as little as 1K for fires that covered a portion of a pixel ($p \leq 0.0001$) could produce errors in fire area attribution by a factor of 100 or more, with a less significant error in temperature retrieval of $> \pm 200\text{K}$. This is of major concern for the use of geostationary sensors for detection, as fires in their early stages make up far less a proportion of a pixel from a geostationary sensor than is the case with a LEO sensor.

The most common method of deriving background temperature for a fire pixel is through the use of brightness temperatures of pixels adjacent to the target pixel [68]. By identifying a number of pixels in the immediate area that are not affected by fire or other occlusion such as smoke and cloud, an estimate can be found by aggregation of the brightness temperature of these pixels. The assumption is made that the adjacent pixels used are of a similar nature in terms of reflectance and emissivity to the target pixel. This background characterisation is then used in comparison to the target pixel in order to identify whether the fire signal is different enough from the background to constitute a fire return. Problems occur with this method when the background temperature is misrepresented. In a study by [25] estimation of the background temperature from adjacent pixels in approximately 22% of cases produced a background temperature that was higher than the brightness temperature of the detected fire, based upon the surface variability of the area surrounding the detection. This study also analysed the general performance of the MODIS bi-spectral fire detection algorithm [20], and found that only 7% of the fire identified by the product could be accurately characterised for fire temperature and area. With the coarser spatial resolution of geostationary sensors, the authors noted that larger potential errors will affect the retrieval of

fires in comparison to LEO sensors when using these methods.

A promising method for background temperature determination is through the use of time series data from geostationary sensors. This time series data can be utilised based upon the premise that upwelling radiation can be predicted based upon incident solar radiation, which varies chiefly by time of day, with some variation due to weather effects and occlusion. Modelling of this Diurnal Temperature Cycle (DTC) has been approached using many different techniques. Earlier work on the modelling of the DTC looked to provide a parameter-based description based upon fitting to discrete mathematical functions, such as the model proposed by [27]. This work applied the modelled DTC estimate directly to measured brightness temperatures using empirically derived parameters. This approach tended to be insensitive to functional variation due to synoptic effects, and performed inadequately during the period of rapid temperature change in the early morning. The work of [84] was the first to utilise a set of prior observations as training data for a signal fitting process, using the mean of previous observations as a state vector for a Kalman filter, which due to the sensitivity of a mean-based estimate to outlying observations, application was limited to cloud-free data only.

The influence of outliers on the training data used for signal fitting was addressed in part by the study of [65], who looked at a selective process whereby previous days DTCs were included in the training data of a pixel based upon the amount of disturbance in the day's observations, with a limit of six cloud or fire affected observations out of a 96 image DTC permitted. These limits eliminated much of the effects of outliers on the subsequent single value decomposition (SVD) used for the initial fitting of background temperature. Issues occurred in areas where there were insufficient anomaly-free days for a fitting to be performed, even with a sampling size of the previous thirty days, in which case DTCs were selected from a library of known anomaly-free DTCs from a similar area. The process was reliant on an accurate cloud mask to determine which days were anomaly-free, and the training data derivation was data intensive, with DTC vectors having to be extracted and calculated for each individual pixel prior to fitting using the SVD process. These issues lead to training data fragility, and introduced some of the issues that are common error sources in contextual algorithms for fire detection.

In order to address the issues caused by sampling training data from a pixel-based approach, this paper presents a new method for deriving training data based upon a broad-area method. This method exploits similarities in incident solar radiation found at similar latitudes to derive training data for a pixel. Geostationary sensor data is aggregated by latitude and an area's local solar time, and formed into a time series based upon a sensor's temporal resolution.

This paper will compare the results obtained using this method to training data derived from individual pixels, such as in the study by [65], and with contextual methods of background temperature determination, to compare the accuracy, efficiency and availability of each method.

3.2 Materials and Methods

3.2.1 Himawari AHI-8 Data

The Japanese Meteorological Agency (JMA) launched the Himawari-8 geostationary satellite in October 2014, which replaced the MTSAT-2 sensor as the main source of meteorological data for Australia. The sensor onboard the Himawari-8 satellite is the Advanced Himawari Imager (AHI), a sensor capable of measuring 16 bands of imagery, ranging from visible light through to thermal infrared. The radiometric resolution of the sensor is significantly improved compared to previous iterations of MT-SAT, with quantization of 11 bit in the visible and near infra-red and up to 14 bits for the MWIR Band 7, which is the band most heavily utilised for fire detection purposes. The sensor boasts a spatial resolution of 2×2 km at nadir for infra-red channels, with visible imagery available at resolutions as high as 0.5×0.5 km. Scanning of the full disk generally takes place every ten minutes, with the area of Japan and surroundings scanned every 150 seconds. The coverage area of images provided by the sensor ranges from western India to Hawaii longitudinally, encompassing much of China, Siberia, Japan, South East Asia, Indonesia and Australia. Characteristics of AHI and MTSAT-2 for fire detection can be found in Table 3.1.

Table 3.1: Comparison of MTSAT-2 and AHI-8 sensors for fire detection using a MWIR ($\sim 4\mu\text{m}$) channel.

Sensor	MTSAT-2	AHI-8
Temporal Resolution	30 min	10 min
Spatial Resolution (nadir)	4×4 km	2×2 km
Medium wave infra-red channel saturation temperature	330 K	400 K
Quantisation (MWIR)	10 bit	14 bit
Noise equivalent delta temperature ($NE\Delta T$)	0.09 @ 300 K	≤ 0.16 @ 300 K

The increase in image frequency from the AHI sensor in particular allows for greater utilisation of time series data for temperature estimation, with gaps in the data caused by station keeping, or sensor recording and processing issues having less of an effect on derived values. Improvements in radiometric resolution allow a greater range of measurements to be recorded, improving

the fidelity of anomaly detection, and a greater saturation temperature allows for the characteristics of larger fires to be determined. Preliminary work with this sensor in the fire detection space [30, 88] has demonstrated the ability to detect and track fires using multiple bands of imagery from this satellite sensor.

Inclusive of the housekeeping periods of the satellite, which occur twice daily, 142 full disk images are available from AHI per day. This study utilises image products published by the Australian Bureau of Meteorology (ABOM), derived from raw AHI images by the process outlined in [37] for Band 7 brightness temperature corrections from raw satellite imagery. The study also makes use of the clear sky probability product, adapted by ABOM from the CLAVR-x cloud mask algorithm outlined in [35] with some minor alterations [54] for verification of cloudy periods in the temporal data stream. Cloud product data for AHI is generally available at ten minute intervals, which is similar to raw satellite images.

3.2.2 Training Data Derivation

Upwelling radiation in the medium wave infra-red consists of two main components: reflection of solar radiation from the earth's surface, and emission of blackbody radiation from the surface due to temperature. The skin surface temperature of the earth is influenced by a number of factors. Rain reduces upwelling radiation by both lowering surface albedo, and reducing surface temperatures and consequently emission. Convective cooling and heating due to air masses can influence surface temperatures, and land cover composition can affect the magnitude of temperature change. However, the most notable influence is heating by solar radiation. For any area of land on a given day at a given latitude, the amount of incident solar radiation should be similar barring obscuration. This should lead to a similar relative diurnal response of upwelling radiation from the land surface at a given solar time t_s . By breaking the land surface into $0.25^\circ \times 0.25^\circ$ blocks, we can use the rotation of the earth to construct an amalgamated time series, with each block translating to a separate minute of local solar time.

Images from the sensor have an ocean mask and a rudimentary cloud mask (all pixels below 270K discarded) applied, and the remaining pixels in each block are aggregated as a median brightness temperature. This swath of blocks from a latitude is assigned a time according to the UTC time of the image capture and the longitude of the block as a function of its offset from the

Greenwich meridian:

$$t_s = UTC_{image} + longitude \times 240(secs) \quad (3.1)$$

Each line of the original image also has a scan line time associated with it, which is the offset of the line capture relative to the nominal image time. This is important, especially when comparing areas of relative spatial proximity that may be scanned at different times in the image capture cycle. Blocks are assigned their own scan line time according to the median scan time of the captured pixels from the original image. Once this is added, each block's apparent time is rounded to the nearest minute and training data processing can commence.




Figure 3.1: Time series diagram for a swath of latitude 25.75°S–26°S from 135°E–150°E longitude on 2015 day 319 (15 November 2015). Each square represents the median temperature of the 0.25° block at the image time on the y-axis. These blocks represent one minute of training data that can be fed into the brightness temperature aggregation process.

An example of the variability in training data can be found in Figure 3.1, which is a swath of land running half the width of the Australian continent at 26°S latitude. This swath demonstrates the diurnal variation of the daytime period of each block as the image time advances. The timing of the peak temperature of the day advances from right to left as the images progress, which should be the expected behaviour as the Earth rotates. Most noticeable in this

3. A Broad-Area Method for the Diurnal Characterisation of Upwelling Medium Wave Infrared Radiation




Figure 3.2: Figure 3.1 visualised relative to local solar time. Each of the swaths of block values extracted from each image are shown as a grey line. The four coloured lines depict the trajectory of individual block temperatures at 135°E, 140°E, 145°E and 150°E as the day passes.

figure is the differing magnitude of temperature changes in the most easterly three degrees of longitude. In this region the land cover changes from the sparsely vegetated desert found in central Australia to more vegetated cropping land interspersed with forest. This composition of land cover displays far less upwelling radiation during the middle of the day than much of the rest of the swath, necessitating the step of temperature standardisation for each block to account for this reduced range of radiation values.

Figure 3.2 shows the relationship of the swath to local solar time as time passes. Each of the grey lines on this graph represents median brightness temperatures relative to local solar time for each block. The coloured lines represent the temperature trajectories of a selection of blocks from the swath. In an ideal situation these lines would sit on top of one another, but variations in land cover and partial obscuration of blocks by cloud, especially on the most easterly vector depicted in green, have prevented it in this instance. Nonetheless, the daily maximums of the other trajectories occur almost at the same point, between 12:30 and 13:00 local time. One characteristic this diagram demonstrates is the overlapping redundant information available in the ten minute time series provided by the AHI sensor. Swaths of sufficient width produce large amounts of redundant measurements at each local solar time




Figure 3.3: An example of the training data fitting process on a swath for a 24-h period. Grey lines represent the raw median values for each swath from the time of each image, the blue data points are representative of the median brightness temperature of the training data at each local solar time, and the red trend line represents the filtered medians of the training data. (Due to the nature of the filtering process, the lack of data at each end of the data results in anomalous fitting and as such the swath sampling has been extended one hour either side of the 24-h period to minimise these errors.)

depending on the imaging frequency of the satellite sensor in question.

Both Figures 3.1 and 3.2 demonstrate a basic issue with this block aggregation method, which is the variation of responses from the land surface in each block due to factors such as differences in surface emissivity and reflection, and the influence of cloud shadowing. In order to minimise the effects of this over the time series, each of the blocks is standardised (mean $\mu = 0$, standard deviation $\sigma = 1$) over a 24 hour period and these standardised values are used in the signal generation process. The data is then merged into one large time series and medians of the deviation from the daily mean of each block are taken at each minute. This information is then fed through a Butterworth low pass filter (fifth order, cutoff frequency three hours) to smooth some of the minute-to-minute variation and provide an ideal temperature curve for the day in question. The result of an example of this process is shown as the red line in Figure 3.3. This filtered data subsequently forms the Broad Area Training (BAT) data set for the fitting of a diurnal temperature curve to the target pixel.

3.2.3 Fitting to Pixel Brightness Temperatures

This study makes use of the robust matching algorithm described in [65], which adapted the SVD method described by [3] for its fitting method. In a similar fashion to the training data set, each pixel brightness temperature value in the 24 hour period to be modelled is categorised by local solar time. These times are rounded to the nearest minute to more easily correspond with the training dataset. The time offset between the first brightness temperature measurement of the pixel and the training data is calculated, and the values are extracted from the training dataset corresponding to the image times of the pixel dataset. This leaves the training dataset consisting of a number of daily vectors that correspond directly to the times that brightness temperatures were measured for the pixel's DTC.

A SVD decomposes the training data matrix A into a number of principal component vectors U which describe the training data as a series of orthogonal vectors, along with the diagonal matrix Σ which contains sorted decreasing eigenvalues for each component and the matrix V which contains coefficients for the reconstruction of the basis DTCs, as shown in Equation (3.2):

$$A = U\Sigma V^T \quad (3.2)$$

Given the datasets involved with the brightness temperature fitting process used in this paper, a significant number of the component vectors in U contribute little to the DTC fitting process. An examination of the relationship between the maximum variance in the training data and the DTC is given by the eigenvalues σ_i found in Σ . To minimise the effect of overfitting caused by these extra degrees of freedom in the component vectors, values of σ_i are tallied until the examined eigenvectors account for at least 90% of the total variance of the SVD (per [64]), with subsequent vectors discarded. This generally leads to a U matrix consisting of between two and seven vectors for fitting in this study.

For a vector of observations from a pixel e , an approximation of the DTC e^* can be created from a reconstruction of the principal components:

$$e^* = \sum_{i=1}^K c_i U_{it} \quad (3.3)$$

where K is the number of basis vectors used and c_i is a series of scalar values derived by taking the inner product of the observations from the pixel and the principal components ($e^T U$). c_i essentially describes the contribution of each of the component vectors in U to the fitted estimate of the DTC e^* .

At this point, the root mean square difference of the estimate of the DTC is compared to the raw temperature measurements. If there is a sufficient fit between the two at this point, the model answer is accepted without adjustment. If not, to minimise the effects of outliers on the robust determination of the DTC, a least squares minimisation of Equation (3.3) occurs, utilising a robust error norm which identifies outliers that contribute excessively to error:

$$\sigma(x, s) = \frac{x^2}{s + x^2} \quad (3.4)$$

x is the value of $e - e^*$ at time t , and s is a scale factor which can be reduced iteratively to improve the solution given by the error minimisation. Care should be taken when selecting suitable s values to ensure appropriate outlier elimination and that the model does not become unstable. This leads to the application of an outlier mask m , such that:

$$m_t = \begin{cases} 1, & |e_t - e_t^*| \geq \sigma/\sqrt{3} \\ 0, & \text{otherwise} \end{cases} \quad (3.5)$$

which leads to the minimisation of the error function $E(c)$ as shown below:

$$E(c) = \sum_{t=1}^N m_t \rho((e_t - (\sum_{i=1}^K c_i U_{it})), \sigma) \quad (3.6)$$

In a set of randomly selected DTCs, the number of anomalies produced by cloud far outweighs that of fire. In order to minimise the influence of sporadic clouds to the temperature fitting process and keep the modelled curve closer to higher temperatures that are more reflective of ground conditions, the weights of all negative residuals are halved during the least squares minimisation process.

This process refines the values of c_i applied to the principal components U_i in order to reconstruct an approximate value for e_t . At this point, outliers from the ideal DTC can be identified and attributed according to their characteristics—whether they are positive or negative anomalies, and whether the change occurs suddenly or gradually over time.

3.2.4 Algorithm Evaluation

3.2.4.1 Method Accuracy

In order to test the accuracy of brightness temperature modelling using BAT data, a number of comparisons are made with currently used methods for deriv-

3. A Broad-Area Method for the Diurnal Characterisation of Upwelling Medium Wave Infrared Radiation

ing brightness temperature. Focus will be placed upon the method described by [65], which gathers training data vectors from the pixel's recent history, ranging up to thirty days prior to the time period being fitted. This method uses a rejection criteria based upon the number of pixel returns in a 24-h set affected by cloud, with rejection of a daily vector based upon more than 6.25% of pixel returns in the period being cloud-affected (six of the 96 images available). A pixel can be classed as able to be modelled using this training data if at least ten days out of the previous thirty are classified as successful and only successful vectors are utilised in the SVD fitting process. A similar process is used to apply this method to pixel history in this study, with a small alteration in the number of bad pixels causing a rejection. As the AHI sensor takes a maximum of 142 images in a 24 hr period, a rejection of a daily vector occurs when there are more than nine instances of clear sky probability equal to zero in the corresponding AHI cloud mask data.

For comparison, a temperature fitting has been applied to each sampled pixel utilising BAT data derived from the previous thirty days and the previous ten days prior to the fitting. These sampling points simulate the minimum and maximum amount of training data available to the pixel-based fitting approach. Comparison is also made to contextual based temperature derivation based upon a similar method used for the MODIS fire products [41], with an expanding kernel based upon 25% cloud and fire-free returns in the target vicinity, with a minimum of six successful returns. In the case of the MODIS products, the kernel is allowed to grow to 21×21 pixels in size before reaching a failed state. Due to the larger areas covered by AHI pixels, the divergence of land cover and geology over these distances could be extreme, so a limit of 5×5 pixels was proposed instead. For this evaluation, a set of 19,916 pixels were randomly selected over the Australian continent for the month of November 2015, with fitting of the raw brightness temperature performed using the four different methods noted above.

3.2.4.2 Method Efficiency

The nature of the BAT method lends itself to pre-processing—once a swath has processed from a particular image time, the data is utilised in any fitting process on any pixel in the swath for the encompassing period (either ten or thirty days in the case of this study). For each image, an ocean mask and basic cloud mask (such as described in Section 3.2.2) is applied before blocks are aggregated by median temperature. This data is stored in files for later access in the block-based fitting method.

For evaluation of the efficiency of both algorithms, the Python code for both

methods was written to take advantage of an initial common data gathering process (location, start time, daily brightness temperature, available cloud). A random location and time in Australia during the month of November 2015 was selected, with the pixel-based training data calculated first. An extraction routine for training data was run once the suitability of an individual day of data was assessed. This routine would then check whether sufficient training periods for a pixel existed, and ran the fitting routine in the event of success. The time taken for both successful and unsuccessful routines using pixel-based training data was recorded.

Next, the BAT fitting was undertaken using thirty training days and ten training days, utilising the pre-processed block data, with the processing time for each recorded. Processing took place utilising code written for Python 2.7, using the Numpy (v1.12) mathematical processing module and the built-in multiprocessing module, along with the python-netCDF4 (v1.2.1) module for data reading and writing. The computer used for processing utilises an Intel(R) Core(TM) i7-5820K CPU, allowing 12 threads of processing to take place at once on the six available cores. Images and pre-processed training data were stored on and accessed from an external hard disk drive using USB-3 for data transfer.

3.2.4.3 *Training Data Availability*

In order to calculate the availability of pixel-based training data as per the [65] method, we used the AHI cloud mask data to evaluate cloud cover at the time of each image, with all values of clear sky probability greater than zero counting as clear sky. The data was aggregated for each pixel for a 24 hour period corresponding with a UTC day, and pixels with greater than 132 cloud-free returns judged to be suitable for training purposes. Ten or more such successes within the thirty day period evaluated made the pixel suitable for pixel based training data collection. As the BAT datasets rely on redundant measurements from adjacent blocks on the swath—with a consequent reduction on reliance upon an individual block for training data—it was difficult to provide a direct comparison case back to the pixel-derived data for evaluating availability. Absences in measurements from a specific block lead to reductions in the accuracy of the training data model for the at-large swath, but complete loss of model performance only occurs where a significant length of an individual swath is completely obscured by cloud for an extended period, which is unlikely outside of situations where the swath is very narrow. For a comparison in this case, we used the availability of pixel training data on a per-pixel basis and the availability of BAT data on a per-block basis for the months of October and November

3. A Broad-Area Method for the Diurnal Characterisation of Upwelling Medium Wave Infrared Radiation

2015 for the Australian continent. The BAT availability figures in this case may be construed as a minimum, and effective training data availability is generally much higher.

3.3 Results

3.3.1 Algorithm Accuracy

For comparison purposes between the pixel-based and BAT methods, we processed solutions for a selection of 19,916 pixels with random locations on the Australian continent with fitting times during the month of November 2015. Of this larger selection, 5747 pixels (approximately 28.9%) had sufficient training data available for use of pixel-based fitting.

Table 3.2: Comparison of fitting techniques to brightness temperatures recorded by the AHI sensor using root mean square error after eliminating incidences of Clear Sky Probability (CSP) of less than one from the evaluation.

Fitting technique Incidences of CSP < 1	RMS Error (K)				
	≤ 10	11 – 30	31 – 50	51 – 70	> 70
Pixel-based training	0.78	1.01	2.28	3.25	10.40
BAT (30 days)	0.94	0.94	1.11	1.48	4.19
BAT (10 days)	1.15	1.21	1.40	2.10	6.31
Contextual temperature	0.33	0.42	0.41	0.40	0.42
<i>Number of samples</i>	<i>903</i>	<i>741</i>	<i>768</i>	<i>851</i>	<i>2345</i>

Table 3.2 shows the relative accuracy of each of the temperature fitting processes in comparison to the raw brightness temperatures recorded by the AHI sensor after measurements flagged as cloud are eliminated. In clear sky conditions (≤ 10 cloud instances) the pixel-based training method performed more accurately than the BAT method for either ten or thirty days of training data, with differences of 0.30 and 0.15K respectively. This was expected, as the data derived from an individual pixel will perform far better with regard to localised effects on pixels such as land cover composition. As the number of cloud instances increased, the pixel-based training accuracy degraded, but the 30 day BAT method showed a steady relationship with the cloud free brightness temperatures up to 30 cloud affected returns, with a more gradual loss of accuracy on heavily cloud affected days. The 10 day BAT method showed a similar performance dynamic, with an approximate 20% loss in accuracy up to the point where accurate temperature determination became difficult (> 50 clouds).




Figure 3.4: Examples of model fitting using the four training data derivations. Figure (a) shows a typical day with less than ten cloud instances, (b) shows a day with between 20–30 instances of cloud, and (c) is typical of a day with more than 70 identified cloud periods.

The context-based temperature derivation appears to perform especially well regardless of the number of outliers encountered in the temperature dataset. This is mainly due to the spatial autocorrelation of cloud and weather effects, which see contextual based temperatures closely track anomalous temperature as measured in each pixel, rather than tracking the base model of upwelling radiation without anomaly. Comparison is also inhibited by the instance of a contextual temperature not being available, which is more likely at times where anomalies would be detected.

Figure 3.4 demonstrates some examples of the fitting process over a few of the random pixels selected. Figure 3.4a shows the typical performance of each temperature estimation technique on a clear sky day, with the fitted curves from the BAT showing greatest variance from the raw temperatures during the night-time period and at the peak of the day. The pixel based training data is better able to handle tracking of brightness temperature during the daytime, but also shows some variance during the night to a lesser extent. Figure 3.4b shows an example of a day with between 20 and 30 cloud instances. This figure shows the fitting curves from both of the BAT datasets effectively ignoring the cloud induced anomalies in the early evening (0900 UTC) and during the peak of the following day (0200 UTC). This is in contrast to the pixel-based trained fitting, which follows the anomalies in the early evening more closely and underestimates the rising temperatures in the morning period (2000–0000 UTC). This figure also demonstrates the spatial autocorrelation issue with the contextually derived background temperatures. The contextual fitting follows the cloud affected temperatures in the afternoon and early evening period, ignoring the significance of the temperature peaks during this period which are more indicative of the true background temperature.

Figure 3.4c shows an example with around 95 instances of non-clear sky conditions. This example shows some of the fragility inherent in all of these methods when a significant portion of the DTC is cloud obscured. Due to the standardisation process that all pixels undertake before fitting of the training data is undertaken, the large temperature drops caused by the evening clouds (0700–1100) skew the normalisation by lowering the mean and increasing the standard deviation. This issue sees both the pixel-based fitting and the BAT fittings underestimate the background temperature throughout the course of the day, with the largest variations in the peak of the day and towards the end of the fitting in the early morning. This also demonstrates one of the limitations with using a model with greater degrees of freedom. Extra degrees of freedom are necessary in the case of sudden drops in the background temperature caused by weather effects, but on days where brightness temperatures do not follow the indicative shape of the DTC, the SVD fitting process tends to produce curves that wander about with little relation to the true background temperature.

3.3.2 Algorithm Efficiency

Table 3.3 shows the average time taken to perform the computations by each of the temperature fitting methods from a processed pool of 19,916 target pixels. The time taken to provide a temperature fitting based on BAT data for a thirty day period is approximately one-tenth of the time taken to derive a similar fitting from the pixel data. Part of the increase in speed comes from pre-processing of the pixel training data, which takes approximately 15 minutes per day of images. The difference between the pixel-based and block-based methods coincide at around 170 pixels evaluated, with the block-based method performing significantly faster beyond this. The pre-processed blocks can also be used in calculations for training data in adjacent time periods ± 10 or 30 days from the start time of the original assessment, speeding up subsequent processing further. The largest improvement in time of processing between the two methods is with regard to acquiring data about the state of cloud in the training data period, where evaluating the suitability of 24-h periods using the cloud mask data makes up over 75% of the time taken to process a solution using this method.

The time taken to calculate a fitting from the 10 day BAT data is even shorter, taking around 40% as long as the thirty day BAT method. A significant improvement in processing time can be achieved for a small decrease in accuracy compared to the fitting provided by the 30 day BAT data. This is of importance in situations where processing time may be more critical than

Table 3.3: Comparison of time taken to provide a temperature fitting using the pixel-based training technique and the BAT fitting process.

Fitting Technique		Time Taken (secs)
BAT	<i>Pre-processing</i>	<i>15 mins/day</i>
	BAT (30 days)	18.4
	BAT (10 days)	7.3
Pixel-based	Sampling cloud statistics	141.8
	Training & fitting	35.5
	Total time	177.2

accuracy.

3.3.3 Training Data Availability

Training data blocks were derived for the entirety of the Australian continent for the months of October and November 2015, from 4379 and 4245 AHI Band 7 images respectively. Block method training data was deemed available when at least one pixel brightness temperature in a 0.25° block was observed above 270 K. Using the BAT method training data availability for the continent averaged 96.9% of all images for October and 95.4% of all images for November. This compared to the pixel-based method, which yielded 70.2% of data in October and 67.8% of data in November. Given the assumption that a minimum of ten days of training data is sufficient to derive a DTC, 91.4% of blocks demonstrate a sufficient number of clear days to derive training data in October, and 81.8% in November. In comparison the pixel based training method supplies sufficient days of training data to a pixel in 43.8% of cases in October and 37% of cases in November. From an availability standpoint, the block based method of training data derivation is a marked improvement over pixel-based methods, and a significant amount of redundant data on each line of blocks extends this availability further.

Figure 3.5 shows the spatial distribution of training data availability over the study area for October and November. The figure demonstrates the overall increase in training data availability when using the BAT method. Areas of limited availability demonstrated similarity in spatial distribution between the two methods, with south eastern Australia and Tasmania in particular suffering from a lack (indeed in some areas a total absence) of pixel-based training data for the evaluation period. In comparison, the majority of blocks retained availability for the required time during the period, with some deterioration in coastal areas. The BAT method demonstrates a significant improvement in temperature fitting ability, especially in the highly populated coastal areas in

3. A Broad-Area Method for the Diurnal Characterisation of Upwelling Medium Wave Infrared Radiation

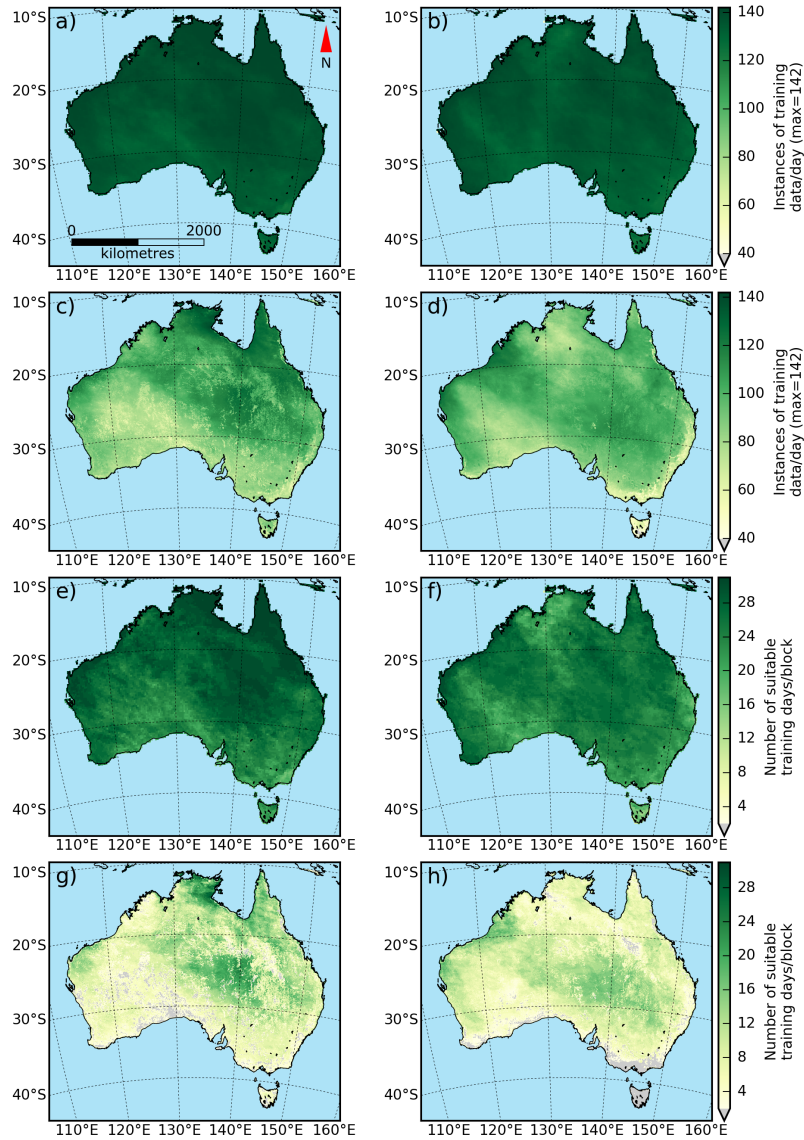


Figure 3.5: Availability of training data from the block and pixel based methods. (a,b) show the mean instances of training data available using the BAT method for October and November respectively; (c,d) show the training data available using the pixel method for the same months. (e,f) demonstrate the number of 24-h periods that could be utilised as training data for each block in October and November, and (g,h) show this same criteria using the pixel based method.

the south and east of Australia.

3.4 Discussion

The advantages of using the BAT method for deriving training data for temperature fitting include robustness against localised cloud, especially in areas with persistent standing cloud such as coasts and mountainous areas; the ability to minimise the number of training days required for deriving a brightness temperature fitting due to the increased availability of training data; and a reduction in the storage of data and processing time of training data for temperature fitting. Whilst in this study the training data is used to feed an SVD fitting of the DTC, the data could easily be applied to other fitting techniques, such as a random forest classifier or as a state vector for Kalman filtering. The nature of the fitting process removes the need for tracking locations that have standing hotspots, as the fitting process is completely context independent, and eliminates errors that may be caused by large variations in response to solar radiation and emission between adjacent pixels. As the broad area method relies on as few as one cloud-free pixel per block from which to derive a median temperature, the method is far more robust in response to occlusion than the pixel-based method. When banks of cloud associated with large weather systems are present, a single block may be totally covered by cloud in one or several images. However, the redundancy associated with evaluation on a continental scale means that this lack of data has a minimal effect on the training data for the same time period.

The quality of the training data used in the process relies on a couple of factors. For example, the width of the longitudinal swath at any given latitude is affected by the amount of ground available to sample that swath. In Australia, latitudes between 25°S–30°S have the full width of the continent to sample temperatures from, with anywhere up to 160 blocks of data per image. This provides a large amount of redundancy in the training data, reducing the affect of outliers in the training set. In contrast, a smaller swath width results in fewer blocks to formulate the training data, with a greater risk of anomalous returns affecting the resultant fitting processes. A smaller swath width also increases the influence of edge cases such as coastlines, which tend to moderate land surface temperature variance if they are not handled by an adequate ocean mask. A buffer of two pixels (between 3 and 6 km) was used to eliminate these edge cases for ocean boundaries, but water bodies such as lakes and reservoirs may also contribute to erroneous training data if the number of blocks used for training data is low. Discontinuous and inadequate areas to derive training data from may prove a challenge, and evaluating the performance of the BAT method in a region like Indonesia where cloud cover is high and land areas are discontinuous would properly test the limits of the method.

3. A Broad-Area Method for the Diurnal Characterisation of Upwelling Medium Wave Infrared Radiation

Deriving training data using the BAT method is not without its issues. The method involves the use of a median value for the entirety of a $0.25^\circ \times 0.25^\circ$ block without taking into account factors such as land cover, land use type, slope and aspect, and surface emissivity, all variables that can vary significantly between pixels in a block. The application of the training data back to the pixel brightness temperature using the SVD process also omits consideration of differing albedo between adjacent pixels in a block. These issues may both be resolved by using a weight from each pixel at different times of the day to take into account the differing emissivity and reflection and how they affect the DTC of each location. The training dataset also has a minimum cutoff temperature of 270 K to minimise the influence of cloud affecting the training data in lieu of an operational cloud mask which may produce large areas of missing data in places where surface temperature and reflection components sit under this value for long periods of time. If a method such as this is used to track surface upwelling radiation in areas that have sustained brightness temperatures below 270 K, a cloud mask could be substituted in this case to eliminate major outlying temperatures instead. The effect of snow on brightness temperature tracking using this method has not been explored, mainly due to the study site and time of year chosen.

Considering the evaluation of the various methods for accuracy of fitting, the block-based training data based upon thirty previous days performed reasonably well for fitting accuracy in comparison to the pixel-based training data for all anomaly classifications, and the results derived from the 10 day block training data were of similar accuracy. If a threshold for suitability of fitting is placed on the results, such as would be for a fire detection model, the 30 day block-based method is far more robust with respect to anomalies and would continue tracking the expected upwelling radiation with up to one quarter of the time series obscured. A full sensitivity analysis using in-situ upwelling radiation data along with a verified cloud mask would be of value to confirm the shape of the diurnal model in order to provide an independent confirmation of the method's accuracy. Of note here is the apparently extremely good performance of the contextual method for providing background temperature, which is an unfortunate side effect of methods for evaluating the accuracy of the model fitting process. The high spatial autocorrelation demonstrated by upwelling radiation results in a contextual based temperature tracking extremely closely to the raw temperature measurements regardless of whether the surrounds themselves are affected by anomalies such as cloud, smoke or fire.

Whilst the accuracy of the fitting method used in this study for days of limited thermal anomalies is quite good, the accuracy of fitting provided by the

SVD process breaks down once a large number of anomalies are encountered in the fitted data. This behaviour, which manifests in the type of “wandering” curves displayed in Figure 3.4c, is caused in part by the standardisation process applied to pixel brightness temperatures, where significant numbers of cloud incidences, especially from thick cold clouds, can act to drop the mean and increase the standard deviation of temperatures across the fitted period. This affects the initial estimation given from the SVD process to the point where the outlier elimination process disregards correct temperature measurements by mistake. One of the limitations of using the SVD fitting method is that outlying measurements in the raw brightness temperatures cannot be easily eliminated from the function evaluation. Whilst applying a more rigorous cloud mask or deriving a standard model for brightness temperatures at a given location and date could be ways of eliminating this mainly low temperature biasing; investigation of other methods for applying the training data for temperature fitting should be a priority.

Overall, the BAT method described in this paper performed adequately from an accuracy standpoint, but the real benefits of the method lie in the improvements in processing time and availability. The BAT method processed an individual fitting at about ten times the speed of the pixel-based training data method, mostly due to the lack of need for cloud mask evaluation of the training data vectors prior to fitting. This issue with the pixel-based training method could be alleviated somewhat by smaller file sizes, as the major issue with processing of the pixel training seems to be the bulkiness of cloud data produced for this sensor (typical file sizes for the AHI cloud product are approximately 90–100 MB). Given the ten day BAT method performed similarly to the pixel-based training method for temperature accuracy, using this data set instead of the thirty day block training set could be justified in situations where processing time of large image sets is of greater importance than extreme accuracy, especially when used for initial anomaly detection purposes.

From a fitting availability standpoint, the BAT method significantly increases the distribution of areas that are able to have a temperature fitting applied to them. In comparison to pixel-based methods, the increase of fitting availability is especially marked in areas such as the east and south east coasts and the island of Tasmania. These areas are heavily populated in comparison to much of the Australian continent, and are at significant risk from rapidly changing events, such as fire and flood. The extended utility and application of the BAT method will be of great interest to land management authorities in these areas.

The BAT method formulated in this paper is designed specifically for the process of providing data to inform fitting processes for positive thermal anomaly detection, such as for fire detection. An enhanced understanding of back-

ground temperature behaviour in the MWIR space could lead to improvements in the determination of fire detection thresholds, potentially leading to delineation of fire thresholds using time of day, latitude and solar aspect, along with a greater understanding of how the mix of solar reflection and thermal emission in a pixel contributes to the minimum detectable characteristics of a fire from a particular sensor. Applications of this technique could also look to provide ongoing monitoring of fires using metrics such as area, temperature and fire radiative power utilising the improved estimation of this background temperature.

In addition to improvements in the fire detection space, this method could have applications in a number of other fields that require change detection over a short period of time. Likely applications could see aggregation of data based upon land cover classification along with local solar time and latitude to provide a baseline for mapping soil dryness changes, or for tracking the spread and severity of phenomena such as flooding and volcanic activity from geostationary imagery.

3.5 Conclusions

This study demonstrates the formulation of a broad-area training data derivation method for temperature fitting, for estimation of the background temperature of a pixel measured by a geostationary sensor whilst obscured by cloud, smoke or fire. In comparison to pixel-based methods of deriving training data for temperature fitting, the BAT improves fitting resilience during periods of light to medium obscuration, with fitting errors reduced by more than 50% on days with between 31 and 70 obscured instances using thirty days of training data. The BAT method also demonstrates significant improvement in processing times compared to pixel-based training methods, especially when dealing with a large number of fittings over similar timeframes. The use of the BAT method also increases the availability of training data for fitting purposes, with fitting availability increasing to between 80% and 90% for the time period covered by this study, compared to approximately 40% availability from pixel-based methods. This improved availability of training data will assist in tasks related to the accurate understanding of background upwelling radiation, such as in fire detection and monitoring.

3.6 Thesis Context

This chapter introduced the Broad Area Training method of diurnal fitting for brightness temperature estimation. The work, which was a variant of a method implemented by [65], demonstrated comparisons in coverage and availability of estimations in comparison to that method of diurnal fitting. The next chapter takes these model fitting estimates of brightness temperature and applies them to a theoretical fire detection implementation, with comparisons to other commonly used fire products.

Advances in Active Fire Detection Using a Multi-Temporal Method for Next-Generation Geostationary Satellite Data

4

4.1 Introduction

Wildfires are a major catalyst for change in the environment, creating essential environmental disturbance to assist natural vegetation cycles, and disruptive change with intense fires causing damage to both property and the environment [80]. With the increased propensity towards more intense fires due to the increased drying of fuels and more extreme weather conditions [51], it is of vital importance to be able to detect wildfires as early as possible, in order to drive potential mitigation strategies, minimise the impact of fires on strategic assets in the built environment and maximise the time available to inform the public of impending fire activity. Using satellite sensors for the detection of fire has proven an effective way of producing information about the likelihood of active fire, along with providing metrics such as the size and intensity of fire activity [68].

In order to obtain timely information about fire ignition, the use of geostationary satellites to provide imagery takes primacy over low earth orbiting (LEO) active fire (AF) products, such as those obtained from the Moderate Resolution Imaging Spectroradiometer (MODIS) and the Visible Infrared Imaging Radiometer Suite (VIIRS) sensor systems [18]. Whilst the LEO sensors have advantages over geostationary satellites, such as the reduced sensor instantaneous field-of-view (IFOV) from their far lower orbits which makes smaller fires easier to detect, the LEO sensors are restricted in their ability to provide rapid temporal updates of active fire information due to the limited number of sensor views of a target area in a 24 hour period [45]. This limit to revisit time makes LEO fire products particularly vulnerable to errors of omission when cloud covers an active fire during the sensor overpass, and when the ephemeral nature of fire causes a drop in fire activity coincident with the image time. These vulnerabilities in fire detection are offset by the use of geostationary sensors, which can use their superior revisit time to sample fire activity during gaps in cloud cover and smoke, and to capture the changing state of an active fire more effectively [66]. Geostationary satellites to this point in time have been effective in the detection of larger area fires, but have suffered from major omission errors when fire sizes are small, due to the lower spatial and radiometric resolution of the geostationary sensors involved [45]. In order to capture fire ignition time more accurately, a focus on improving the capability of detecting these smaller fires is required.

With the launch of new geostationary satellites, with sensors such as the

This chapter was published in a peer-reviewed journal as: Hally, B., Wallace, L., Reinke, K., Jones, S., & Skidmore, A. (2018). Advances in active fire detection using a multi-temporal method for next-generation geostationary satellite data. *International Journal of Digital Earth*, 1–16. <https://doi.org/10.1080/17538947.2018.1497099>

Advanced Himawari Imager (AHI), the gap in active fire attribution capability between geostationary sensors and the LEO sensors is decreasing [45]. The AHI sensor, launched onboard the Himawari-8 satellite in 2014 by the Japan Meteorological Agency (JMA), provides enhanced capabilities for fire detection in comparison with previous geostationary sensors, with 2km spatial resolution at nadir and an increased 14-bit quantization in the sensor's 3.9 μm band 7 in comparison to previous sensors in the area. These enhancements make the AHI sensor theoretically capable of detecting far smaller fires, both in absolute size and in fire radiative power, and these improvements have been demonstrated in initial studies based upon AHI data, such as those undertaken in [94], [88] and [31]. The real advantage of using geostationary sensors for fire detection is the increased temporal resolution in comparison to LEO sensors. Whilst a satellite constellation such as MODIS may record up to four images of a location per day, sensors such as the AHI are able to record at least 144 images a day over the full disk visible from the sensor location, with much higher temporal resolution in AHI's case over areas such as Japan. This higher revisit rate of imagery enables opportunities not just to provide quasi real-time detection capabilities, but also allows for us to gain a better understanding of the characteristic nature of the earth's reflectance and emittance in the all-important 4 μm region of the electromagnetic spectrum.

Of vital interest to any attempt to detect and quantify fire activity in a location is an accurate estimation of the location's background temperature in the absence of fire. The difference in the background and elevated pixel temperature is the basis of most hotspot algorithms, starting with the initial work on satellite fire detection by [11], which lead to commonly used active fire products such as the MODIS fire product [26] and the VIIRS active fire product [75] used at LEO scale, and the WF-ABBA fire product [45] and the LSA SAF Meteosat products [89] which utilise geostationary data. These methods had their genesis in the evaluation of individual images for fire activity – a necessary evil given the temporal gap between revisits using LEO sensors – and looked to provide the background temperature of suspected fire pixels from the context of the surrounding non-flaming area. This method of background temperature determination relies on accurate differentiation of cloud and smoke-affected pixels from valid occlusion-free pixels, along with elimination of adjacent fire-affected pixels and adjustments for a sensor's point spread function. In difficult detection conditions with large amounts of occlusion, the area required for obtaining sufficient clear pixels to facilitate this estimation process expands significantly. The sampled areas surrounding a target pixel may bear little resemblance to the emissive and reflective properties of the central pixel, due to factors such as landscape cover, relief and weather effects [20, 77]. Context-

tually derived temperature can also suffer from spatial autocorrelation issues, where the omission of thin cloud or smoke reduces brightness temperatures across wider areas and reduces the temperature contrast required to detect thermal anomalies in the first instance. These issues with contextual determination are exacerbated by the sensitivity of fire algorithms to accurately determined background temperature. [25] highlighted some of the difficulties with contextual temperature determination methods, including the estimate of background temperature from the surrounding context being higher than the target pixel's flaming temperature, which occurred in 22% of cases with the MODIS fire product between 2003-2012. Issues also arise when the detected fire temperature is marginally above the estimated background temperature, with attribution errors in background temperature of 1 K causing errors in the estimated fire area of a factor of 100 [25], though these types of errors are less likely to affect initial fire detection.

These issues with applying contextual temperature derivations in difficult conditions have led to investigation of the estimation of background temperature using temporal-based methods, which focus upon the estimation of background temperature over a time period leading up to pixel ignition. This estimation, based upon the diurnal variation of a pixel's temperature, attempts to leverage the higher temporal resolution of sensors to provide an accurate background temperature, based upon the expected behaviour of the pixel when no anomaly occurs. Initial work evaluated the use of absolute descriptive models [27, 38] and Reproducing Kernel Hilbert Space models [83], which both suffered from the inability of the model to allow for adjustments due to emissivity change, for instance after rainfall. Kalman filtering methods for temperature estimation were first utilised by [85] and were refined by [65], and these generally provided good results for fitting the diurnal temperature cycle whilst accounting for gaps and anomalous data in the pixel record. The work by [65] in particular showed promise for obtaining more accurate background temperature data, and this method utilises training data from the target pixel for a 30-day period prior to the day being modelled. Training data derived in this way for temperature fitting is vulnerable to gaps in the data - in areas of persistent standing cloud or periodic obscuration, sufficient training data for a pixel fitting may not be obtainable, and in this case estimates are provided by pixels from an anomaly-free library of pixels. [31] investigated this method and found that at in Australia only approximately 40% of pixels had sufficient training data available to utilise the pixel-based training data method. This study also found significant impediments in the form of processing time required to produce fittings, which was also mentioned in the original study by [65].

Recent work in the field of background temperature estimation has investi-

gated the use of a broad-area time-series approach for deriving training data for temperature fitting, as described in [31]. Raw image data from the AHI-8 sensor was aggregated based upon median temperature values of the land surface at 0.25° by 0.25° intervals for a period of 30 days prior to the fitting start time. This dataset was then aggregated by local solar time and medians were again taken in order to create an idealised diurnal temperature cycle for a strip of latitude at one minute intervals. This model was then sub-sampled for the corresponding values at the times where brightness temperature was available for the location in question. A fitting process was then applied in order to provide an optimised background temperature, based upon the training data and the brightness temperature data available from the location for the corresponding 24 hour period. This method produced similar or improved accuracy of background temperature determination when compared to pixel-based methods such as those in [65], with a significant increase in training data availability over all areas examined.

The increased availability and accuracy of the background temperature determination method described by [31] provides a starting point for a fire detection algorithm. The robustness of the model suggests that the periods of negative temperature anomalies such as cloud and smoke, and positive temperature anomalies such as fire activity, have a reduced effect on the subsequent fitted temperatures in comparison to the ideal diurnal temperature cycle that should be demonstrated by the pixels being examined. This paper proposes the use of this fitting technique with a series of temperature thresholds for thermal anomaly attribution, with a view to examine the abilities and limitations of the temperature fitting method with regard to fire detection. The paper will examine the performance of this threshold algorithm against commonly used LEO products such as the VIIRS and MODIS active fire products, in order to determine the completeness of anomaly capture of the various methods, and also to examine the capability of the AHI sensor to detect fire in its early stages and give an estimate of the expected improvements in fire detection times using this method.

4.2 Data

4.2.1 Sensor data

Geostationary satellite sensor information for this study was obtained from Japan Meteorological Agency via the Australian Bureau of Meteorology (ABOM). The study makes use of Band 7 ($3.9\mu\text{m}$ wavelength) images from the AHI-8 sensor for the month of August 2016, which were corrected both radiometri-

cally and geographically according to the procedures outlined in [60]. Data for the month was preprocessed according to the method described in [30] in order to provide training data for the temperature fitting process. To determine the availability of cloud-free pixels for false attribution evaluation, a clear sky probability product based upon the CLAVR-x cloud mask algorithm [35] with some minor alterations [54] was used.

The study considers two commonly used active fire products from low earth orbiting sensors; the MODIS Collection 6 (MOD/MYD14) 1 km active fire product as outlined in [26], and the VIIRS 375 m (VNP14IMG) active fire product as described in [75], making use in both cases of the geographic position of the detected hotspots and the time of satellite overpass. Both of these products show low commission rates for fire activity in areas with similar geographic and land cover characteristics, especially with the larger fires this study area tends to produce. In addition, to determine the extent of the study area affected by fire and to provide a method against which to explore fire disturbance, a MODIS burned area 0.0025° product [53], produced by the TERN Auscover project, was used. This dataset describes the area affected by fire and attributes areas with the date and time of the first recorded disturbance in the month in question. This product is used for this study primarily because of its effectiveness at accurately describing fire-induced change in geographic regions similar to those used in this study [53].

4.2.2 Study Area

The area selected for the study is a section of north western Australia bounded between 15°S and 20°S in latitude, and 125°E and 130°E longitude as shown in Figure 4.1, during the month of August 2016. The area is predominantly comprised of savanna woodland, tending to drier hummock grass conditions in the southern portion [71]. The time of year selected corresponds to Mid Dry Season type fire behaviour [12], with a high number of fires and larger fires able to take hold due to reduced fuel moisture and humidity. The low prevalence of cloud cover in the area compared to other periods of the year provides ideal conditions for any fire detection methods used. This lack of obscuration, coupled with the remoteness of examined fire activity from population centres, which mean that fires continue to spread naturally after initial detection, allows for an ideal test-bed for comparison between various fire products. The study area itself comprises of 64,374 unique AHI-8 pixels covering approximately 295,000 km².

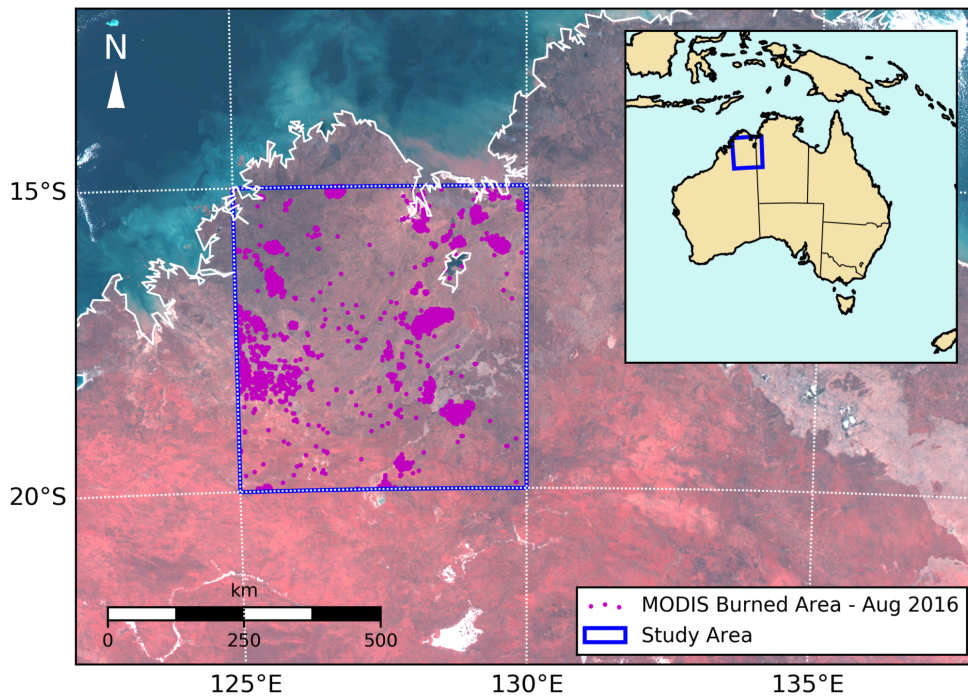


Figure 4.1: Diagram showing the location of the selected study area in north Western Australia, along with the areas affected by fire between 1st - 28th of August 2016 as determined by the MODIS burned area product [53].

4.3 Method

4.3.1 Broad Area Training method

This study utilises previous work conducted upon the AHI sensor in the application of a multi-temporal method of background temperature estimation, known as the Broad Area Training (BAT) method [31]. This method involves a two-step process for geostationary data - a preprocessing step, where AHI Band 7 images are aggregated based upon the cloud-free median of data at the 0.25° by 0.25° spatial scale; and then a fitting step, where this spatially aggregated data is stacked temporally using local solar time, and a standardised diurnal model is produced at 0.25° by 0.25° scale for use in individual pixel fitting using a single value decomposition (SVD) process. These fittings at pixel level can then be compared to the raw brightness temperature data as measured by the satellite sensor, to identify thermal anomalies such as those caused by active fire.

4.3.2 Fire Detection and Threshold Selection

In order to isolate positive thermal anomalies, basic thresholds above the temperature estimate of a location are used. Taking into account the relatively young age of the AHI sensor and the minimal existing published work on its fire detection capabilities [59, 94], the nature of the sensor's response to fire is largely unassessed. The findings of [31] describing the error involved in fitting temperatures to clear-sky scenarios, and results from [66] using the SEVIRI sensor, provide a potential starting point. These studies suggest a lower bound of 2 K as a starting point for threshold examination, with a higher bound of 5 K based upon an assessment of the response of known fires in comparison to low earth orbiting active fire products, with supplementary values selected in between at integer level. Whilst by no means being a comprehensive set of values for evaluation, the range of thresholds evaluated seeks to gain an understanding of the relationship between the commission and omission error associated with using a thresholding method of this type.

4.3.3 Anomaly Detection Rate Evaluation

To evaluate the validity of using assorted thresholds for thermal anomaly detection, an analysis of the difference between recorded and modelled temperatures was conducted for a 28 day period (2016, Julian days 214 - 241) during August 2016 over the study area (see Figure 4.1). The study area was broken into 400 blocks, each $0.25^\circ \times 0.25^\circ$ in size, which typically comprise of between 150 - 170 AHI pixels in this part of the full disk. Blocks were then selected at random and coupled with a date and time selected from the study period. Dates of analysis for each block were randomised, whilst the hour at which fitting commenced was divided evenly between samples, to minimise the effect of time of day on fitting accuracy evaluation. From the 400 blocks and 672 unique starting times (24 hours x 28 days) available, a total of 960 block samples were selected for evaluation. Training data for the pixels in each block was derived using the BAT method, and temperature fitting was performed on the brightness temperature data for all pixels in the block using a single value decomposition method.

In order to ensure that the effect of cloud on the temperature fitting was minimised for the set of results used for threshold determination, the clear sky product described in Section 4.2.1 was evaluated for pixels to be fitted. Pixel diurnal vectors were discarded when less than 133 instances of clear sky occurred in the day examined from the 142 daily images. As the experimental AHI clear sky product used had significant issues with false attribution of cloud cover during the night, a temporal mask was applied to its use. The clear sky

product was used for cloud cover determination during daylight hours (approx. 2300 – 1000 UTC in this case), whilst at night cloud cover was flagged by use of a rudimentary mask of brightness temperature less than 280 K. Once a specific pixel diurnal vector has passed the cloud requirement, the differences between the modelled and the measured brightness temperatures were recorded for analysis.

As a further analysis to determine the rate of detections associated with actual fire activity, the AHI fitting dataset was also compared to burned area detections from the MODIS burned area product [53] for location and time. If an AHI detection was observed, the burned area product was checked for a recording for a period of 48 h either side of the AHI detection, and active fire was attributed as the cause of the detection if this criteria was met. This large evaluation window was selected to ensure that fire-induced detections were not missed due to cloud cover blocking one or more of the daytime imaging passes of the MODIS satellite.

4.3.4 Active Fire Product Intercomparison

The MODIS burned area product described by [53] was used to determine instances of fire-induced change in the study area. The positional information in this dataset was converted to AHI image coordinates, and any duplicate returns in both location and time were eliminated. Once these locations were identified, a similar analysis of the MODIS and VIIRS active fire products was undertaken. Hotspots identified by each of the products had their image time and equivalent AHI locations recorded. Incidences where active fire products recorded a hotspot detection in the AHI pixel area in the 23 hours preceding the burned area disturbance were noted. This led to the evaluation of four intercomparison types – (1) fires which were recorded by the burned area product, but not by either of the active fire products; (2) fires which were detected by the VIIRS active fire product and the burned area product, but not by the MODIS active fire product; (3) fires that were identified by the MODIS active fire product and the burned area product, but not by VIIRS; and (4) fires discovered by all three of the LEO products used.

From these intercomparison classes, a random sample of 150 incidents were selected for further analysis. For each of the selected instances, a temperature fitting window was selected based upon the time of the MODIS burned area product. The temporal window for the diurnal fitting of pixel background temperature was fixed beginning 23 h prior to the detection of change from the burned area product (Figure 4.2), in order to maximise the temporal coverage of fitting prior to the burned area detection. Training data for each of the

4. Advances in Active Fire Detection Using a Multi-Temporal Method for Next-Generation Geostationary Satellite Data

pixel fittings was derived using the BAT method [31], and temperature fittings applied to each of the raw brightness temperature sets using a single value decomposition method. The thresholds described in Section 4.3.2 were then applied to the difference between the fitted background temperature and the recorded brightness temperature, and instances where these thresholds are exceeded are recorded with the time of AHI detection noted. Success of the threshold detection algorithm was considered as at least one anomaly being detected during the 24 hour period of fitting, with synchronous detections being classified as instances where an AHI anomaly occurs within 20 minutes of any active fire product detection in the pixel in question.




Figure 4.2: An example illustrating the shifted temporal window described in Section 4.3.5 in operation. In this instance, a significant portion of the original fitting window, which is the time period to the right of the grey vertical line, is affected by active fire in the pixel, and this results in a higher fitted background temperature than expected during the period between 2016-218 09:00 and 2016-218 18:00, as shown by the line in dark green. By moving the start of the fitting earlier, based upon detection by the VIIRS active fire product in this case, the fitting process is less affected by the elevated temperatures associated with the ongoing fire. This produces a fitting such as that shown by the light green line, leading to fire-related anomalies being identified earlier, dependent upon the threshold set.

4.3.5 Determination of Fire Ignition Time

The time required for fire activity to grow from ignition to being visible via remote sensing in the study area is minimal, given the general lack of tree cover in the landscape, fuel moisture and weather conditions [80], and the theoretical minimum sampling size of the sensors [75]. With fires growing to dimensions

visible from geostationary sensors in minutes, the sporadic temporal nature of LEO image capture means a significant portion of fires manifest on geostationary imagery prior to detection via LEO algorithms. As highlighted by [31], the accuracy of the SVD temperature fitting method decreases as the number of temperature anomalies increases. This poses a problem with regard to the use of a low earth orbiting detection product such as the MODIS fire product as a reference data set - if the point in time at which the target anomaly can be detected by the geostationary sensor precedes the LEO sensor overpass by a significant amount of time, the subsequent temperature fitting of the geostationary pixel data will be skewed by this period of artificially elevated temperatures, resulting in an over-estimation of the background temperature over the 24 hour fitting period. Depending on the size of the initially detectable fire, this could lead to errors in the estimation of the initial ignition time, and could also lead to a pixel being misattributed as non-fire during flaming periods if the diurnal signal is sufficiently disturbed.

In order to mitigate the effects of ongoing fire activity on the background temperature estimation process, an adjusted temporal window for the temperature fitting is applied. Initially, the time of the first detection by a low earth orbiting detection algorithm within the original period (23 h prior to the burned area detection) is noted, and compared to the recorded instances of fire detection from the threshold algorithm. If the first active fire detection from the threshold algorithm is within two hours of this first LEO active fire detection, the time of the detection by the LEO product is fixed as the seed for the new temporal window, which starts the temperature fitting process for the target pixel 22 hours prior to this resolved first detection. In instances where the threshold algorithm detects fire more than 2 hours prior to a LEO fire product detection, the time of the first detection of the AHI product seeds the temporal window time instead. Despite the change in fitting start times, the same 24 hour window is evaluated for fire detection as in Section 4.3.4, with any change in the initial time of AHI detection recorded.

4.4 Results

4.4.1 Threshold Selection and Anomaly Detection Rates

Using the selection method for clear sky pixels as described in Section 4.3.4, a total of 93,906 unique locations and starting times were identified as suitable for analysis, which corresponded to approximately 63% of the pixels selected initially. Table 4.1 displays the rates of detection of individual thermal anomalies each of the thresholds selected, where an anomaly detection is defined

4. Advances in Active Fire Detection Using a Multi-Temporal Method for Next-Generation Geostationary Satellite Data

as at least one brightness temperature being more than the threshold above the temperature fitting for the 24 h period fitted. The thresholding algorithm displays a very high number of anomaly detections at 2 K, with a low number of these being related to detected fire activity. As the threshold is increased, detection rates drop but the number of detections associated with burned area instances increases, to a point where at 5 K non-burned-area thresholding detections occur in less than 0.1% of all sampled pixels. The commission rate of the anomaly detection relative to burned area product detections remains relatively high at all thresholds, with the lowest rate of 23.4% at 5 K.

Table 4.1: Raw anomaly rates for the threshold algorithm for the selected temperature thresholds, and the proportion of these anomalies which have an associated disturbance detected by the MODIS burned area product, from a total selection of 93,906 cloud-free diurnal fittings.

Threshold (K)	AHI fittings with positive anomalies identified	Anomalous fittings with associated MODIS burned area
2 K	12.17%	7.9%
3 K	2.22%	22.0%
4 K	0.64%	53.7%
5 K	0.37%	76.6%

4.4.2 Active Fire Product Intercomparison

The selection process outlined in Section 4.3.3 identified 7625 separate instances of possible fire-induced disturbance across 2765 unique AHI pixels during the period of the study. The larger number of total fire incidences compared to AHI pixels affected is explained by the spatial disparity of the datasets used – the burned area product is of a far higher spatial resolution than the AHI imagery, so a number of separate incidents may occur in one AHI pixel area over the period of the study. Table 4.2 shows the rates of detection for the four fire incident groups as described in Section 4.3.4. Of note is the effect of the spatial resolution of the various sensors on the rate of complementary detection. Notwithstanding fires that have not been detected by the LEO active fire products, the type of classification roughly correlates with fire size – events that are detected by the VIIRS 375m product only are identified by AHI at a lesser rate than those found by the MODIS active fire product only, and both of these have lower detection rates than when both active fire products are triggered. Despite the spatial resolution disparity between the VIIRS active fire product and the AHI Band 7 images, the threshold algorithm is capable of detecting fire in these areas in over 75% of cases when using the largest of the

four thresholds evaluated.

Table 4.2: Detection results of the thresholding algorithm on 150 fire incidents in each detection grouping per temperature threshold. Detections occur where at least one brightness temperature measurement exceeds the fitted brightness temperature by the selected threshold. Synchronous fire detections are classified as where an anomaly detected by one or both of the active fire products has at least one corresponding detection from the threshold algorithm within twenty minutes of the LEO detection.

Group\Threshold <i>n=150 for all</i>	2 K		3 K		4 K		5 K	
	Detected	Synchronous	Detected	Synchronous	Detected	Synchronous	Detected	Synchronous
Burned area only	75.3%	N/A	63.3%	N/A	56.0%	N/A	50.0%	N/A
VIIRS AF only	95.3%	38.7%	88.0%	27.3%	84.7%	22.0%	77.3%	17.3%
MODIS AF only	97.3%	60.7%	97.0%	58.0%	91.3%	52.7%	86.0%	48.0%
Both AF products	99.3%	68.0%	98.3%	58.7%	92.0%	51.3%	89.3%	46.0%

Importantly, the threshold algorithm is detecting positive anomalies in pixels that have identified as being burned, but no LEO active fire product detection. In 50% of cases the 5 K threshold is able to identify a fire in a pixel where both of the LEO active fire products used here have failed, with higher attribution rates for the smaller thresholds.

4.4.3 Time of Initial Detection

Table 4.3 shows statistics concerning the detection rates of each threshold for the incident classes examined for the forward-shifted temporal window (Figure 4.3), along with the average times of the AHI fire detection before a low earth orbit detection. These mean time differences are shown both for detections from fittings with the temporal window affixed to the burned area product, and for detections from fitting windows affixed to the leading fire detection determined in the burned-area-fixed fitting. The three groups of incidents examined all show improved detection times with use of the threshold algorithm as opposed to the LEO active fire products. Of particular note is the improvement in detection times when dealing with fires that have been detected by both active fire products – these fires are generally greater in duration and radiative output than fires detected by a single active fire product. Even with multiple LEO detections available, the threshold algorithm is on average detecting a fire 210 minutes prior to the subsequent LEO detection using the 5 K threshold.

The addition of the temporal window shift increases the differences in detection time between AHI-based detections and the LEO products even further. The improved results using this shifted window demonstrate the effect of long-lived fire on the fitting method, with the larger set of background temperatures available earlier contributing to a better fit of background temperature and earlier detection. This is also more reflective of results that may be achieved using

4. Advances in Active Fire Detection Using a Multi-Temporal Method for Next-Generation Geostationary Satellite Data

Table 4.3: Time of detection of fires using the threshold algorithm in comparison to times of first detection using the two LEO active fire products at each temperature threshold. Times shown are the average time of detection prior to LEO active fire detection, with numbers shown for both the diurnal temporal window commencing 23 h prior to burned area detection, and for the shifted temporal window commencing 22 h before initial active fire detection. Times shown are in hours and minutes.

VIIRS Detection only (n=150)	2 K	3 K	4 K	5 K
Original detection rate	95.3%	88.0%	84.7%	77.3%
Shifted detection rate	95.3%	88.0%	85.3%	76.0%
Mean detection time before first LEO AF with original window	4h 48m	2h 41m	2h 07m	1h 55m
Mean detection time before first LEO AF with shifted window	6h 47m	6h 08m	6h 06m	5h 43m
MODIS Detection only (n=150)	2 K	3 K	4 K	5 K
Original detection rate	97.3%	94.0%	91.3%	86.0%
Shifted detection rate	91.3%	84.0%	82.0%	82.7%
Mean detection time before first LEO AF with original window	8h 06m	6h 28m	5h 42m	4h 49m
Mean detection time before first LEO AF with shifted window	9h 36m	7h 34m	6h 34m	5h 39m
Both AF Detected (n=150)	2 K	3 K	4 K	5 K
Original detection rate	99.3%	95.3%	92.0%	89.3%
Shifted detection rate	95.3%	89.3%	88.0%	84.7%
Mean detection time before first LEO AF with original window	5h 25m	4h 27m	3h 54m	3h 31m
Mean detection time before first LEO AF with shifted window	7h 26m	6h 09m	5h 35m	5h 24m

this method in a psuedo-realtime application, as the fitting window ends closer to the actual fire ignition time. The shift of the temporal window forward in time removes at most 10% of the fire detections, as these detections may occur on developing fire that becomes visible more than two hours after the original LEO active fire product detection. This drop in detections from the thresholding algorithm seems to be exclusively associated with the group of detections solely from the MODIS active fire product, with only two fires lost in the 5 K threshold from fire detections associated with VIIRS. This may suggest an increased rate of false positives in the area of study from the MODIS product is being isolated by this selection criteria.




Figure 4.3: Associated fittings applied to a pixel at 15.5409°S, 129.2377°E, with a MODIS burned area product detection at 2016-219 05:20 UTC (shown by the red vertical line). The algorithm detection threshold set is 4 K. This figure shows ongoing fire activity in the AH1 Band 7 brightness temperatures, shown here in blue, surrounding a single VIIRS active fire detection at 2016-218 16:40 UTC (in orange). With the temporal window based upon the time of the burned area product minus 23 hours, the first AH1 detection at the 4 K threshold occurs 90 minutes after the VIIRS active fire detection (initial window detections are black circles). However, the shifted temporal frame based upon the time of this VIIRS detection produces a lower fit for background temperature during this night-time period, and the initial fire detection from AH1 moves to 190 minutes before the VIIRS overpass.

4.5 Discussion

Whilst the AH1 sensor has shown impressive utility for fire detection in this study, the lack of validated products from the sensor data remains an issue for adoption of this thresholding method. As noted in [31], the ability to accurately evaluate rates of false detection is impeded by the lack of an accurate, verified cloud mask product for the AH1 sensor. Commission errors for cloud cover are particularly prevalent in coastal areas and during night-time periods. The adjustment made to the clear sky product for night-time use in this study is by no means a comprehensive fix, but considering the average of overnight temperatures in the area at the time of the year of the study, coupled with the tropical climate and lack of high elevation areas, brightness temperatures at the low level specified are rarely not associated with cloud cover. There are ongoing issues with cloud omission from the cloud product, especially with thinner clouds that may not be detected by a cloud product but still influence the brightness temperatures recorded by the sensor and decrease daily average

temperatures. If sufficient cloud is present for a period, the sudden absence of cloud will cause temperature spikes and lead to false positives. Cloud mask accuracy improvement will allow for a rigorous assessment of the method's commission rate in less than ideal conditions. The AHI sensor also suffers from some image registration issues which cause periodic shifting in brightness temperature values, which is especially notable in areas of high temperature contrast, such as coastal areas.

In order to investigate the distribution of likely causes for anomaly detection, a further visual analysis of the algorithm detections from the 4 K dataset from Section 4.4.1 was conducted. Incidents were broken into three categories – detections caused by sudden elevation of brightness temperatures, which reflect fire activity; detections caused by excessive cloud not captured by the clear sky product, which lead to poor temperature fitting; and anomalies with no immediately apparent cause. Examples of these assessed classes can be seen in Figure 4.4. In instances where an AHI pixel had a recorded MODIS burned area detection within the 48 hour period after fitting commenced, the cause of attribution in 96% of cases was due to fire activity detectable from AHI. Where no burned area detection occurred, 67% of detections from the thresholding method were anomalies caused by cloud, and these detection instances would have been eliminated from selection with better cloud masking. In the remaining threshold-based detections not associated with a burned area detection, a high proportion of detections occurred due to poor model fitting in the early morning period, which was also noted in [89], and these types of detection errors may also be reduced by better cloud detection during these times of the day.

Detection rates for thermal anomalies using the different threshold values are shown in Section 4.4.1, with raw detection rates decreasing as the thresholds are increased. Evaluation of the appropriate threshold that accurately describes fire activity in the area is difficult. The burned area product used in this study only gives an indication of the initial time of fire disturbance, and does not account for areas revisited by fire or with lingering fire activity over many days. There is also limited *in-situ* fire extent data available with which to verify the areas covered by fire activity during the study. The high commission rates shown in this table are of concern, but as discussed in the previous paragraph, much of the commission error can be explained by the inadequate performance of the clear sky product used.

The method presented in this paper for fire detection using the AHI sensor shows impressive rates of agreement with fire activity detected by the LEO active fire products used for comparison. Given the similarity of absolute detection rates between the MODIS active fire set and the instances where both




Figure 4.4: Examples of temperature fitting and sources of perturbation. a) shows an example of anomalies in pixel brightness temperature caused by fire activity, b) is an example of negative temperature anomalies causing false detections, in this case cloud cover, and c) shows an example of false detection caused by improper fitting of the diurnal model.

active fire products are triggered, these two groups could be considered as one set of results for comparison to MODIS. These larger fires (detectable at a 1 km^2 resolution) are almost completely detected by the thresholding algorithm at the 2 K threshold, with a detection rate closer to 88% at the larger 5 K threshold. For fires that are smaller than the detection capability of MODIS, but can be detected by the 375 m VIIRS product, detection rates drop to around 95% at 2 K and around 77% at 5 K. These numbers suggest that despite the smaller sizes of the fires detected solely by the VIIRS active fire product, fire activity in these pixels will increase at some time before or after the LEO detection to a magnitude that can be detected from the larger area pixel recorded by the geostationary satellite. This leads to the sensor's increased revisit rate being more than capable of making up for the spatial resolution shortfall of this geostationary sensor in a fire detection role, and reinforces geostationary detection behaviour noted by [89] in their study on SEVIRI active fire data.

Of most interest from the detection results is the increase in detection capability for fires which are not detected by either of the LEO active fire products. Given fire-driven disturbance detected by the burned area product used here, the threshold algorithm detects active fire in 75% of instances at 2 K, down to 50% of instances at 5 K. Fires detected in this set generally take the form of short-lived, intense flaming periods, with typical rises in recorded brightness temperature of between 10-20 K from the typical background fitting and active durations of four hours or less. A further inspection of this dataset revealed that of the 75 detections from the no active fire set of pixels at 5 K, 41 of

these instances had a detection from the threshold algorithm at a point in time synchronous (± 20 minutes) with a MODIS sensor overpass. Despite potential issues such as viewing angle, obscuration due to smoke and cloud, and the potential differences in the attribution of background temperature considering the contextual nature of the MODIS active fire algorithm, this is a seemingly high number of omissions by this active fire product in this area, although these numbers agree with those from [33] concerning the MODIS product. Application of the thresholding fire detection technique is required over a wider range of areas and weather conditions with similar comparisons to LEO overpass times in order to determine whether this heightened LEO omission rate occurs due to the physical geography of the selected area or whether it is more systematic in nature. Once again, the detection instances where no LEO active fire detection exists demonstrate the effectiveness of an increased rate of imaging for detection of active fire.

A number of potential sources of error may affect the results obtained in the detection analysis. The accuracy of the burned area product used [53] is affected by rainfall events in sparsely vegetated areas, which reduce surface reflectance, albeit to a much lesser extent than fire. If there is a significant amount of cloud cover causing burned area product commission errors, this may cause instability in the fitting model and lead to thresholding detection errors. Lack of accuracy in determining the time of first disturbance in a pixel can also lead to active fire activity being missed, especially in the case of cloud obscuration of the MODIS images used. There is also a small chance due to scaling issues that fire disturbance is detected in the MODIS pixel prior to the active fire being able to be detected at geostationary scales. The temporal fitting window settings used in this study are by no means definitive, and the nature of the results suggests that flexibility in the time period of fire detection selected leads to earlier detection of thermal anomalies. Issues may also manifest with the classifications used to analyse the threshold algorithm with the various LEO fire products. Locations of all three LEO-based products have been converted to AHI coordinate space before comparison has occurred, and this may lead to collocation errors between products due to satellite tracking, view angle and the re-sampling of LEO products.

With regard to the ability of AHI to detect fires at the same time as LEO sensor products, the number of successful synchronous detections compares well to synchronous detections from [94], although there are some differences. In this study the relevant comparison classes are those which have MODIS active fire detections, which includes the class that contains detections from both LEO active fire products. When evaluating the detections for synchronicity where both LEO active fire products are used, no distinction has been made with re-

gard to which LEO detection is synchronous to the AHI detection, and many of these synchronous returns will relate to VIIRS active fire detections instead. The [94] study used a spatial window of 5 x 5 pixels to evaluate concurrent detections, whereas this study registered all LEO data in the AHI coordinate system, and examined the indicated AHI pixel only. Of note also is the time window used for judging synchronicity in the [94] paper, which restricts examination to ± 6 minutes from the LEO detection, or at most one AHI image either side of that point in time. Whilst this is a rigorous method to assess simultaneous observations, this undersells the ability of the geostationary sensor to detect fire at a similar time due to the ephemeral nature of fire activity. An example of this can be found in Figure 4.3, where fire activity is quite easily detected before and after the LEO active fire product detection, but the relative lack of elevated temperatures in the AHI record around the time of the LEO overpass leads to the pixel instance being classified as non-synchronous even when fire activity is present. This behaviour is especially evident with fires in the VIIRS-AF-only dataset, which tend to comprise of fires with lower activity during LEO capture than other groups, and this may explain the reduced rate of synchronous capture for this group. Given the changing fire activity observed in many examples of this type and with refinement of this method, the ability of the AHI sensor to detect fires in instances such as these may be under-represented in the results reported here.

As has been discussed previously by [31], the accuracy of background temperature fitting reduces with an increased number of thermal anomalies. By nature of the fitting method used for determining background temperature in this instance, a large number of positive anomalies in the data to be fitted will cause the mean temperature of the day to rise, which will raise the curve fitted accordingly. Given the evaluation methods used in this paper this is unlikely to manifest as a source of omission – unless the brightness temperatures of an entire 24 hour period are uniformly raised by fire activity, generally peak periods of fire activity will occur during a diurnal cycle and be subsequently identified. The model fitting errors caused by increased fire activity preclude use of this method as a fire monitoring tool at this point. Given the introduction of a more resilient fitting technique which follows the ideal diurnal cycle and the theoretical background temperature curve more effectively, this method may be able to be extended into providing a measure of fire radiative power for the initial period of active fire incursion, but would most likely be limited to examination of the first 24 hours of fire activity. Whilst ongoing monitoring of a fire's size and intensity is of interest to both ecologists and land managers, this capability is beyond the scope of this particular method of fire detection.

Understandably, with the far higher number of images of active fire recorded by the geostationary sensor, the most marked improvement in detection ability is in determining fire ignition time as previously demonstrated using both contextual [88, 45] and temporal [65] based detection algorithms with geostationary data. In this study, mean values for detection time for the window fixed by the burned area product range from just under two hours in the case of VIIRS and the conservative 5 K threshold, up to beyond 8 hours earlier with MODIS active fire at 2 K, with detection times improving across the board with threshold temperature reduction. Improvements in the accuracy of fitting will likely increase the time differences between the LEO products and the threshold algorithm further.

The extreme detection time improvements associated with the threshold algorithm in the MODIS active fire detection dataset most likely has more to do with the dataset selection criteria than anything else. Fires that are detected using the MODIS active fire product exclusively are more likely to occur on days of sporadic cloud activity, where the 3 – 5 measurements of MODIS give a greater chance of detection than the two images available through VIIRS. The fitting process used for the threshold algorithm is vulnerable to increased model instability on days of increased cloud cover, which may lead to increased commission error in conditions of this type. Given the very small total of MODIS-only fires with respect to the total burned area related dataset, it is possible that this subset of pixels also contains an inflated number of MODIS active fire commissions, although this tends not to be borne out in the detection rates presented.

Of note here also is the significant improvement of the threshold algorithm in detection time of fire activity for fires solely detected by the VIIRS active fire product once the shifted temporal window is applied. Notwithstanding the random element of cloud obscuration, from visual examination of the AHI temperature data the majority of fires found exclusively by the VIIRS active fire product in this case are areas where the peak fire activity occurs well before the VIIRS overpass, and the VIIRS product detects lingering burning and smouldering activity in the post-burn pixel. Shifting the temporal window helps the initial detection of fire by significantly increasing the amount of fire-free diurnal signal the fitting process can use, which leads to a more accurate estimation of the background temperature at the fire's ignition time.

By nature of the timing of the active fire and burned area products used, there will be issues with the exact comparison of the AHI active fire product and the various active fire and burned area products produced by the LEO sensors in the time periods specified. For instance, a burned area product generally requires the use of visible band imagery to detect changes in surface

reflection, which is not possible to detect at night. This restricts the triggering of the temporal window for temperature fitting either to the morning or afternoon MODIS overpass, with morning burned area detections corresponding roughly to fire activity during the previous day and afternoon detections corresponding to fire activity in the immediate hours prior to the sensor overpass. This invariably leads to the ends of the fitting process occurring during the more rapidly changing daytime period, and may manifest in larger temperature attribution errors than if the fitting times were more evenly distributed across the diurnal cycle.

This study has shown the results of anomaly detection using the BAT-based temperature fitting technique at a number of temperature thresholds. Assessing the ideal temperature threshold to use for this particular area at the time of the study is an exercise in balancing potentially heightened false detection rates associated with lower temperature thresholds with their superior detection capabilities compared to higher thresholds. Improvements in the accuracy of the fitting method used will reduce the detection errors associated with this technique, and may allow for further reductions in detection threshold temperatures. Application of this method of fire detection over differing landforms, land cover types, viewing angles, and times of the year will further assist in the appropriate selection of detection thresholds over wider areas of the AHI full disk.

4.6 Conclusions

This paper outlines a method of fire detection utilising the training data acquisition method outlined in [31] to provide fitting information for 4 μm brightness temperatures derived from the AHI-8 sensor over part of northern Australia. This fitting of the idealised background temperature is then used to provide a baseline for thresholding of the sensor brightness temperature information to provide a method of determining the time and propensity of thermal anomalies. The method is capable of detection of between 75 - 99% of thermal anomalies reported by associated LEO products during the period of study, depending upon the threshold selected, with similar frequency of synchronous detections in comparison to other studies of this type. The method also detects between 50 - 75% of fires which have a detection from burned area records, but which have no detection using either of the LEO active fire products examined. Significant improvements in the initial time of detection of fire activity have also been achieved, with detection time differences of 5 - 7 h in favour of this method over detections from the examined LEO fire prod-

ucts. This study supports the conclusions of previous work on the AHI sensor that highlight its ability to accurately detect and attribute smaller fires than previous geostationary sensors, but also demonstrates the effectiveness of the new generation of geostationary sensors in detecting thermal anomalies that may not be detected by LEO active fire products. With the similarities between the AHI sensor and the newly launched GOES-ABI sensor, this work with some adjustments should be directly applicable to fire detection studies using this new sensor in the western hemisphere.

4.7 Thesis Context

This chapter examined the use of the Broad Area Training technique for anomaly detection purposes in comparison to commonly used fire detection products from low earth orbiting sensors. Considering the complexity of the BAT process, and the difficulties involved with applying such a technique to images captured from low earth orbiting sensors, potential uses of the BAT method are currently more limited than desirable. The next chapter introduces a change of direction with regard to brightness temperature estimation, which is less based upon stochastic modelling of upwelling radiation trends and more upon leverage of real measurements from prediction images.

A New Spatio-Temporal Selection Algorithm for Estimating Up-welling Medium-Wave Radiation

5

5.1 Introduction

Remote sensing is a powerful tool often used to investigate changes in the landscape over time, and has been employed in this way across applications including vegetation change [36], urban growth [2], and disaster response [6]. The mapping of change in the physical environment requires some knowledge of the state of the landscape pre-change, in order to determine the nature and magnitude of such changes [55]. In an ideal case, change tracking would make use of data that spans the temporal domain, as well as the spatial domain. Given knowledge about how the earth's surface reacts to nominal physical phenomena over time, predictions can be made about a landscape's expected behaviour at a subsequent point in time. This information can then be leveraged to provide a method of isolating and identifying anomalous landscape-level behaviour, along with identifying obscuring influences such as cloud, smoke and fire, by comparing a predicted image to data recorded "in reality".

Fire detection is a well established application of remote sensing, with many commonly used products produced from both low-earth orbit and geostationary sensors [61, 26, 89, 45]. The most important element of the electromagnetic spectrum for these purposes is the Medium-Wave InfraRed (MWIR) (3–4 μm), where peak emission from the fire energy source occurs. Excess radiation at this wavelength allows fire to be detected when it constitutes only a small portion of an image segment (down to $\approx 10^{-5}$ of a pixel) [8]. Estimating up-welling radiation at this wavelength is complicated by the dual source nature of electromagnetic energy, with components made up of thermal emission and solar reflection [5].

Detection algorithms such as those in [26] rely upon the difference of the candidate fire pixel from a reference background value. Generally this reference value is derived from a convolution-style filtering approach, where values from surrounding pixels are averaged to provide the temperature estimate. A comprehensive examination of this practice was undertaken in [32], which highlighted shortcomings in the use of such a method, especially in areas of high spatial frequency (i.e. heterogeneous landscapes). The study showed the contextual derivation of temperature acted in a manner similar to an edge detector in areas exhibiting rapid spatial changes, increasing errors in temperature estimation variability. These areas are often of vital importance for fire detection purposes, as they are often found on the edges of areas of anthropomorphic change, boundaries between forests and urban areas especially.

This chapter has been submitted for publication as: Hally, B., Wallace, L., Reinke, K., & Jones, S. (2019) A New Spatio-Temporal Selection Algorithm for Estimating Up-welling Medium-Wave Radiation. to *IEEE Transactions in Geoscience and Remote Sensing*

The practice of using convolution filtering for estimation of MWIR background radiation is based upon the spatial autocorrelation effect, where areas near to a specific location are assumed to exhibit characteristics more similar to that location than those further away [82]. The top of atmosphere solar radiation component of an up-ward welling signal will strongly adhere to behaviour similar to this, due to its strong relationship with solar azimuth and zenith angles [55]. Surface characteristics such as the slope and aspect of a surface, land cover type, and water bodies can all affect the resultant signal emanating from a pixel. However, in areas where human activity has created frequent discrete and abrupt changes in a landscape, the strong similarities found in untouched landscapes are less likely to hold true. In many cases, the likelihood is high that areas exist outside of the immediate proximity of a target pixel that more comprehensively characterise the signal of a location than those closer by. Given that typical phenomena that may obscure up-ward welling signal from a candidate pixel (fires for instance) will display a high level of spatial-autocorrelation themselves and are more likely to influence areas proximal to a potential target, having a method of estimation of radiation that is less reliant on the local area may also provide greater robustness in a solution.

This study seeks to introduce a new method of brightness temperature estimation in the MWIR, based upon determination of locations that more closely resemble the behaviour of MWIR radiation of the target pixel with a defined local region. This Spatio-Temporal Selection (STS) method (Section 5.2) tracks the history of brightness temperatures for a given period, selects training pixels based upon their statistical fit to the target pixel, and makes predictions for brightness temperature based upon values from a prediction image. Section 5.3 will cover the assessment of the method's validity against the common estimation method and Section 5.4 examines both the results and the potential of the method for improving image reproduction for other purposes.

5.2 Spatio-Temporal Selection

This section describes the proposed method implementation over the selected dataset (Section 5.2.1), followed by a formal description (Section 5.2.2) clarifying the main framework of the method, and outlining the parameters that can be modified for refinement of solutions.




Figure 5.1: Locations of the case study areas selected for analysis in this paper, depicted on the AH1 full disk.

5.2.1 Test Data Example

For this study, we have used data from the Japan Meteorological Agency's AH1-8 sensor upon the Himawari-8 geostationary satellite. This satellite, positioned in geostationary orbit at 140.7°E longitude, provides coverage over eastern Asia, the western Pacific and Australia with a 16 band multispectral imager (three VIS, 3 NIR, 1 MWIR, 9 TIR bands) [60]. Whilst the main purpose of the sensor is for meteorological forecasting, the high temporal resolution (full disk recorded every 10 min) has encouraged use of the imagery for detection and monitoring of rapid change in the environment. Images captured from AH1's MWIR Band 7 (3.7 μm) were masked for removal of water pixels using the ancillary land-sea masks supplied by the Australian Bureau of Meteorology. A cloud mask was also applied to the images from this band, based upon the mask applied to AH1-8 outlined in [94], which was adopted from a similar mask applied to GOES imagery [95].

To facilitate the study, a number of 200 x 200 pixel case study areas were identified across the AH1 full disk, which are detailed in Table 5.1 and shown in Figure 5.1. These areas underwent an analysis of fire activity for the year of

2016 using the VNP14IMGML VIIRS active fire product [75], with the peak of fire activity in a 30-day rolling window over the year adopted as the central day of a 31 day period of examination. These case study areas were then divided into 50×50 pixel regions, areas that mainly consisted of sea pixels were dismissed, and from the remaining areas seven 50×50 regions were randomly selected for the analysis. Times selected for the pixel predictions were at the local times corresponding to 09:00, 12:00, 15:00, and midnight, with a random offset of (0:50:10) minutes to provide some statistical independence for overlapping training sets. The study areas selected cover about 18.5 Mha of land in total.

The procedure for providing training data for the image reconstruction was to select a circular area of 50 pixels in radius surrounding each of the target pixels in the region, stack the previous 48 images at 2 hourly intervals together, and calculate the difference between all pixels in the surrounding radius and the target. A root mean square error could then be calculated from all of these temperature differences in the training set of images, similar to that shown in the left images in Figure 5.3. Any training pixels that had less than four coincident observations with the target over the 48 images were deemed to have insufficient data to determine correlation, and were eliminated from further analysis. From the remaining potential training pixels, the RMSEs were sorted, and the 24 pixels with the lowest error compared to the target were selected for the prediction phase (shown in the right images in Figure 5.3).

At the prediction step, a minimum of six training pixels had to be available in order to provide a STS estimate of brightness temperature - targets with less than six were discarded from analysis. The unmasked values of the training pixels from the prediction image were then filtered for outliers (removal of $\geq \pm 2\sigma$), and the remaining values were averaged for the target estimation. Comparisons of brightness temperatures for analysis purposes were undertaken where solutions for STS estimation and the contextual estimate were coincident with a raw brightness temperature from the prediction image. Contextual estimates were calculated based upon the guidelines for computational accuracy set out in [32], which determined that 5×5 contextual estimates with at least 65% of adjacent pixel availability were the minimum to ensure contextual estimation accuracy.

5.2.2 Formal Description

Figure 5.2 demonstrates the STS method from first principles. Important variables for altering the method's implementation depending upon the conditions of use are as follows:

- n_{image} - number of images in the pixel training stack.

5. A New Spatio-Temporal Selection Algorithm for Estimating Up-welling Medium-Wave Radiation

Table 5.1: Specifications for the timeframes, area of the AHI disk and UTC time offsets for each of the case study areas examined.

CS area	Start Date	End Date	UTC Hours	Sectors used
sea	2016-03-30	2016-04-29	02:00-02:50 05:00-05:50 14:00-14:50 23:00-23:50	sea_b: [4400, 4550, 3000, 3150] sea_c: [4450, 4600, 3000, 3150] sea_e: [4350, 4500, 3050, 3200] sea_f: [4400, 4550, 3050, 3200] sea_h: [4500, 4650, 3050, 3200] sea_j: [4350, 4500, 3100, 3250] sea_k: [4400, 4550, 3100, 3250]
nwa	2016-10-23	2016-11-22	00:00-00:50 03:00-03:50 06:00-06:50 15:00-15:50	nwa_b: [3600, 3750, 1950, 2100] nwa_c: [3650, 3800, 1950, 2100] nwa_e: [3550, 3700, 2000, 2150] nwa_f: [3600, 3750, 2000, 2150] nwa_g: [3650, 3800, 2000, 2150] nwa_p: [3600, 3750, 2100, 2250] nwa_q: [3650, 3800, 2100, 2250]
bor	2016-02-14	2016-03-15	01:00-01:50 04:00-04:50 07:00-07:50 16:00-16:50	bor_a: [2550, 2700, 1350, 1500] bor_f: [2600, 2750, 1400, 1550] bor_g: [2650, 2800, 1400, 1550] bor_h: [2700, 2850, 1400, 1550] bor_j: [2550, 2700, 1450, 1600] bor_l: [2650, 2800, 1450, 1600] bor_p: [2600, 2750, 1500, 1650]
thl	2016-02-28	2016-03-29	02:00-02:50 05:00-05:50 08:00-08:50 17:00-17:50	thl_a: [1750, 1900, 750, 900] thl_c: [1850, 2000, 750, 900] thl_j: [1750, 1900, 850, 1000] thl_k: [1800, 1950, 850, 1000] thl_m: [1900, 2050, 850, 1000] thl_p: [1800, 1950, 900, 1050] thl_q: [1850, 2000, 900, 1050]
chn	2016-08-27	2016-09-26	01:00-01:50 04:00-04:50 07:00-07:50 16:00-16:50	chn_a: [950, 1100, 1550, 1700] chn_b: [1000, 1150, 1550, 1700] chn_e: [950, 1100, 1600, 1750] chn_g: [1050, 1200, 1600, 1750] chn_k: [1000, 1150, 1650, 1800] chn_m: [1100, 1250, 1650, 1800] chn_q: [1050, 1200, 1700, 1850]
jpn	2016-05-03	2016-06-02	00:00-00:50 03:00-03:50 06:00-06:50 15:00-15:50	jpn_b: [900, 1050, 2450, 2600] jpn_e: [850, 1000, 2500, 2650] jpn_f: [900, 1050, 2500, 2650] jpn_j: [850, 1000, 2550, 2700] jpn_k: [900, 1050, 2550, 2700] jpn_n: [850, 1000, 2600, 2750] jpn_p: [900, 1050, 2600, 2750]
sib	2016-05-10	2016-06-09	01:00-01:50 04:00-04:50 07:00-07:50 16:00-16:50	sib_b: [200, 350, 1950, 2100] sib_d: [300, 450, 1950, 2100] sib_j: [150, 300, 2050, 2200] sib_l: [250, 400, 2050, 2200] sib_m: [300, 450, 2050, 2200] sib_n: [150, 300, 2100, 2250] sib_r: [300, 450, 2100, 2250]

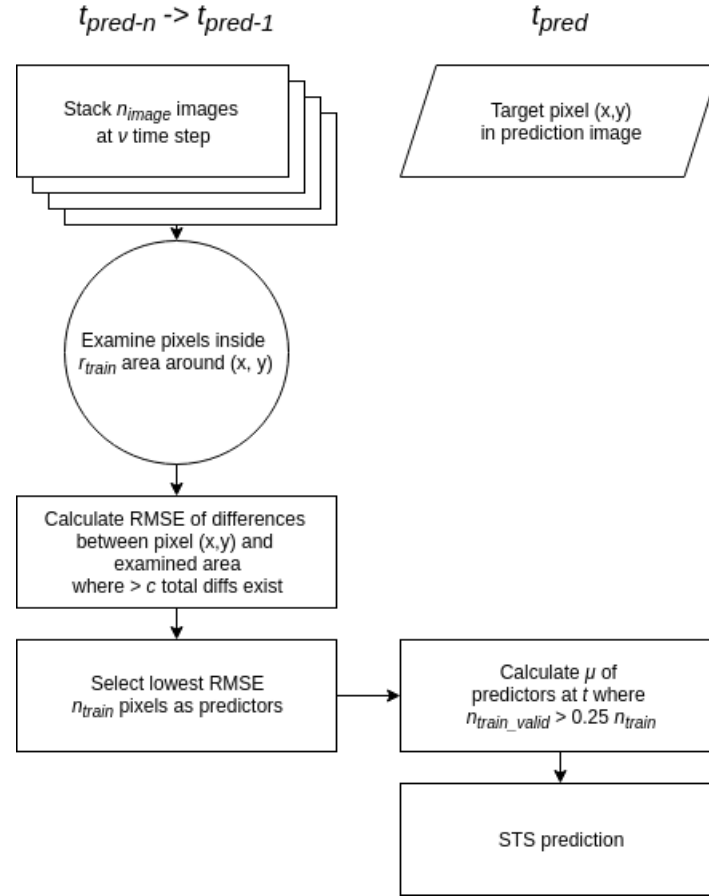


Figure 5.2: Flowchart of the STS selection and estimation process.

- ν - time gap between images in the pixel training stack.
- r_{train} - the radius of restraint for searching for training pixels around a target pixel.
- c - the total number of coincident measurements between potential training pixels and the target in the training stack.
- n_{train} - the number of training pixels selected for use in the target estimation.
- λ - the time period from the prediction stack creation whereby training selections remain valid.

By altering these parameters from those set in this study, a number of theoretical affects may improve the solution derived. For instance, n_{image} sets the number of images over which to assess the validity of training pixels. Generally the larger the number of assessed images to judge training suitability the

better, but this value could be traded off for a larger search radius to maintain processing efficiency. n_{image} is also closely related to the altering of the value of ν - in this study we have selected values of these two parameters at 48 images and two hours respectively in order to find a balance of measurements over the diurnal cycle of the training area. The alteration of these values to affect training accuracy is related to cloud conditions in the training area during a predictive assessment. Lengthening the time over which pixels are assessed for their suitability should mitigate for weather systems and their associated cloud, but may not account for major alterations that occur in a landscape due to fires and rainfall events. Training should also occur over a range of diurnal conditions to mitigate the effects of developing convective clouds, which are a constant feature in some of the assessed landscapes during the afternoon daytime period.

r_{train} and n_{train} speak to the suitability of the surrounding area around the target for providing sufficient training data. The search radius r_{train} should be increased or decreased dependent upon the likelihood of correct characterisation, noting that in certain cases (like in fig. 5.3c) suitable pixels may not display the typical spatial autocorrelation pattern usually associated with temperature estimation. Increasing the number of training pixels selected may improve the likelihood of obtaining a valid estimation, especially in cloudy prediction images, but will likely result in less accurate estimation when the prediction image is clear, as the extra values for the target estimation will be coming from less correlated pixel values. The c value set in the process relates to the expected accuracy of predictor pixels - setting this too low may result in the selection of training pixels that are highly correlated for a very short portion of the training stack, but otherwise have little in common, whereas setting this too high may reduce the effectiveness of the training pixel search, especially if the number of values approaches the number of valid measurements on the target pixel over the training period.

5.2.3 Overall Accuracy Assessment

The typical method of assessment of estimation accuracy is based upon the variation of the respective estimation methods from the recorded values. Whilst this is a sensible approach to a perfect landscape with no occluding features, the inclusion of comparisons to anomalous temperature values from the imagery, whether from fire activity or misattributed cloud, can lead to this type of assessment being flawed. A situation where obvious image contamination has occurred would penalise the accuracy of an estimation method that is correctly identifying an error, with the level of penalty increasing with the effectiveness

of the error identification. As such, results for this study have been produced both with these obvious errors, for the purposes of comparison with previous studies, and without these anomalous values. Upon visual examination of the images, an anomaly rate of 2 % was adopted as the standard level of error in the brightness temperature images, in order to not penalise the accurate identification of these errors.

Two sets of accuracy assessment figures will be presented in the results - a comparison of the STS and context estimates to the image temperatures, both including and without the largest 2 % of anomalous differences from the image temperatures, measured in absolute terms. The mean and standard deviation of differences from the image temperature will be reported along with the percentage difference in standard deviation between the two methods.

5.3 Results

5.3.1 Training Pixel Selection

Figure 5.3 displays a typical set of training data selections and subsequent estimations for a number of adjacent target pixels on the east coast of Kalimantan. The maps to the left of this figure depict of the RMSEs of each of the potential training pixels relative to the target, with lighter colours in areas that appear more like the target pixel. In the depicted cases, there is an obvious trend whereby noise values increase with the distance from the target pixel, which is not unexpected behaviour. Of interest though is the propensity of selection of training pixels to occur in areas of similar makeup. In fig. 5.3a lower RMSEs occur in the strip of land immediately inland from the coast - this area is heavily cleared and contains some urban areas in comparison to other parts of this region. The second pixel selected more heavily favours those areas, and the selected training pixels stretch out along the coastal fringe. The pixel depicted in fig. 5.3c is coastal in nature, and the lowest RMSEs of training pixels reflect that - the pixels selected for training purposes by STS are strung out in a line along the coast reflecting this characterisation.

The right half of Figure 5.3 shows the various pixel trajectories over the STS training period. The red target value is compared to the temperature values of all pixels in the training set, but only pixels selected for fitting, along with the pixels describing the 5×5 area surrounding the target, are shown here. Figure 5.3a has the least noisy contextual temperature estimation shown here, with most disagreement between the target and the context values occurring in the middle of the day. The blue training values tend to stick closer to the target, with some mild deviations mostly in the night-time period where at times no

5. A New Spatio-Temporal Selection Algorithm for Estimating Up-welling Medium-Wave Radiation

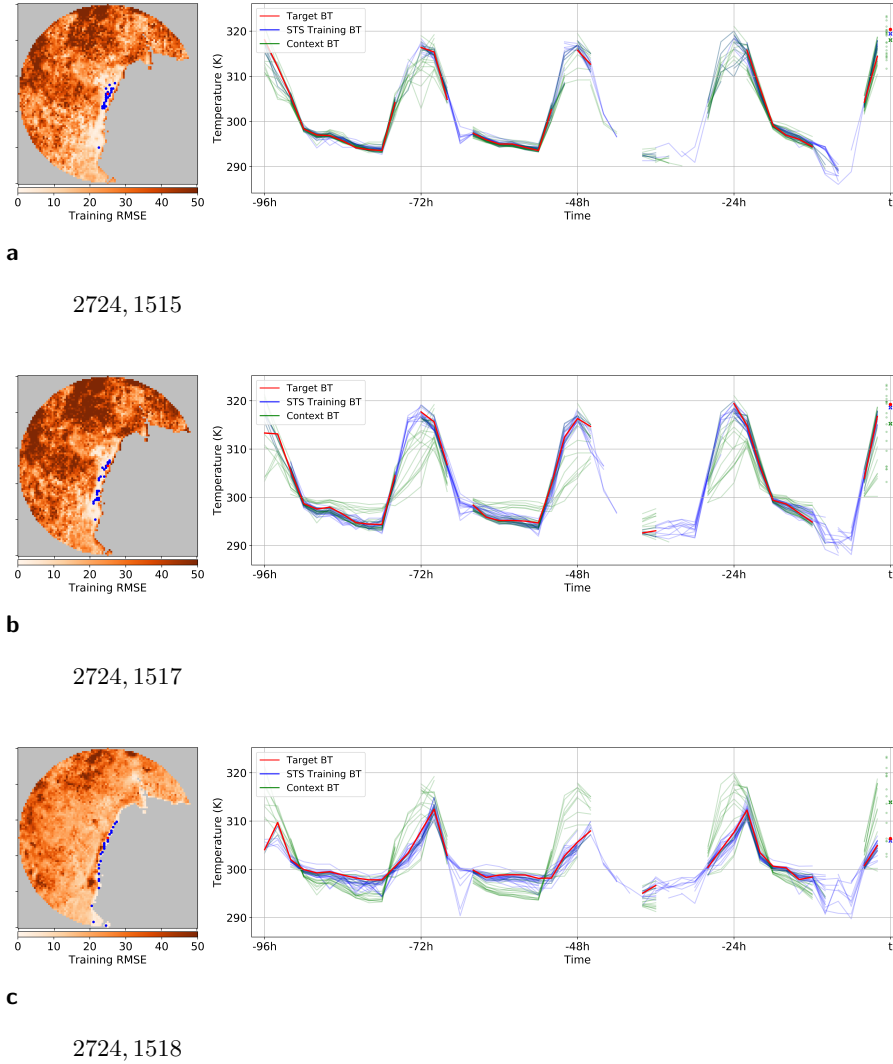


Figure 5.3: Pixel training comparisons for selected pixels in the bor_l group. (Left) shows the spatial distribution of points selected during the training process relative to the 50-pixel radius selection area; and (Right) depicts the pixel trajectories over the image set examined for training, with the prediction target pixel value shown in red, the STS training pixel values shown in blue, and the surrounding context pixel values in green. Shown at time t is the distribution of values in the prediction image from both prediction methods, with their respective means shown as coloured crosses in comparison to the recorded brightness temperature shown as a red dot. Pixels are shown as labelled with prediction time t at 2016-067 04:20 UTC.

comparison takes place. At time t the contextual values adjacent are shown in green, whereas temperatures from the STS training group are shown in blue.

Table 5.2: Accuracy of estimation techniques against brightness temperature values from the assessed images by case study area. $\Delta\sigma$ is the percentage change between the standard deviations of the context and STS estimation methods.

Site	Anomalies retained					Anomalies removed				
	Context		STS		$\Delta\sigma\%$	Context		STS		$\Delta\sigma\%$
	μ (K)	σ (K)	μ (K)	σ (K)		μ (K)	σ (K)	μ (K)	σ (K)	
sea	0.004	1.328	0.006	1.066	-19.7	0.014	1.144	0.011	0.849	-25.8
nwa	0.020	1.348	0.052	1.292	-4.1	0.039	1.147	0.078	1.025	-10.6
bor	0.034	1.239	0.142	1.429	15.3	0.032	1.045	0.149	1.203	15.1
thl	0.010	1.274	0.045	0.849	-33.4	0.032	1.105	0.048	0.665	-39.8
chn	0.011	0.937	0.047	0.757	-19.2	-0.006	0.814	0.065	0.635	-22.0
jpn	0.024	1.576	0.027	1.213	-23.0	-0.013	1.413	0.018	1.040	-26.4
sib	0.008	1.541	-0.021	2.139	38.8	0.017	1.327	-0.005	1.779	34.0

In this case the range of values seen in contextual pixels is around 7 K larger, with the contextual mean approximately 5 K below the target figure. The blue STS pixels are more closely grouped, with the mean value of the STS solution far closer to the recorded value. In this case, the STS training set shares nine pixels with the contextual surrounds of the target, with the values obtained from pixels outside of this region strengthening the resulting solution.

The two figures fig. 5.3b and fig. 5.3c show more extreme examples of potential pixel trajectories from a highly variable landscape. The contextual pixel values shown in fig. 5.3b once again show the most variation during the day, but also demonstrate variation during the night. This type of lagging temperature variation is a hallmark of coastal locations, where the pixels are often a mixture of land and water. The water portion of these pixels tends to retain heat during the night in comparison to the adjacent land areas, and during the day the reduced reflection coupled with delay in heating of the water compared to land areas results in the peak temperatures being later in the day. Using the STS pixels reduces the effects of these pixels on the estimation, once again with a tighter grouping of values and less low-value outliers. Figure 5.3c shows the fitting of a coastal pixel, with the lagging of the STS pixel values compared to the higher contextual temperatures. The resultant set of pixel values to base estimates on is much more tightly grouped in the case of STS, with the resultant mean estimates providing a decrease of estimation error of 15 K when compared to the contextual estimate. Given the appearance of the pixel trajectories in this case, it is likely that the attributed error in the contextual method in this region is permanent in nature, with seasonality and tidal effects being the major contributing factors to potential variation in this error.

5.3.2 Overall Accuracy Assessment

Table 5.2 describes the errors and standard deviations of the two assessed estimation methods against the measured brightness temperature values for each case study area. Biases of both methods tend to be fairly low, with all but one study area having mean differences of less than 0.1K. Increases in the means of the STS estimation tend to be due to the omission error of the cloud mask used, which drives estimates from STS lower due to the prevalence of misattributed cloud. The four study areas showing the most improvement in variation were **thl**, **chn**, **jpn**, and **sea**, with decreases in estimate variation of between 19–34% in the dataset including anomalies. Of particular note is the improvement in estimate error in the **chn** study error, where contextual estimation performs significantly better than in the other areas covered.

Temperature estimation using the STS method remains a challenge in the **bor** and **sib** case study areas, with increases in the variation of estimates of 15% and 39% respectively. Errors in these areas with STS estimation are generally caused by the large amount of cloud occlusion present in these image sets, both with regard to the training period, which causes mis-selection of training pixels, and in the prediction image, where cloud values reduce the number of training pixels that can contribute to a solution for a target.

The error rates for the estimates reported with anomalous temperatures removed shows a decrease in variance for both estimation methods, with a 10–17% drop in contextual estimation error, and a 16–28% drop in STS estimation error. On the whole the elimination of outliers treats the remaining STS variations favourably, with an decrease in comparison to the context estimates across all sites. The selection of the 2% anomaly rate seems to be supported by these numbers, although decreases in reported variance were much lower in the **jpn** area than others, due to a much higher rate of anomalous temperature differences resultant from the extreme landscape and land use variability in the area. If this "outliers removed" dataset is assumed to be a better account of temperature estimation for the bulk of temperature values, temperature characterisation is improving using the STS method by up to 40% in favourable conditions.

5.3.3 Image Assessment

Figure 5.4 shows a series of images that demonstrate both the STS and contextual methods over a subset of the **thl** region. The area shown is centred over the Loei province of northern Thailand, with the northern half of the image over western Laos. From left to right, images shown are the brightness temperature image from AHI Band 7, the STS estimation of the region, and

the contextual estimation of the region. The two rightmost images show the differences between the first image and the second and third images respectively, giving an overall assessment of temperature differences. At first glance, in fig. 5.4a there appears a marked similarity in the landscape depiction produced by STS in comparison to the sensor image. Fine details in the image are retained – the silhouette of the Mekong River shown in the upper centre of the brightness temperature image is also evident in the STS estimation, along with the stratification of temperature zones in the north and west of the region. In contrast, the contextual image performs the role of a smoothing filter, reducing the contrast of the image overall and dulling the finer details of temperature change. These effects can be seen most markedly in the differencing images with an overall decrease in both the high and low temperature magnitudes seen on the STS difference image compared to the equivalent for context. Despite the overall lowering of temperature variation, the STS method is still identifying some anomalous pixels in the original image, in the west and south of the STS difference image.

The increased error experienced at the midday period can be seen in fig. 5.4b, which shows the result of an estimation at noon local time. A marked feature of this particular image time is a number of anomalies shown in the STS difference image in the area. The large amount of red pixels that may signify agricultural burning in the area are easy to identify on the STS difference image, but tend to get mixed with the landscape edge effects in the contextual difference image. One notable effect of these high temperature anomalies in the context difference is the subsequent effect these higher temperatures have on estimation at the edge of the anomalies. The high temperatures tend to result in a ring of low temperature results surrounding this activity – generally fire detection algorithms remove this problem by flagging the fire pixels as anomalous before estimating surrounding temperatures, but a treatment such as this has not been attempted on the data presented here. Figure 5.4c and fig. 5.4d show the subsequent progression of the temperatures in the area through the afternoon and into the next night. Landscape patterns seen during the day tend to change at night due to differences in surface emissivity and subsequent heat retention/loss, but the STS estimation has no issue with maintaining image reproduction quality even with changes in the distribution of temperature gradients. Some anomalous temperature activity in this area continues throughout the night – the contextual estimation highlights the anomalies more strongly at night, but also provides a far noisier solution for pixels not undergoing anomalous activity.

5. A New Spatio-Temporal Selection Algorithm for Estimating Up-welling Medium-Wave Radiation

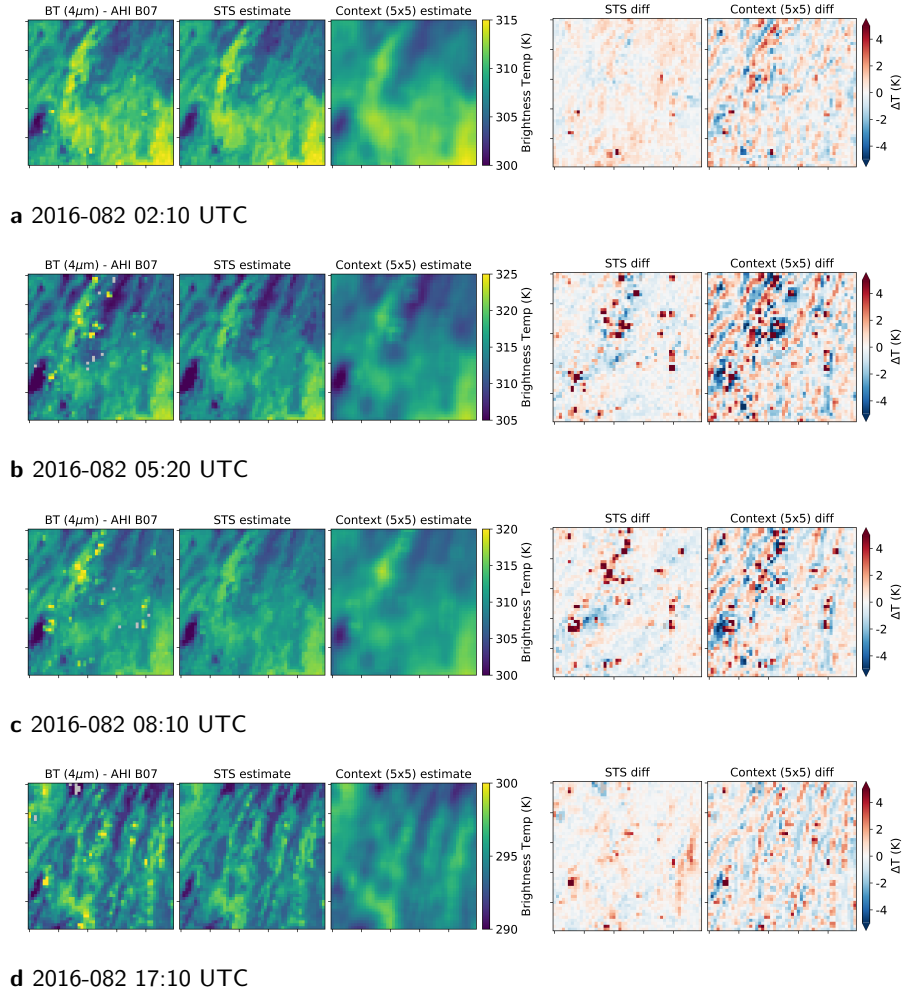


Figure 5.4: A series of brightness temperature images and related estimations for the thl_j region. From left to right, the AHI B07 brightness temperature at the prediction time, the STS prediction image of the area, the contextual estimation of the area, and the differences between the AHI image and STS estimates, and the AHI image and context, are shown. The differences shown highlight positions where the recorded image value is higher than the estimation (red) and vice versa (blue). Prediction times are shown next to each figure.

Table 5.3: Breakdown of the availability of temperature values using the two estimation methods against total image pixels present. n BT Obs gives the number of cloud free image pixels out of the total possible shown in Total Pixel Obs.

Site	Total Pixel Obs	n BT Obs	% context-image	% sts-image	% sts-context
sea	2142472	1253867	93.96	120.17	127.90
nwa	2170000	1634662	96.65	116.40	120.44
bor	1694460	1028491	90.68	131.03	144.49
thl	2162560	1715681	97.23	112.96	116.18
chn	2149044	1153593	95.10	118.34	124.45
jpn	1754724	725493	90.68	119.36	131.63
sib	2168264	801329	90.19	117.36	130.12

5.3.4 Estimation Availability

Table 5.3 shows a breakdown of the availability of estimations using the STS and context methods against the recorded brightness temperatures from images. The first two columns of this table report the total amount of land pixels assessed from the 7 50×50 regions from each case study for the 124 images of the assessment period, and the number of pixels that recorded a brightness temperature value once the cloud mask was applied. Cloud as assessed by the cloud mask was least prevalent in the **thl** and **nwa** areas, and most prevalent in the **jpn** and **sib** areas with more than half of all pixels affected by cloud. Contextual estimation can occur in 90–97% of cases where brightness temperature values exist, which is a higher rate of estimates available than reported in [32], but this rate is highly dependent upon the cloud identified by the cloud mask used. With the STS training selection process, the diversity of the sampling area means that estimation can take place in more fragmented images where contextual estimation may not be possible. Of course, due to the 50 pixel buffer applied to each temperature estimation process, the area that can be potentially sampled is far higher than for contextual temperature estimation, but if the assumption is made that average cloud conditions are similar over the buffered area to those over the assessed area, the STS method provides estimates of temperature in around 20% more cases than the brightness temperature images themselves. It is beyond the scope of this study to assess the accuracy of those extra pixels provided by the STS method, but there is potential to provide a source of validation for missing values in an image using STS in cases such as these.

5.4 Discussion

With regard to parameter selection for STS in this study, it is important to note that only one set of parameters have been examined, and this affects the results coming from the STS estimation in a number of ways. The training of pixels in this instance occurs over the preceding 48 images at a two hour spacing - a total training period of 96 hours. Training accuracy is reliant on pixels gathering variation through the time series in order to filter out less accurate prediction pixels - significant periods of cloud during this training process will lead to a reduction of coincident measurements between the target and ideal prediction pixels within the search radius. This weakness to cloud cover during the training period could be mitigated by variation of the time between training images - lengthening the time between images would result in the lessening influence of weather effects in the short term. Care must be taken though to ensure that training happens within a reasonable period before the prediction step, as changes to the landscape will increase the error of those pixels when attempting to find matches. The opportunity also exists to shorten the time step if a known period of clear sky exists before the prediction step. A fuller reckoning of the method's accuracy when altering the time step is required to make recommendations on these effects. Another opportunity the STS method affords to estimation of temperatures is the ability to select persistent training pixels, which could also assist in periods where training data is less accurate. Training pixel locations can be held fixed over subsequent prediction steps if the noise in the nearer-term time series of predictors is too great, and this is also another potential topic for investigation.

It is less certain whether other parameters set in this implementation of STS require further iteration. The number of training pixels set for prediction is fixed in this instance at 24, and this was done to provide a direct comparison to the number of pixels usually available for contextual estimation. The pixels selected in the training step are the pixels most like the target available - adding extra training pixels will add some extra robustness to estimation in the face of cloud in the prediction image, whilst adding extra noise to the solution in more favourable conditions. The likelihood is high that users will prefer most robust solutions on cloudy days, and the exact number of training pixels required to supply this requires further examination. The radius of prediction set in this study has generally resulted in an accurate set of training pixels for estimation in most cases, with notable exceptions being in some coastal and urban areas, and where major landscape change has occurred during the training period. Locations such as these may benefit from a wider range of pixels to provide training data, but increasing the search radius does involve a quadratic leap in

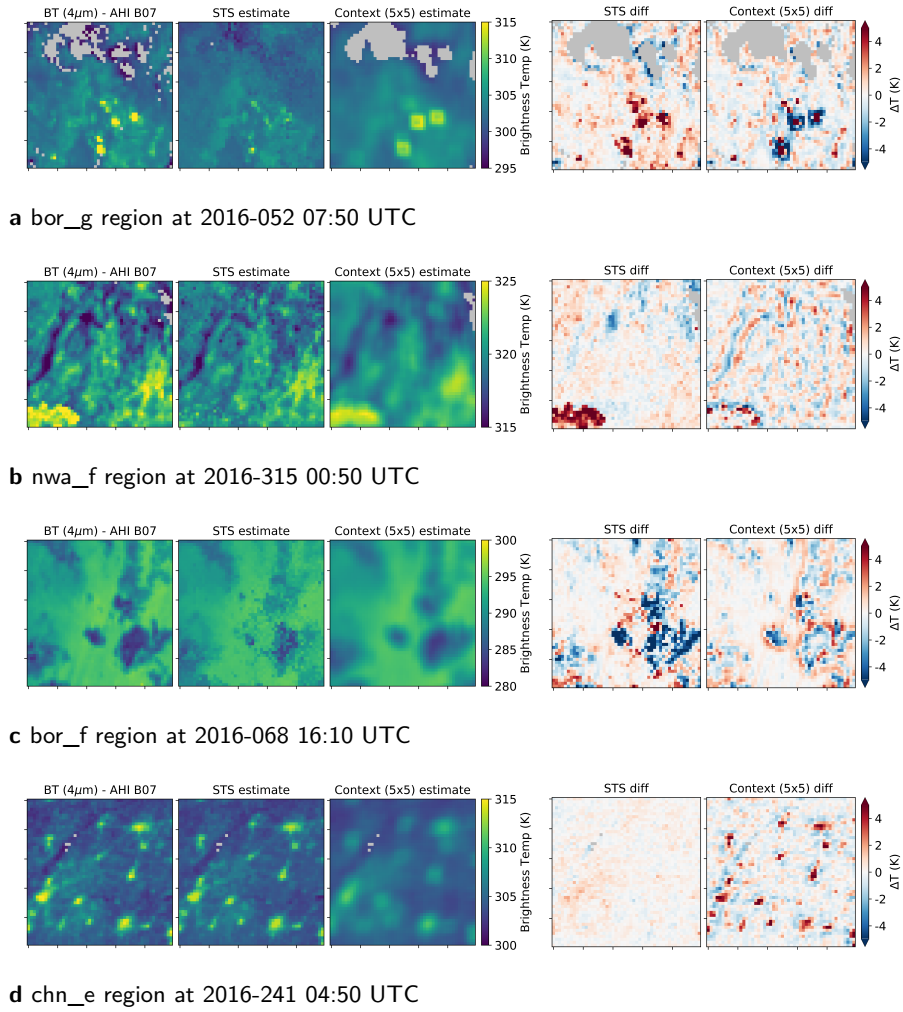


Figure 5.5: Examples of common error in contextual brightness temperature estimation and the results using STS in similar conditions.

processing time. Areas displaying these temperature behaviours may be better served by a concerted classification scheme, in order to identify ideal training pixels outside of the search bounds, without unduly increasing a search radius that is adequate for most solutions.

In general, use of the STS method for estimation results in a 20–30 % improvement in global variation of brightness temperatures, with most notable improvements of up to 40 % in the Thailand sites examined. Given the nature of the STS implementation here, there may be still room for significant improvement in temperature estimation beyond those figures. Susceptibility

to cloud cover remains a major concern, as it does with any method of background estimation. The cloud mask used in this study seems to have trouble with omission errors in these case study areas, which causes erroneous pixel values to slip through and contribute to the training evaluation. These errors can affect both the training data, which will eliminate more accurate training pixels from the selection, and the target pixel itself, with temperature errors caused by cloud hampering the comparison of all pixels to the target. These target errors can lead to mis-selection of training pixels that share the same errors at the same times over more accurate selection candidates. The periods of high cloud cover experienced in the Borneo and Siberian study areas may not explain all of the errors that occurred in these locations, but this extra cloud cover coupled with the poor performance of the cloud mask used in this study suggests that significant improvement in performance in areas such as these may be as simple as the adoption of a more appropriate cloud product.

With regard to the range of contextual variation shown here in comparison to the results reported in [32], selection of the time of day of images analysed is an important consideration – the previous study selected only one time point during the day, whereas predictions in this study are made at four time points over one diurnal cycle. Notwithstanding this, values calculated by context estimators agree well with the previous study apart from decreases in variation in the **jpn** area and increases in variation in the **bor** and **sib** areas.

Comparison sets used in this study have only focused upon areas where brightness temperature from the original AHI image and estimations from both the contextual and STS methods have been available. This omits the portion of pixels that achieve a STS estimate with sufficient robustness that lack a coincident contextual estimation. The threshold set for this study on contextual availability stem from the study by [32], which concluded that 65 % of available adjacent pixels was the minimum for achieving sufficient accuracy in estimation. As the error in STS estimation similarly relates to the number of predictors available and is not adjacency-based, it is possible that some or all of the STS pixel estimations outside the valid contextual area are healthy estimates of temperature.

Notwithstanding the issues with cloud demonstrated in this study, one potential adaptation to the STS that may make cloud less of a factor could be internal implementation of a cloud masking procedure. The STS process in a clear sky state has been shown to give highly accurate background predictions, and has demonstrated the ability to pick out misattributed cloud (see Figure 5.5c). Given enough confidence in estimations from a clear period, anomalous predictor pixels in a prediction set could be flagged as such, with estimations within the same image steering clear of using these marked training pixels. This could

also feed into an updated training image for use in subsequent prediction activity further in time. There are limits to the effectiveness of such a process - a persistent period of bad weather may break the continuity of cloud propagation through the training series, and demonstrated issues with landscape change during the training period may still not be adequately addressed by such an approach. Nevertheless, such an extension if made to work could enhance results coming from this estimation method further.

Further examination of the STS method in action can provide insights into common artifacts seen in contextual temperature estimates, and can highlight problems that need additional attention. Examples such as areas of low temperature anomalies around high temperature anomalies (fig. 5.5a), high temperature anomalies engulfed by neighbouring anomalous values (fig. 5.5b), low temperature anomalies suffering similar treatment (fig. 5.5c), along with standing anomalies such as urban heat islands (fig. 5.5d) demonstrate the instability of using contextual temperature estimation in proximity to anomalous behaviour. Results coming from these types of estimations fail in the situations that they are most heavily relied upon, and fire algorithms have been changed to accept the follies in the contextual estimation process rather than the other way around.

Whilst the simplicity of contextual temperature estimation will ensure its use will continue through many applications, the context method suffers demonstrable flaws in areas of high spatial heterogeneity, with standing anomalies of up to 5 K seen in areas of this study. More sophisticated image reconstruction methods, such as the STS, display much better performance in these areas, and as computational costs and hurdles reduce in size, there should be a willingness to explore use of methods such as these to augment existing temperature estimation practice. Short of embracing other methods, there are simple things that can be done during the implementation of contextual estimation, such as outlier elimination, that can reduce errors and lead to more accurate products. The time of simply applying a smoothing filter to a satellite image and calling it a day should be long gone - we can do better.

5.5 Conclusion

A new MWIR background temperature estimation method has been developed which uses similarities between a target pixels and pixels within a search radius to provide training data for pixel prediction. The Spatio-Temporal Selection (STS) demonstrates an improvement of between 10–40% in brightness temperature estimation over the typically utilised contextual estimation method, with

an overall improvement in the fidelity of image reproduction. The STS method also allows for temperature prediction in areas that are not suitable for contextual estimation, with an increase of between 16–45 % of available estimations, and can act as a pixel estimator in cases where the target pixel is obscured by cloud (12–31 % increase in pixels with valid temperatures over the base images). Whilst the initial implementation of this method focuses upon geostationary sensor imagery, the study demonstrates scope for further investigation into adaption for LEO imagery, and proposes extensions of the method into cloud masking.

5.6 Thesis Context

This chapter introduced the spatio-temporal selection method of brightness temperature estimation, which is based upon selection of training pixels that most resemble the target pixel in the immediate area of the target. The method is an amalgamation of both temporal fitting for the comparison of candidate pixels, and contextual methods, considering that often STS will select candidates from the target pixel context. This chapter highlighted the potential of this method for high fidelity reconstruction of brightness temperature images. The concluding chapter will not only summarise the content of the preceding chapters, but it will also examine theoretical improvements available to STS users and the potential for use in other applications.

Synthesis

6

6.1 Summary of Results

Satellite remote sensing has a long and successful history of use for fire detection purposes, and with changing climatic conditions and the increase of human-induced landscape change, more accurate and timely fire detection is increasingly important. An intrinsic property of many of the fire detection algorithms in common use today is the reliance on knowledge of the unperturbed state of the candidate pixel for detection. As this cannot be measured directly, due to the outsized contribution of fire to the overall radiation budget, methods of estimation of this unperturbed state must be employed in order to accurately identify anomalous behaviour from image pixels. The favoured method by far for achieving this estimation, in most operational products, is use of an estimation based on the brightness temperatures of the pixels immediately surrounding the candidate pixel, in a convolution filtering style process.

The accuracy of contextual-based brightness temperature estimates are restricted by how similar the surrounding landscape is to the candidate pixel. Factors such as the slope and aspect of the landform, land cover and land use heterogeneity, and the presence of water in the surrounding pixels, all contribute to resultant errors in temperature estimation using context. Further perturbing influences such as cloud cover, and in the case of fire the presence of smoke, also influence the potential energy budget attributed to the surrounding pixels and therefore the target estimate. Despite the widespread use of this type of estimation, no previous study has quantified the extent of errors that result from the use of such a method. The central drivers of this thesis study are twofold - first, to quantify the inherent errors pertaining to use of contextual estimation of background brightness temperature; and second, to derive new methods of obtaining background temperature that base their estimates not only upon similarities in location, but upon short-term temporal relationships that each pixel has with potential candidate estimators.

Recent developments in satellite-sensing technologies make this study timely. New geostationary platforms, such as the AHI-8 sensor from which much of this study's data has been supplied (and GOES-16 ABI), have been launched in the last four years carrying sensors that have increased our coverage of the globe and, critically for fire detection, increased our temporal resolution. Coupled with this, the enhanced spatial and radiometric resolutions of these sensors, now provide us with comparable image quality of moderate resolution low earth orbit images from only 10-15 years ago. Never before have we been presented with the earth's diurnal 'heartbeat' in as much detail as is being supplied by these sensors. The current paradigm for brightness temperature estimation is based upon methods that treat images as discrete and independent,

whereas these sensors allow us to do more than that, being able to leverage information not only from the spatial surrounds of the single image, but from the images immediately preceding it in time.

Chapter 1 of this thesis outlines the current fire detection paradigm, and the role brightness temperature estimation plays in describing the existence of fire in the landscape. Chapters 2-5 of this thesis focus on addressing the research objectives outlined in section 1.3, of which the key findings are discussed in this section. This final chapter provides a summary and synthesis of the major research outcomes, along with a roadmap for future extension of the methods described in this work for fire detection and other fields.

Question 1. What is the effect of systematic and structural errors caused by the use of contextual estimation in common fire detection techniques?

Contextually generated background temperature estimation was shown to be a sufficiently accurate method in situations where little to no occlusion of the target context occurs, but this relationship degenerated once more than one third of the target context is occluded. From a full disk analysis of images from the AHI-8 sensor, it was shown that pixels with less than 65 % contextual availability showed a 61 % increase in estimation variability compared to those pixels that suffered no adjacent cloud coverage, with this rate of variation increasing with less available context pixels. Given the accuracies as determined in table 2.5, it was deemed that at least 65 % of context pixels should be available in order to accurately estimate the background temperature of a target pixel from context. This availability percentage trades off accuracy of estimation versus availability of pixel temperature estimates, but even at this rate of acceptance more than one in seven pixels (14.5 %) will not receive a temperature estimate.

This study further showed that the use of expanding contextual windows for background temperature determination, as commonly found in many fire products, is fundamentally flawed. The error of temperature estimation increases by nearly 50 % when increasing the search radius to just 7×7 in situations where no solution was present at 5×5 (table 2.5), and the recovery rate of temperature estimations at higher window sizes following these rules was poor. For instance, when using the aforementioned 65 % context threshold, expanding the window only increased the total count of estimates by 3.1 %, with most of these estimates displaying standard deviations above 3 K. These estimates are too noisy to be of any real use - any smaller fires would be completely eliminated given the tolerances that would need to be applied to have surety in detection of anomalies.

Analysis of contextual estimation also focused upon a number of case study

areas selected to draw upon a wide range of latitude, climate and landform scenarios in order to identify perturbing influences. The sources of variation highlighted in these case study areas were almost as diverse as the areas themselves. Sources of high contextual estimate variation included fire, smoke, snow, seasonal and spatial variations in land cover, urban-rural interfaces, slope and aspect, snow lines, and land-water interfaces. Variability showed significant clustering around landscape features, with the effects heightened when depicted at larger window sizes. Areas with high land cover heterogeneity were particularly susceptible to high estimate variation, and these are generally critical areas for fire detection and monitoring as they are heavily associated with human activity.

Question 2. How can we use the common diurnal variation of upwelling radiation to estimate brightness temperature in a robust fashion?

Harnessing the full disk to fill gaps in a model fitting technique proved an effective tool for generating brightness temperature estimates. The developed method, referred to as the Broad Area Training method, aggregated the median temperature of $0.25^\circ \times 0.25^\circ$ blocks by local solar time into an idealised diurnal model of the latitudinal swath. This diurnal model, which was standardised based upon the mean and standard deviation of each block, was then filtered to remove outliers. Fittings to individual pixels then occurred using this diurnal model as a basis for a single value decomposition over the preceding 24 h before the prediction time.

In comparison to a single pixel fitting method used in [65], the BAT method utilised sufficient data for pixel training in around 86% of cases compared to around 40% for the single pixel method. The BAT method also showed resilience to distortion of the fitting in pixels with increased cloud cover in the preceding diurnal period, with major improvements in estimate error when between 20–50% of the diurnal period experienced cloud cover. The method demonstrated increased fitting viability in coastal areas and areas of south-eastern Australia that were poorly covered by the single pixel fitting technique.

Question 3. How effective is the broad area fitting method at identifying fire-related brightness temperature anomalies in comparison to other fire detection methods?

In comparison to commonly used polar orbit based fire products, BAT derived estimations provided a solid baseline to identify anomalies in geostationary imagery. For situations with fire detections from both the VIIRS and MODIS active fire products, anomalous activity was recorded in 89.3% of cases at a cutoff temperature threshold of 5 K using BAT, with synchronous detec-

tion in almost half of those cases. Considering the much larger pixel size of the geostationary sensor and the ephemeral nature of fire, the omission rate was considerably lower than expected. The commission rate of detection using the BAT method at the 5 K threshold was around 23% in comparison to a burned area product adapted for the study area, which was also lower than expected.

The largest benefit of the BAT method was for first detection. Improvements in detection time when using the BAT fittings on AHI temperature measurements were between 2 h and 6 h, dependent upon the proximity of the time of fitting to the first anomaly detected in the diurnal cycle for the 5 K threshold. This was expected behaviour given the much higher temporal coverage of the geostationary sensor, but was still a promising result in the context of the fitting method and given the coarse spatial resolution. The study showed the viability of using such a method in a near-real-time capacity, as the accuracy of temperature fittings improved when the estimation ended closer to the time of first detection.

Question 4. How can we use similarities in image characteristics over time to improve temperature estimation over a single-image contextual approach?

Using a defined search radius and discrete set of training images, we identified candidate training pixels based upon the magnitude of differences from the target, and then used these training pixels to provide values for pixel estimation in a later image, a technique known as the spatio-temporal selection (STS) method. The STS method demonstrates a significant improvement in temperature estimation accuracy, with impressive reproduction of images in comparison to the raw sensor images. The STS method demonstrates an improvement in estimate variation of between 10–40% in areas with sufficient clear sky for pixel training purposes compared to 5×5 contextual estimation. The STS method also allows for temperature estimates in more locations, including under cloud, with an increase in estimate availability of between 16–45% when compared to contextual estimates.

Results from this study also further highlighted deficiencies in contextual estimates when faced with anomalous temperature behaviour, with STS estimates more readily identifying these anomalies with lower noise and highlighting larger areas of anomalous temperature behaviour compared to the contextual estimates. The STS method also has scope for improvement with regard to candidate pixel selection, with a number of criteria that may be modified to improve the resultant temperature estimation further.

6.2 The Background Behind Background

In hindsight, a study with a focus upon background temperature estimation seems a long way from the grand design of fixing satellite fire detection from geostationary sensors. The study itself has very little to do with the absolute determination of active fire in the environment, partially due to the difficulties presented by validating the ephemeral perturbing influence caused by fire. While conducting this research, with the focus on ever-increasingly complicated algorithms and false alarm detectors, the lack of complexity involved in the background temperature estimate used by almost all of these fire detection techniques was striking. Why of all things was this measure, which seemingly held the key to both finding the fire in the first place, and then in some methods providing estimates of fire size and severity, based upon something as crude as an adjusted smoothing filter? Whilst the practicality of using such a method for estimation, based upon the spatial-autocorrelation principle [82], is unquestionable, especially in the earlier days of remote sensing when simplicity ruled all, the general acceptance in the present of this estimation method despite its obvious flaws drew me in to look deeper.

The first piece of work presented in this thesis attempts to quantify the magnitude of the error associated with the use of such a method. Significantly, the results are critical of the expansion method used to provide data when insufficient adjacent contextual information is available for calculation. This widely accepted method to augment data derived from spatial context doesn't provide satisfactory levels of accuracy for calculation of fire characteristics, and barely adds to the number of pixel estimations anyway, for the price of being heavily computationally intensive. If only one thing emerges from this thesis as a significant contribution to the active fire community, it's that the current paradigm for anomaly detection is flawed in conditions most likely be associated with fire detection - that is, areas that may be cloud covered and undergoing lightning strikes, and areas of potential fire spread that are occluded by smoke. Estimates taken from spatial contextual measurements during these times are a muddled mess. Commonly fire detection techniques (for example the MODIS AF product [26]) have taken a two-pass approach to anomaly identification, identifying potential anomalies using a series of thresholds, and then ruling false detections out using a series of tests. The findings of the work in Chapter 2 strongly suggest that a universal set of thresholds for anomaly identification is the wrong approach - the dynamic range of recorded temperatures in some landscapes is such that just scooping the brightest portion of pixels as a first pass will miss a significant portion of fire activity, and this is reaffirmed in some of the results of the STS study in Chapter 5.

The flaws seen in estimates derived from spatial context drove the investigation into the use of other methods for background temperature estimation, but they were not the sole driver for investigation of the multi-temporal techniques covered in the subsequent chapters. Recent launches of new generation geostationary satellites, such as the Himawari AHI-8 sensor (used predominately in this study) and GOES-16 ABI, have seen a general improvement in the level of spatial, temporal and radiometric resolution available from geostationary imagery. A real opportunity presented to use the data from such sensors in new and exciting ways, and to potentially break away from the flawed paradigm of contextual estimation by using time as a dimension in anomaly detection. This being said, adopting a new sensor and new workflows also poses a series of challenges to efficient and effective research outcomes. Our greatest challenge with using AHI-8 was the lack of validated cloud mask information, which is of vital importance to successful isolation of negative anomalies in any group of estimation values. We found mixed results with a number of unpublished masking products, and had trouble understanding the decisions made in the development of some others that had been applied to this sensor. For instance, the AHI-8 work published by [94] applies a cloud mask based upon a series of absolute thresholds developed for other sensors, with a high omission rate reported for use on the GOES sensor [95]. This omission rate, when reconciled with the percentage of total cloud cover reported in Chapter 2, suggests that large numbers of cloud pixels are making it past this masking process, with potentially 15% of all unflagged pixels potentially being cloud affected in some way. In spite of this, a lack of other options drove us to the use of such a mask, with the acceptance that the errors associated with it could heavily influence the subsequent accuracy of derived estimates. Our concerns about cloud masking drove the adoption of the BAT process, which looked to provide a multi-temporal solution to temperature estimation without the reliance on cloud masking.

The BAT method creates a new way of determining training data for a signal fitting process - in no other work have I found reference to the use of aggregated broad area measurements to provide a time-corrected signal. Maybe no one else was crazy enough to look at it - figuring out how to reconcile local solar time with image recording intervals and sensor scan offsets to provide the standardised signal was the result of many months work. Whilst in this case the BAT was used to isolate a relatively high frequency signal (diurnal variation), there is potential for its use to map seasonal and annual trends in data sets and subsequently isolate changes. Application in wavelengths other than the MWIR has great potential for a method such as this, to map changes in sea surface temperature and plant phenology for instance. Of course, there

are obvious improvements that should be made if these applications are to come to pass. The fitting method, which used a principal component analysis to derive the fitting trajectory over the prediction vector, didn't take into account our knowledge of what the resultant signal should look like enough. The very strong diurnal signal found in the standard model was often unable to be reconciled in cloud-affected prediction vectors, even though we knew what the resultant fitting should look like. Reconfiguring the BAT method of signal derivation to encompass both external cloud mask data, and to weight the strong diurnal signal more heavily, would vastly improve what is already an impressive use of the volumes of geostationary sensor data available for such a task.

The application of BAT fittings to the task of anomaly identification was the topic of Chapter 4 of the thesis, and whilst the results derived from this study showed improvements in detection times especially, the strength of these results vindicates the choices made in the design of the sensor used, rather than the BAT method. Given that the relevant information with regard to isolation of potential thermal anomalies identified in Chapter 2 was not available at that time (due to issues with cloud masking products), a series of thresholds was examined for identifying thermal anomalies, instead of tailored thresholds based upon a given study area. The case study area used in Chapter 4 very closely overlapped that of the north western Australia region covered in the contextual analysis study (Chapter 2), and the standard deviations reported in section 2.3.3 probably suggest that the thresholds used in this study were too aggressive to definitively categorise fire activity. Nevertheless, the study showed the practical effectiveness of using such a fitting technique to isolate anomalies, notwithstanding the potential improvements that could be made as mentioned earlier in this section. The work also highlighted the importance of proper analysis of data sources such as the AHI-8 sensor - the vast improvements in detection time seen in this study would have been just as apparent in a spatial contextual study of the same fire activity, and the study provided an opportunity to be measured favourably against active fire products in common usage.

6.3 A New Direction for Real-Time Anomaly Detection

A particularly exciting development in the process of completing this thesis has been the implementation of the spatio-temporal selection method (Chapter 5) and the results derived from it. This method, which uses selections of candidate pixels based upon a set of training images in the recent history

before the prediction image, provides images of estimated pixel values with a striking resemblance to the sensor recorded prediction images, certainly in comparison to the contextual images presented alongside. The global estimate of error for most sites is an improvement of around 15-20% based upon the standard deviation of differences from the target value. The true potential in this method comes from its flexibility in the choice of training criteria, which is only briefly touched upon in the thesis chapter it was introduced in.

One of the common themes of this study is the challenge imposed by cloud-based occlusion and its effect upon accurate temperature reproduction. The STS method demonstrates a reasonably accurate method of temperature estimation, which would be expected to further improve with the addition of a quality cloud mask. The potential lies in the STS method to circumvent the need for an externally provided cloud mask. If enough clear sky ground temperatures are available at a particular time point, the accuracy of the training estimators relative to the target pixel can reveal information about potential anomalous temperature readings in the training pixels. Given enough variety in training pixels in comparison to target pixels (and with an increase in the value of n_{train}), outliers in the training stacks can be isolated, and with an ensemble aggregation process based upon the suspected confidence in each training pixel, a map of potential anomalies could be created based upon this. Dependent upon how the model is initialised, anomalies on the fringe of clear sky areas could be flagged as such to assist in both eliminating erroneous training values for subsequent estimates, and to identify other anomalous values. In this way, this the cloud (and fire) isolation process could act in lieu of an explicit external cloud mask, which may well result in further improvements to method accuracy.

Another unexplored facet of STS is the ability to preserve predictor pixels over time. To a certain extent, the training phase of STS is independent of the prediction phase. The exact amount of time that selected training pixels are valid as estimators is unknown, but the likelihood is that major changes in the location of training pixels would only occur where there is a major change in the composition of land cover in either the target pixel or the training pixel. The magnitude of this change is also of importance of course, but based upon the the relatively coarse spatial scale that has been examined with imagery in this study (~ 2 km), disturbances such as fire, land clearing and flooding are the only factors likely to change the temporal signature of pixels over time. In an area of relative temperature stability, this means we may be able to preserve a set of training pixel locations from a good solution to fill in gaps in estimations when the training data set is severely affected by cloud occlusion.

A more complex task, but not out of the realms of possibility, is the appli-

cation of the STS method to non-continuous time series imagery, such as that from polar orbiting sensors. Considering the STS provides only the location of suitable estimator pixels, and carries forward no information about pixel properties for fitting purposes, these locations should be able to be applied to individual captures for prediction purposes, even with image sets that lack a complete diurnal definition. A major stumbling point with this adaptation may be discrete pixel locations, which are readily available and regularly repeated in geostationary imagery, but vary wildly based upon orbital geometries in the case of polar orbiters. This could be counteracted by the use of confidence measures for the given training locations, which could then be applied in a nearest neighbour estimation or something more complex, keeping in mind that in most cases the available statistical strength of temperature estimates is high.

This study has focused upon the estimation of surface temperatures in the context of fire detection, but at their heart these estimates provide a baseline for testing all sorts of anomalous behaviour in satellite imagery, and mapping short and long-term change based upon this type of anomaly identification. Results from chapter 5 showed the potential to highlight change due to burnt area in the training set of certain pixels. Such techniques could be used to highlight land cover change due to clearing, rainfall events and localised flooding, and in longer term work, urban encroachment and heat island effects. Due to this study focusing on fire, use of techniques such as STS has only been applied using the MWIR, and these other types of change may be more easily identified when these techniques are applied to other wavelengths. Whilst the methods described in this thesis are not by any means the panacea for all change detection related issues in remote sensing, there is potential for some parts of this work to be extended to provide potential solutions to a range of different problems.

The technological side of remote sensing continues to push forward, with more advanced sensors and greater volumes of data available from platforms both in polar and geostationary orbits. With this proliferation of satellite information, especially in the temporal space for geostationary imagery, we possess the capability to monitor and measure change with a greater capacity than ever before. Polar orbiting sensors remain the gold standard for research in remote sensing due to their superior spatial characteristics, but their placement in sun-synchronous positions to aid visible light measurements has kept research from properly tackling the role of time in prediction techniques. With satellite launching becoming cheaper than ever, there will soon be a demand for the spatial resolution of products we see from polar orbiters beyond the one or two daily images we currently enjoy. As the current constellation of polar-orbiters

increases the temporal density of observations from these sources will increase concurrently. The applicability of multi-temporal techniques outlined in this thesis that are effective with geostationary sensors will soon be applicable to data obtained from an increased number of polar orbiting sensors.

Ultimately the best outcome of this research would be the integration of these techniques in a data assimilation framework utilising imagery from both polar and geostationary orbiters, as long as the challenges of scaling between data sources and coverage variability can be overcome. The need for more research in this fire detection space remains, and contributions made by this body of work will hopefully inspire future solutions.

References

- [1] J. K. Balch, B. A. Bradley, J. T. Abatzoglou, R. C. Nagy, E. J. Fusco, and A. L. Mahood. Human-started wildfires expand the fire niche across the United States. *Proceedings of the National Academy of Sciences*, 114(11):2946–2951, mar 2017.
- [2] B. Bhatta, S. Saraswati, and D. Bandyopadhyay. Urban sprawl measurement from remote sensing data. *Applied Geography*, 30(4):731–740, 2010.
- [3] M. Black and A. Jepson. Eigenttracking: Robust matching and tracking of articulated objects using a view-based representation. *International Journal of Computer Vision*, 26(1):63–84, 1998.
- [4] D. M. J. S. Bowman, J. K. Balch, P. Artaxo, W. J. Bond, J. M. Carlson, M. A. Cochrane, C. M. D’Antonio, R. S. Defries, J. C. Doyle, S. P. Harrison, F. H. Johnston, J. E. Keeley, M. A. Krawchuk, C. A. Kull, J. B. Marston, M. A. Moritz, I. C. Prentice, C. I. Roos, A. C. Scott, T. W. Swetnam, G. R. van der Werf, and S. J. Pyne. Fire in the Earth system. *Science*, 324(5926):481–4, 2009.
- [5] D. S. Boyd, G. M. Foody, and P. J. Curran. The relationship between the biomass of cameroonian tropical forests and radiation reflected in middle infrared wavelengths (3.0-5.0 μ m). *International Journal of Remote Sensing*, 20(5):1017–1023, 1999.
- [6] D. Brunner, G. Lemoine, and L. Bruzzone. Earthquake damage assessment of buildings using VHR optical and SAR imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 48(5):2403–2420, 2010.
- [7] A. Calle and J. L. Casanova. Forest Fires And Remote Sensing. In H. G. Coskun, H. K. Cigizoglu, and M. D. Maktav, editors, *Integration of Information for Environmental Security*, pages 247–290. Springer Netherlands, Dordrecht, 2008.
- [8] A. Calle, J. L. Casanova, and F. González-Alonso. Impact of point spread function of MSG-SEVIRI on active fire detection. *International Journal of Remote Sensing*, 30(17):4567–4579, aug 2009.
- [9] A. Calle, J. L. Casanova, and A. Romo. Fire detection and monitoring using MSG Spinning Enhanced Visible and Infrared Imager (SEVIRI) data. *Journal of Geophysical Research: Biogeosciences*, 111(G4):n/a–n/a, dec 2006.
- [10] A. Carvalho, M. D. Flannigan, K. A. Logan, L. M. Gowman, A. I. Miranda, and C. Borrego. The impact of spatial resolution on area burned and fire

- occurrence projections in Portugal under climate change. *Climatic Change*, 98(1-2):177–197, jan 2010.
- [11] J. Dozier. A method for satellite identification of surface temperature fields of subpixel resolution. *Remote Sensing of Environment*, 11(C):221–229, 1981.
 - [12] A. C. Edwards, S. W. Maier, L. B. Hutley, R. J. Williams, and J. Russell-Smith. Spectral analysis of fire severity in north Australian tropical savannas. *Remote Sensing of Environment*, 136:56–65, sep 2013.
 - [13] C. Filizzola, R. Corrado, F. Marchese, G. Mazzeo, R. Paciello, N. Pergola, and V. Tramutoli. RST-FIRES, an exportable algorithm for early-fire detection and monitoring: description, implementation, and field validation in the case of the MSG-SEVIRI sensor. *Remote Sensing of Environment*, 186:196–216, dec 2016.
 - [14] M. D. Flannigan, B. M. Wotton, G. A. Marshall, W. J. de Groot, J. Johnston, N. Jurko, and A. S. Cantin. Fuel moisture sensitivity to temperature and precipitation: climate change implications. *Climatic Change*, 134(1-2):59–71, jan 2016.
 - [15] S. P. Flasse and P. Cecatto. A contextual algorithm for AVHRR fire detection. *International Journal of Remote Sensing*, 17(2):419–424, jan 1996.
 - [16] R. Frouin and C. Gautier. Calibration of NOAA-7 AVHRR, GOES-5, and GOES-6 VISSR/VAS solar channels. *Remote Sensing of Environment*, 22(1):73–101, jun 1987.
 - [17] C. Gao, X. Jiang, Z.-l. Li, and F. Nerry. Comparison of the Thermal Sensors of SEVIRI and MODIS for LST Mapping. In *Thermal Infrared Remote Sensing*, volume 17, pages 233–252. Springer, 2013.
 - [18] L. Giglio. Characterization of the tropical diurnal fire cycle using VIRS and MODIS observations. *Remote Sensing of Environment*, 108(4):407–421, 2007.
 - [19] L. Giglio, I. Csiszar, A. Restas, J. T. Morisette, W. Schroeder, D. Morton, C. O. Justice, Á. Restas, J. T. Morisette, W. Schroeder, D. Morton, and C. O. Justice. Active fire detection and characterization with the advanced spaceborne thermal emission and reflection radiometer (ASTER). *Remote Sensing of Environment*, 112(6):3055–3063, jun 2008.
 - [20] L. Giglio, J. Descloitres, C. O. Justice, and Y. J. Kaufman. An enhanced contextual fire detection algorithm for MODIS. *Remote Sensing of Environment*, 87(2-3):273–282, oct 2003.

- [21] L. Giglio and C. O. Justice. Effect of wavelength selection on characterization of fire size and temperature. *International Journal of Remote Sensing*, 24(17):3515–3520, 2003.
- [22] L. Giglio and J. D. Kendall. Application of the Dozier retrieval to wildfire characterization a sensitivity analysis. *Remote Sensing of Environment*, 77(1):34–49, 2001.
- [23] L. Giglio, J. D. Kendall, and C. O. Justice. Evaluation of global fire detection algorithms using simulated AVHRR infrared data. *Int. J. Remote Sensing*, 20(10):1947–1985, 1999.
- [24] L. Giglio, J. T. Randerson, and G. R. van der Werf. Analysis of daily, monthly, and annual burned area using the fourth-generation global fire emissions database (GFED4). *Journal of Geophysical Research: Biogeosciences*, 118(1):317–328, mar 2013.
- [25] L. Giglio and W. Schroeder. A global feasibility assessment of the bi-spectral fire temperature and area retrieval using MODIS data. *Remote Sensing of Environment*, 152:166–173, 2014.
- [26] L. Giglio, W. Schroeder, and C. O. Justice. The collection 6 MODIS active fire detection algorithm and fire products. *Remote Sensing of Environment*, 178:31–41, jun 2016.
- [27] F. M. Göttsche and F. S. Olesen. Modelling of diurnal cycles of brightness temperature extracted from METEOSAT data. *Remote Sensing of Environment*, 76(3):337–348, jun 2001.
- [28] Y. Govaerts, M. J. Wooster, A. Lattanzio, and G. Roberts. MSG SEVIRI Fire Radiative Power (FRP) Characterisation Algorithm Theoretical Basis Document. Technical Report January 2010, EUMETSAT, 2010.
- [29] M. Griffin, H.-h. Burke, and J. Kerekes. Radiative transfer in the midwave infrared applicable to full spectrum atmospheric characterization. In *IEEE International IEEE International IEEE International Geoscience and Remote Sensing Symposium, 2004. IGARSS '04. Proceedings. 2004*, volume 6, pages 4191–4194. IEEE, 2004.
- [30] B. Hally, L. Wallace, K. Reinke, and S. Jones. Assessment of the Utility of the Advanced Himawari Imager to Detect Active Fire over Australia. In *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, volume XLI-B8, pages 65–71, jun 2016.
- [31] B. Hally, L. Wallace, K. Reinke, and S. Jones. A Broad-Area Method for the Diurnal Characterisation of Upwelling Medium Wave Infrared Radiation. *Remote Sensing*, 9(2):167, feb 2017.

- [32] B. Hally, L. Wallace, K. Reinke, S. Jones, C. Engel, and A. Skidmore. Estimating Fire Background Temperature at a Geostationary Scale—An Evaluation of Contextual Methods for AHI-8. *Remote Sensing*, 10(9), 2018.
- [33] S. Hantson, M. Padilla, D. Corti, and E. Chuvieco. Strengths and weaknesses of MODIS hotspots to characterize global fire occurrence. *Remote Sensing of Environment*, 131:152–159, 2013.
- [34] L. He and Z. Li. Enhancement of a fire detection algorithm by eliminating solar reflection in the mid-IR band: application to AVHRR data. *International Journal of Remote Sensing*, 33(22):7047–7059, nov 2012.
- [35] A. K. Heidinger. CLAVR-x cloud mask algorithm theoretical basis document (atbd). Technical report, University of Wisconsin, Madison, WI, 2004.
- [36] S. Hislop, S. Jones, M. Soto-Berelov, A. Skidmore, A. Haywood, and T. H. Nguyen. A fusion approach to forest disturbance mapping using time series ensemble techniques. *Remote Sensing of Environment*, 221(November 2018):188–197, 2019.
- [37] Japan Meteorological Agency. Himawari Standard Data User’s Guide. Technical report, JMA, Tokyo, 2015.
- [38] G. M. Jiang, Z. L. Li, and F. Nerry. Land surface emissivity retrieval from combined mid-infrared and thermal infrared data of MSG-SEVIRI. *Remote Sensing of Environment*, 105(4):326–340, 2006.
- [39] S. Jones, K. Reinke, B. Hally, C. Wickramasinghe, and L. Wallace. Large area validation of Himawari-8 fire active fire products. In *38th Asian Conference on Remote Sensing - Space Applications: Touching Human Lives, ACRS 2017*, 2017.
- [40] C. O. Justice and P. Dowty. IGBP-DIS satellite fire detection algorithm workshop technical report. *IGBP-DIS working paper*, 9:88, 1994.
- [41] C. O. Justice, L. Giglio, S. Korontzi, J. Owens, J. T. Morisette, D. P. Roy, J. Descloitres, S. Alleaume, F. Petitcolin, and Y. J. Kaufman. The MODIS fire products. *Remote Sensing of Environment*, 83(1-2):244–262, 2002.
- [42] C. O. Justice, J. D. Kendall, P. R. Dowty, and R. J. Scholes. Satellite remote sensing of fires during the SAFARI campaign using NOAA Advanced Very High Resolution Radiometer data. *Journal of Geophysical Research: Atmospheres*, 101(D19):23851–23863, oct 1996.
- [43] Y. J. Kaufman, C. O. Justice, L. P. Flynn, J. D. Kendall, E. M. Prins, L. Giglio, D. E. Ward, W. P. Menzel, and A. W. Setzer. Potential global fire monitoring from EOS-MODIS. *Journal of Geophysical Research-Atmospheres*, 103(D24):32215–32238, 1998.

- [44] J. E. Keeley and A. D. Syphard. Historical patterns of wildfire ignition sources in California ecosystems. *International Journal of Wildland Fire*, 27(12):781, 2018.
- [45] A. Koltunov, S. L. Ustin, and E. M. Prins. On timeliness and accuracy of wildfire detection by the GOES WF-ABBA algorithm over California during the 2006 fire season. *Remote Sensing of Environment*, 127:194–209, 2012.
- [46] A. Koltunov, S. L. Ustin, B. Quayle, B. Schwind, V. G. Ambrosia, and W. Li. The development and first validation of the GOES Early Fire Detection (GOES-EFD) algorithm. *Remote Sensing of Environment*, 184:436–453, oct 2016.
- [47] M. A. Krawchuk, S. G. Cumming, and M. D. Flannigan. Predicted changes in fire weather suggest increases in lightning fire initiation and future area burned in the mixedwood boreal forest. *Climatic Change*, 92(1-2):83–97, jan 2009.
- [48] S. Langaas. A parametrised bispectral model for savanna fire detection using AVHRR night images. *International Journal of Remote Sensing*, 14(12):2245–2262, 1993.
- [49] Z. Li, S. Nadon, and J. Cihlar. Satellite-based detection of Canadian boreal forest fires: Development and application of the algorithm. *International Journal of Remote Sensing*, 21(16):3057–3069, 2000.
- [50] Z. L. Li, B. H. Tang, H. Wu, H. Ren, G. Yan, Z. Wan, I. F. Trigo, and J. A. Sobrino. Satellite-derived land surface temperature: Current status and perspectives. *Remote Sensing of Environment*, 131:14–37, 2013.
- [51] Y. Liu, J. Stanturf, and S. Goodrick. Trends in global wildfire potential in a changing climate. *Forest Ecology and Management*, 259(4):685–697, feb 2010.
- [52] L. Lu, V. Venus, A. Skidmore, T. Wang, and G. Luo. Estimating land-surface temperature under clouds using MSG/SEVIRI observations. *International Journal of Applied Earth Observation and Geoinformation*, 13(2):265–276, apr 2011.
- [53] S. W. Maier. Changes in surface reflectance from wildfires on the Australian continent measured by MODIS. *International Journal of Remote Sensing*, 31(12):3161–3176, 2010.
- [54] L. Majewski. Personal Communication. (ABOM, Melbourne, Vic, Australia), personal communication, sep 2016.

- [55] J.-F. Mas. Monitoring land-cover changes: A comparison of change detection techniques. *International Journal of Remote Sensing*, 20(1):139–152, jan 1999.
- [56] G. Masiello, C. Serio, I. De Feis, M. Amoroso, S. Venafrà, I. F. Trigo, and P. Watts. Kalman filter physical retrieval of surface emissivity and temperature from geostationary infrared radiances. *Atmospheric Measurement Techniques*, 6(12):3613–3634, 2013.
- [57] M. Matson, G. Stephens, and J. Robinson. Fire detection using data from the NOAA-N satellites. *International Journal of Remote Sensing*, 8(7):961–970, jul 1987.
- [58] K. Meyer, S. Platnick, L. Oreopoulos, and D. Lee. Estimating the direct radiative effect of absorbing aerosols overlying marine boundary layer clouds in the southeast Atlantic using MODIS and CALIOP. *Journal of Geophysical Research: Atmospheres*, 118(10):4801–4815, may 2013.
- [59] L. Na, J. Zhang, Y. Bao, Y. Bao, R. Na, S. Tong, and A. Si. Himawari-8 Satellite Based Dynamic Monitoring of Grassland Fire in China-Mongolia Border Regions. *Sensors*, 18(1):276, jan 2018.
- [60] A. Okuyama, A. Andou, K. Date, K. Hoasaka, N. Mori, H. Murata, T. Tabata, M. Takahashi, R. Yoshino, and K. Bessho. Preliminary validation of Himawari-8/AHI navigation and calibration. In *SPIE*, volume 9607, page 96072E, 2015.
- [61] P. Oliva and W. Schroeder. Assessment of VIIRS 375m active fire detection product for direct burned area mapping. *Remote Sensing of Environment*, 160:144–155, 2015.
- [62] D. Peterson, J. Wang, C. Ichoku, E. J. Hyer, and V. Ambrosia. A sub-pixel-based calculation of fire radiative power from MODIS observations: 1. Algorithm development and initial assessment. *Remote Sensing of Environment*, 129:262–279, 2013.
- [63] E. M. Prins and W. P. Menzel. Geostationary satellite detection of biomass burning in South America. *International Journal of Remote Sensing*, 13(15):2783–2799, oct 1992.
- [64] A. Rajaraman and J. D. Ullman. *Mining of Massive Datasets*. Cambridge University Press, 2nd edition, 2011.
- [65] G. Roberts and M. J. Wooster. Development of a multi-temporal Kalman filter approach to geostationary active fire detection & fire radiative power (FRP) estimation. *Remote Sensing of Environment*, 152:392–412, 2014.

- [66] G. J. Roberts and M. J. Wooster. Fire detection and fire characterization over Africa using meteosat SEVIRI. *IEEE Transactions on Geoscience and Remote Sensing*, 46(4):1200–1218, 2008.
- [67] J. M. Robinson. Fire from space: Global fire evaluation using infrared remote sensing. *International Journal of Remote Sensing*, 12(1):3–24, jan 1991.
- [68] D. P. Roy, L. Boschetti, and A. M. Smith. Satellite Remote Sensing of Fires. In *Fire Phenomena and the Earth System: An Interdisciplinary Guide to Fire Science*, pages 77–93. Wiley-Blackwell, 2013.
- [69] O. Rozenstein, Z. Qin, Y. Derimian, and A. Karnieli. Derivation of land surface temperature for landsat-8 TIRS using a split window algorithm. *Sensors (Switzerland)*, 14(4):5768–5780, 2014.
- [70] S. W. Running. CLIMATE CHANGE: Is Global Warming Causing More, Larger Wildfires? *Science*, 313(5789):927–928, aug 2006.
- [71] J. Russell-Smith, C. P. Yates, P. J. Whitehead, R. Smith, R. Craig, G. E. Allan, R. Thackway, I. Frakes, S. Cridland, M. C. P. Meyer, and A. M. Gill. Bush-fires 'down under': patterns and implications of contemporary Australian landscape burning. *International Journal of Wildland Fire*, 16(4):361, 2007.
- [72] T. J. Schmit, P. Griffith, M. M. Gunshor, J. M. Daniels, S. J. Goodman, and W. J. Lebar. A Closer Look at the ABI on the GOES-R Series. *Bulletin of the American Meteorological Society*, 98(4):681–698, apr 2017.
- [73] W. Schroeder, I. Csiszar, L. Giglio, and C. C. Schmidt. On the use of fire radiative power, area, and temperature estimates to characterize biomass burning via moderate to coarse spatial resolution remote sensing data in the Brazilian Amazon. *Journal of Geophysical Research Atmospheres*, 115(21):1–10, 2010.
- [74] W. Schroeder, E. Ellicott, C. Ichoku, L. Ellison, M. B. Dickinson, R. D. Ottmar, C. Clements, D. K. Hall, V. Ambrosia, and R. Kremens. Integrated active fire retrievals and biomass burning emissions using complementary near-coincident ground, airborne and spaceborne sensor data. *Remote Sensing of Environment*, 140:719–730, 2014.
- [75] W. Schroeder, P. Oliva, L. Giglio, and I. Csiszar. The New VIIRS 375m active fire detection data product: Algorithm description and initial assessment. *Remote Sensing of Environment*, 143:85–96, 2014.
- [76] W. Schroeder, P. Oliva, L. Giglio, B. Quayle, E. Lorenz, and F. Morelli. Active fire detection using Landsat-8/OLI data. *Remote Sensing of Environment*, 185:210–220, nov 2016.

- [77] W. Schroeder, E. Prins, L. Giglio, I. Csiszar, C. Schmidt, J. Morisette, and D. Morton. Validation of GOES and MODIS active fire detection products using ASTER and ETM+ data. *Remote Sensing of Environment*, 112(5):2711–2726, 2008.
- [78] W. Schroeder, M. Ruminski, I. Csiszar, L. Giglio, E. Prins, C. Schmidt, and J. Morisette. Validation analyses of an operational fire monitoring product: The Hazard Mapping System. *International Journal of Remote Sensing*, 29(20):6059–6066, 2008.
- [79] C. Stephenson, J. Handmer, and R. Betts. Estimating the economic, social and environmental impacts of wildfires in Australia. *Environmental Hazards*, 12(2):93–111, jun 2013.
- [80] B. Teague, R. McLeod, and S. Pascoe. *2009 Victorian Bushfires Royal Commission - Final Report Summary*. Number 332. Parliament of Victoria, 2010.
- [81] D. Thomas, D. Butry, S. Gilbert, D. Webb, and J. Fung. The costs and losses of wildfires: a literature survey. Technical report, National Institute of Standards and Technology, Gaithersburg, MD, nov 2017.
- [82] W. R. Tobler. A computer movie simulating urban growth in the Detroit region. *Economic geography*, 46(sup1):234–240, 1970.
- [83] G. Udahemuka and F. V. D. Bergh. Robust fitting of diurnal brightness temperature cycle. *South African Computer Journal*, 40:1–6, 2008.
- [84] F. van den Bergh and P. E. Frost. A multi temporal approach to fire detection using MSG data. In *Proceedings of the Third International Workshop on the Analysis of Multi-Temporal Remote Sensing Images 2005*, volume 2005, pages 156–160, 2005.
- [85] F. van den Bergh, G. Udahemuka, and B. J. van Wyk. Potential fire detection based on Kalman-driven change detection. In *International Geoscience and Remote Sensing Symposium (IGARSS)*, volume 4, pages 77–80, 2009.
- [86] F. van den Bergh, M. A. van Wyk, and B. J. van Wyk. Comparison of data-driven and model-driven approaches to brightness temperature diurnal cycle interpolation. In *17th Annual Symposium of the Pattern Recognition Association of South Africa*, Parys, South Africa, 2006.
- [87] G. R. van der Werf. Continental-Scale Partitioning of Fire Emissions During the 1997 to 2001 El Nino/La Nina Period. *Science*, 303(5654):73–76, jan 2004.

- [88] C. Wickramasinghe, S. Jones, K. Reinke, and L. Wallace. Development of a Multi-Spatial Resolution Approach to the Surveillance of Active Fire Lines Using Himawari-8. *Remote Sensing*, 8(11):932, nov 2016.
- [89] M. J. Wooster, G. Roberts, P. H. Freeborn, W. Xu, Y. Govaerts, R. Beeby, J. He, A. Lattanzio, D. Fisher, and R. Mullen. LSA SAF Meteosat FRP products – Part 1: Algorithms, product contents, and analysis. *Atmospheric Chemistry and Physics*, 15(22):13217–13239, nov 2015.
- [90] M. J. Wooster, G. Roberts, G. L. W. Perry, and Y. J. Kaufman. Retrieval of biomass combustion rates and totals from fire radiative power observations: FRP derivation and calibration relationships between biomass consumption and fire radiative energy release. *Journal of Geophysical Research*, 110(D24):D24311, 2005.
- [91] M. J. Wooster, W. Xu, and T. Nightingale. Sentinel-3 SLSTR active fire detection and FRP product: Pre-launch algorithm development and performance evaluation using MODIS and ASTER datasets. *Remote Sensing of Environment*, 120(2012):236–254, 2012.
- [92] X. Xiong, J. Butler, K. Chiang, B. Efremova, J. Fulbright, N. Lei, J. McIntire, H. Oudrari, J. Sun, Z. Wang, and A. Wu. VIIRS on-orbit calibration methodology and performance. *Journal of Geophysical Research: Atmospheres*, 119(9):5065–5078, may 2014.
- [93] H. Xu, Y. Yu, D. Tarpley, F. M. Göttsche, and F. S. Olesen. Evaluation of GOES-R Land Surface Temperature Algorithm Using SEVIRI Satellite Retrievals with In-Situ Measurements. *IEEE Transactions on Geoscience and Remote Sensing*, 52(7):3812–3822, jul 2014.
- [94] W. Xu, M. J. Wooster, T. Kaneko, J. He, T. Zhang, and D. Fisher. Major advances in geostationary fire radiative power (FRP) retrieval over Asia and Australia stemming from use of Himawari-8 AHI. *Remote Sensing of Environment*, 193:138–149, may 2017.
- [95] W. Xu, M. J. Wooster, G. Roberts, and P. H. Freeborn. New GOES imager algorithms for cloud and active fire detection and fire radiative power assessment across North, South and Central America. *Remote Sensing of Environment*, 114(9):1876–1895, sep 2010.
- [96] Zhengming Wan and Zhao-Liang Li. A physics-based algorithm for retrieving land-surface emissivity and temperature from EOS/MODIS data. *IEEE Transactions on Geoscience and Remote Sensing*, 35(4):980–996, jul 1997.

Summary

Fire is an integral catalyst for change and regeneration in the environment, along with being a major impact on social and economic activity. In evolving climatic conditions, wildfires are increasing both in intensity and in area impacted in recent years. Remote sensing has been used for many decades to provide insight into fire activity and impact through the use of infrared imagery for active fire detection. Electromagnetic wavelengths at around $4\mu\text{m}$ are particularly sensitive to fire activity, in comparison to nominal conditions which incorporate both solar reflection and thermal emission of the earth's surface. Brightness temperature measurements in these wavelengths isolating such radiative anomalies enable fire detection and attribution from satellite sensors. The processes involved in providing these active fire products require an accurate estimation of the background brightness temperature of the area in question without the influence of fire. These commonly used active fire products generally use a contextual-based estimate to provide this background temperature information. Whilst this estimation technique is widely accepted for use, especially in single time-point polar orbiting products, the introduction of new generation geostationary sensors provides substantial improvements to knowledge about the earth's surface, especially with regard to diurnal variation. These sensors provide an opportunity to not only evaluate the accuracy of context-based estimation of brightness temperature, but to integrate the rich spatio-temporal information provided by such sensors to improve the accuracy and availability of background estimates.

In order to determine an adequate level of accuracy required for the derivation of new temperature estimators, it was important to know the accuracy of the current paradigm of contextual temperature estimation. To date, whilst contextual temperature estimation is widely used, no definitive study of the expected error in temperature estimates had been completed. An analysis of the error involved in contextual estimation was conducted upon medium wave infrared radiation (MWIR) images taken by the AHI-8 sensor onboard the Himawari-8 geostationary satellite. Comparisons were made between contextual estimates and the raw brightness temperature observations over the AHI-8 full disk for 36 images at 0500 UTC across 2016, and across a number of case study areas for 31 days of images surrounding the peak fire period as determined by the VIIRS active fire product. The study found that variation in temperature estimations from context had negligible bias and standard deviations around 1.1 K when the surrounding 5×5 area was clear of cloud, which occurred in 53.9% of cases. Accuracy diminished as the contextual estimation surface was obscured, such that pixels with 65% or more context available experienced

a 56% increase in estimate variation. The common practice of window expansion saw the variation of estimates increase substantially, with 7×7 windows resulting in a 44% increase in variation over the 5×5 results. The study concluded that 5×5 contextual estimation should be limited to using values where at least 65% of the contextual surface is available, with no expansion of the contextual window due to the detrimental effect on estimation accuracy. This resulted in 1 in 7 non-obscured pixels (14.5%) in the examined images not having accurate contextual estimates available. Common causes of increased contextual inaccuracy included coastlines, land cover changes, slope and aspect of the surface, urban heat signatures and land inundation, with these effects highlighted in the examined case studies.

With the identification of inadequacies in the contextual estimation method, investigation of methods to leverage the temporal domain of the geostationary sensor to fill these gaps was undertaken. Particular focus was placed upon the modelling of the diurnal temperature variation of locations, and in particular how gaps in the training data for such models could be filled. Previous studies that had used diurnal modelling in typically cloud obscured areas had identified deficiencies in the use of single pixel data for creating temperature models. A new technique was developed using a standardised model of diurnal temperature variation based upon the latitude of the examined area, and corrected for local solar time. Results from models created by this method, known the Broad Area Training (BAT) method, were compared to a single-pixel derived model and the raw temperature recordings from AHI-8. The comparison found that the RMS error of the BAT-derived models maintained sufficient accuracy for temperature estimation with up to half the estimation days' values obscured by cloud, with errors reduced by more than 50% compared to the single pixel method with between 30 – 70 cloud affected images present in the day of estimation. The method also increased the availability of training data for modelling using this type of multi-temporal method, with up to 90% of pixels across the Australian continent possessing sufficient training data for estimation, in comparison to 40% for the single-pixel model.

The success of brightness temperature estimation using the BAT method led to an investigation of the potential isolation of brightness temperature anomalies using these models. A comparison study was undertaken using active fire information taken from the MODIS and VIIRS active fire products, and burned area information from the TERN Auscover MODIS Burned Area product, and information from these three sources was compared to anomalies isolated from the BAT modelling of the affected locations using a number of rudimentary temperature thresholds. Anomalies were detected the BAT modelling of temperature against raw image temperatures in between 75–99% of cases where

a LEO fire detection took place, with variation based upon the threshold set. Synchronous fire activity was detected between LEO fire products and the BAT anomalies in between 46–68% of case where fires were identified by both active fire products. Using BAT to find anomalies also resulted in an increase of anomalies detected above the LEO products, with between 50–75% of burned area pixels without LEO active fire hotspots resulting in anomalous temperature detections. The comparison of anomaly time of detection versus the LEO fire products was not unexpectedly favourable to the BAT anomalies, with improvements in initial isolation of 5–7 h over the two LEO products used.

Whilst the BAT method of temperature modelling was relatively successful in isolation of anomalies in geostationary imagery, investigations led to the identification of methods of estimation that could theoretically be applied to both geostationary and LEO sensor imagery. A new method of background temperature estimation was developed, this time using the similarities of temperatures measured in a search radius around the target to be estimated, in order to derive suitable candidate pixels for estimation from a single image. This method, known as the Spatio-Temporal Selection (STS) method, was applied to images from a number of case study areas across the AHI-8 full disk, and comparisons were made with values from the prediction image and the 5×5 contextual estimation. The STS method demonstrates a 10–40% improvement in variation over contextually derived temperatures, with marked improvements in visually assessed accuracy. The method also provides more estimates of temperature than the contextual estimator, with between 16–45% more pixels able to have their brightness temperature estimated. The study demonstrated the potential extension of the method into use of LEO imagery and highlighted other deficiencies with the contextual estimation method that the first study of the thesis did not identify.

Samenvatting

Satelliet-teledetectie heeft een lange en succesvolle geschiedenis van gebruik voor branddetectiedoeleinden, en met veranderende klimatologische omstandigheden en de toename van door de mens veroorzaakte landschapsveranderingen, wordt nauwkeuriger en vroegere branddetectie steeds belangrijker. Een intrinsieke eigenschap van veel van de branddetectiealgoritmen die tegenwoordig algemeen worden gebruikt, is de afhankelijkheid van kennis van de onverstoorde staat van de kandidaatpixel voor detectie. Omdat dit niet direct kan worden gemeten, vanwege de buitenmaatse bijdrage van vuur aan het totale stralingsbudget, moeten schattingsmethoden voor deze onverstoorde staat worden gebruikt om afwijkend gedrag van beeldpixels accuraat te identificeren. De verreweg de voorkeur genietende methode voor het bereiken van deze schatting in de meeste operationele producten is het gebruik van een schatting op basis van de helderheidstemperaturen van de pixels die direct rondom de kandidaatpixel liggen, in een convolutiefilterstijlproces.

De nauwkeurigheid van op de context gebaseerde helderheidstemperatuurschattingen wordt beperkt door hoe het omringende landschap is voor de kandidaat-pixel. Factoren zoals de helling en het aspect van de landvorm, landbedekking en landgebruik en de aanwezigheid van water in een deel van of in de omringende pixels, dragen allemaal bij aan de resulterende fouten in temperatuurschatting met behulp van de context. Verdere versturende invloeden zoals bewolking, en in geval van brand de aanwezigheid van rook, hebben ook invloed op het potentiële energiebudget dat wordt toegeschreven aan de omringende pixels en dus op de centrale schatting. Ondanks het wijdverbreide gebruik van dit soort schattingen, heeft geen eerdere studie onderzocht hoeveel de fout in de contextuele schatting kan bijdragen aan foutieve resultaten van het gebruik van een dergelijke schatter. De centrale driver van dit proefschriftonderzoek is tweevoudig - ten eerste om de inherente fouten te kwantificeren die betrekking hebben op het gebruik van contextuele schatting van de achtergrondhelderheidstemperatuur; en ten tweede om nieuwe methoden af te leiden voor het verkrijgen van een achtergrondtemperatuur die hun schattingen niet alleen baseert op overeenkomsten in locatie, maar op kortetermijntemporele relaties die elke pixel heeft met potentiële kandidaatschatters.

De reeks geschikte sensoren van de huidige dag voor de taak heeft ook geleid tot een studie van dit type. In de afgelopen vier jaar zijn nieuwe geostationaire platforms gelanceerd met sensoren die onze dekking van de wereld hebben vergroot en, kritisch voor branddetectie, onze tijdelijke resolutie van beeldverspreiding vergroot. Gekoppeld aan de verbeterde ruimtelijke en radiometrische resoluties van deze sensoren, worden we nu steeds meer over-

spoeld met geostationaire beelden die de beeldkwaliteit van sommige lage aardbaanbeelden van slechts 10-15 jaar geleden benaderen, zo niet gelijk zijn. Nooit eerder hebben we de dagelijkse 'heartbeat' van de aarde in zoveel detail gepresenteerd als deze wordt geleverd door deze recorders. Het huidige paradigma voor schatting van de helderheidstemperatuur is gebaseerd op methoden die beelden als discreet en onafhankelijk behandelen, terwijl deze sensoren ons in staat stellen meer te doen dan dat, waarbij we niet alleen informatie uit de ruimtelijke omgevingen van het enkele beeld, maar ook uit de beelden kunnen gebruiken. onmiddellijk voorafgaan aan de tijd.

Hoofdstuk 1 van dit proefschrift schetst het huidige branddetectieparadigma, en de rolschattintemperatuurschatting speelt bij het beschrijven van het bestaan van vuur in het landschap. De hoofdstukken 2-5 van dit proefschrift richten zich op het aanpakken van de onderzoeksdoelstellingen die worden geschetst in paragraaf 1.3, waarvan de belangrijkste bevindingen in deze paragraaf worden besproken. Dit laatste hoofdstuk biedt een samenvatting en synthese van de belangrijkste onderzoeksresultaten, samen met een routekaart voor toekomstige uitbreiding van de methoden die worden beschreven in dit werk voor branddetectie en andere velden.

Vraag 1. Wat is het effect van systematische en structurele fouten veroorzaakt door het gebruik van contextuele schatting in gewone branddetectietechnieken?

Contextueel gegenereerde achtergrondtemperatuurschatting bleek een voldoende nauwkeurige methode te zijn in situaties waar weinig tot geen occlusie van de doelcontext optreedt, maar dat deze relatie degenereerde zodra meer dan een derde van de doelcontext is afgesloten. Uit een volledige schijfanalyse van afbeeldingen van de AHI-8-sensor bleek dat pixels met een contextuele beschikbaarheid van minder dan 65 % een toename van 61 % in schattingsvariabiliteit vertoonden in vergelijking met die pixels die geen aangrenzende bewolking hadden, met deze mate van variatie stijgen met minder beschikbare contextpixels. Gegeven de nauwkeurigheid zoals bepaald in tabel 2.5, werd aangenomen dat ten minste 65 % van de contextpixels beschikbaar zou moeten zijn om de achtergrondtemperatuur van een doelpixel uit de context nauwkeurig in te schatten. Deze waarde wisselt de nauwkeurigheid van de schatting versus de beschikbaarheid van schattingen van de pixeltemperatuur uit, maar zelfs met deze acceptatiegraad ontvangt meer dan één op zeven pixels (14.5 %) geen temperatuurschatting.

Deze studie toonde verder aan dat het gebruik van uitbreidende contextuele vensters voor achtergrondtemperatuurbepaling, zoals vaak wordt aangetroffen in veel brandproducten, fundamenteel onjuist is. De fout van temperatuurschatting stijgt met bijna 50% wanneer de zoekstraal wordt verhoogd

tot slechts 7×7 in situaties waarin een dergelijke behandeling noodzakelijk werd geacht (tabel 2.5), en de herstelfrequentie van temperatuurschattingen bij hogere venstergrootten volgens deze regels was hopeloos. Bijvoorbeeld, bij gebruik van de bovengenoemde contextdrempel van 65 % verhoogde het uitbreiden van het venster het totale aantal schattingen met 3.1 %, waarbij de meeste van deze schattingen standaarddeviaties boven 3K vertoonden. Deze schattingen zijn te lawaaierig om echt te kunnen worden gebruikt, kleinere branden zouden volledig worden geëlimineerd, gezien de toleranties die zouden moeten worden toegepast om zekerheid te hebben bij het opsporen van anomalieën.

Analyse van de contextuele schatting concentreerde zich ook op een aantal casestudiegebieden die werden geselecteerd om te putten uit een breed scala van breed-tegraad-, klimaat- en landvormen om storende invloeden te identificeren. De bronnen van variatie in deze casestudiegebieden waren bijna net zo divers als de gebieden zelf. Bronnen van hoge contextuele schattingsvariaties omvatten brand-, rook-, sneeuw-, seizoen- en ruimtelijke variaties in bodembedekking, rurale en landelijke interfaces, helling en aspect, sneeuwlijnen en land-water-interfaces. Variabiliteit toonde significante clustering rond landschapseigenschappen, met de effecten verhoogd wanneer afgebeeld bij grotere venstergroottes. Gebieden met een hoge heterogeniteit van de landbedekking waren bijzonder gevoelig voor hoge schattingsvariaties, en dit zijn over het algemeen kritieke gebieden voor branddetectie en -monitoring, omdat deze sterk geassocieerd zijn met menselijke activiteit.

Vraag 2. Hoe kunnen we de gemeenschappelijke dagvariatie van opwindende straling gebruiken om de helderheidstemperatuur op een robuuste manier in te schatten?

Het gebruiken van de baan van de volledige schijf om lacunes in een model-passingstechniek op te vullen, werd bewezen als een effectief hulpmiddel voor het genereren van schattingen van de helderheidstemperatuur. De ontwikkelde methode, aangeduid als de Broad Area Training-methode, aggregeerde de mediane temperatuur van $0.25^\circ \times 0.25^\circ$ blokken door lokale zonnetijd in een geïdealiseerd dagmodel van het breedtegewijs zwad. Dit dagmodel, dat gestandaardiseerd was op basis van het gemiddelde en de standaardafwijking van elk blok, werd vervolgens gefilterd om uitschieters te verwijderen. Koppelingen naar individuele pixels gebeurden vervolgens met het standaardmodel als basis voor een ontleding van een enkele waarde gedurende de voorgaande 24 uur vóór de voorspellingstijd.

In vergelijking met een enkele aanpassingsmethode voor pixels die werd gebruikt in [65], ondervond de BAT-methode voldoende gegevens voor pixeltraining in ongeveer 86 % van de gevallen, vergeleken met ongeveer 40 % voor de methode met één pixel. De BAT-methode toonde ook een sterkere

veerkracht in pixels met een verhoogde bewolking in de voorgaande dagperiode, met grote verbeteringen in schattingsfouten wanneer tussen 20 en 50% van de dagperiode een bewolking doorwiste. De methode toonde verhoogde levensvatbaarheid in kustgebieden en gebieden in het zuidoosten van Australië die slecht werden bestreken door de techniek voor het toepassen van één pixel, die gebieden zijn die zijn aangemerkt als kritisch voor branddetectie en -monitoring.

Vraag 3. Hoe effectief is de methode om breed gebied aan te passen bij het identificeren van brandgerelateerde afwijkingen in helderheidstemperatuur in vergelijking met andere branddetectiemethoden?

In vergelijking met veelgebruikte op polaire omloop gebaseerde brandproducten, leverden BBT-afgeleide schattingen een solide basislijn op om anomalieën in geostationaire beelden te identificeren. Voor situaties met branddetecties van zowel de VIIRS- als de MODIS-actieve brandproducten werd in 89.3% van de gevallen abnormale activiteit geregistreerd bij een drempelwaarde voor cutoff-temperatuur van 5 K using BAT, met synchrone detecties in bijna de helft van die gevallen. Rekening houdend met de veel grotere pixelgrootte van de geostationaire sensor en het kortstondige karakter van vuur, was de omissiesnelheid aanzienlijk lager dan verwacht. De commissie van detecties met behulp van de BBT-methode bij de 5K-drempel was ongeveer 23% in vergelijking met een verbrand product dat was aangepast voor het studiegebied, dat ook lager was dan verwacht.

Vergelijkingen voor de tijd van eerste detectie was waar de BAT methode van anomalie isolatie straalde. Verbeteringen in de detectietijd bij het gebruik van de BAT-fittingen bij AHI-temperatuurmetingen waren tussen 2 uur en 6 uur, afhankelijk van de nabijheid van de tijd van aanpassing aan de eerste afwijking gedetecteerd in de dagcyclus voor de 5K-drempel. Dit was het verwachte gedrag gezien de veel hogere temporele dekking van de geostationaire sensor, maar was nog steeds een veelbelovend resultaat in de context van de aanpasmethode en gezien de grove ruimtelijke resolutie. De studie toonde de levensvatbaarheid van het gebruik van een dergelijke methode in een bijna-real-time capaciteit, omdat de nauwkeurigheid van temperatuurfittingen verbeterde toen de schatting dicht bij het tijdstip van de eerste detectie eindigde.

Vraag 4. Hoe kunnen we na verloop van tijd overeenkomsten in beeldlocaties gebruiken om op een later tijdstip een helderheidstemperatuurbeeld te reconstrueren?

Analyse van de fout in gelijktijdige metingen gedurende een discrete tijdsperiode kan leiden tot de identificatie van geschikte trainingspixels voor daaropvolgende temperatuurschatting, die wordt onderzocht met behulp van de spatio-temporele selectiemethode voor schatting. De STS-methode toont

een significante verbetering in nauwkeurigheid van temperatuurschatting, met indrukwekkende weergave van afbeeldingen in vergelijking met de onbewerkte sensorbeelden. De STS-methode toont een verbetering in schattingsvariatie tussen 10-40% in gebieden met voldoende heldere lucht voor pixeltrainingdoeleinden vergeleken met 5×5 contextuele schatting. De STS-methode maakt ook temperatuurschattingen mogelijk op meer locaties, ook onder cloud, met een toename van de beschikbaarheid van schattingen tussen 16-45% in vergelijking met contextuele schattingen.

De resultaten van deze studie wezen ook op tekortkomingen in contextuele schattingen wanneer ze werden geconfronteerd met afwijkend temperatuurgedrag, met STS-schattingen die deze anomalieën gemakkelijker konden identificeren met minder ruis en die grotere gebieden met afwijkend temperatuurgedrag benadrukten in vergelijking met de contextuele schattingen. De STS-werkwijze heeft ook ruimte voor verbetering met betrekking tot kandidaatpixelselectie, met een aantal criteria die kunnen worden gemodificeerd om de resulterende temperatuurschatting verder te verbeteren.