

Detecting groundwater flows from land to sea: a remote sensing perspective

by

Júlio Caineta

BSc. Mining and Geological Engineering, Instituto Superior Técnico, 2009

MSc. Mining and Geological Engineering, Instituto Superior Técnico, 2010

Submitted to the Graduate Faculty of

the Dietrich School of Arts and Sciences in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2022

UNIVERSITY OF PITTSBURGH
DIETRICH SCHOOL OF ARTS AND SCIENCES

This dissertation was presented

by

Júlio Caineta

It was defended on

July 28th 2022

and approved by

Daniel Bain, Associate Professor, Department of Geology and Environmental Science

Brian Thomas, Senior Lecturer, Civil and Geospatial Engineering, Newcastle University

Eitan Shelef, Assistant Professor, Department of Geology and Environmental Science

Jeffrey Newman, Professor, Department of Physics and Astronomy

John Gardner, Assistant Professor, Department of Geology and Environmental Science

Michael Ramsey, Professor, Department of Geology and Environmental Science

William Harbert, Professor, Department of Geology and Environmental Science

Copyright © by Júlio Caineta
2022

Detecting groundwater flows from land to sea: a remote sensing perspective

Júlio Caineta, PhD

University of Pittsburgh, 2022

Submarine groundwater discharge (SGD) represents an important component of the global water cycle. Improved assessment of SGD occurrences provides information to constrain current estimates of total freshwater discharges, which are crucial for evaluation of changes in the global water cycle.

Groundwater usually exhibits distinct and less variable temperature than seawater. Consequently, thermally anomalous plumes often occur in coastal areas with SGD. Nutrient-enriched groundwater flows have been linked to changes in ocean color as a consequence of algae growth and of the presence of other solids. Under such conditions, SGD presence may also manifest as plumes of different colors. This research demonstrates the improved detection of SGD when combining observations of sea surface temperature (SST) and color. To date, only the thermal response has been used from a remote sensing perspective.

Clustering and spectroscopy techniques were applied to detect SGD plumes. SST data are clustered to map thermal anomalies. Derivative analysis and angular distance are applied to identify a color signature linked to SGD. In a novel approach, SST and color are combined to improve the reliability of remote sensing-based SGD detection. The identified plumes are compared to field surveys of SGD occurrences to verify that the novel methods effectively map SGD. These methods are further applied to build a time series of detected SGD, whereby repeated observations through time reveal evidence of persistent SGD. Whereas field measurements provide the strongest evidence of SGD, these results narrow the gap between the well supported field surveys and model-based assessments, by improving the reliability of remote sensing-derived indicators. An alternative approach based on deep learning is proposed to unravel the complex relations between temperature, color, and SGD flows, while also addressing two constraining factors of the initial methods (reliance on one observed spectrum and requirement of significant thermal gradients between groundwater and seawater).

These methods are applied to provide further evidence for SGD occurrences in south

Ireland and in Hawai'i, and showcase that satellite data can be utilized to better identify SGD occurrences.

Keywords: Submarine groundwater discharge, Sea surface temperature, Ocean color, Coastal waters, Remote Sensing.

Table of Contents

1.0	Introduction	1
2.0	Submarine Groundwater Discharge Detection Through Remote Sensing: An Application of Landsat 7 and 8 in Hawai'i and Ireland	6
2.1	Introduction and Background	6
2.2	Data	10
2.2.1	Case Studies	10
2.2.1.1	Southern Ireland	11
2.2.1.2	West coast, Island of Hawai'i	12
2.2.2	Data Sets	12
2.2.2.1	VNIR and TIR Observations from Landsat 7 and 8	12
2.2.2.2	Reference Spectrum of Water With Suspended Sediments and Chlorophyll Content	13
2.3	Methods	15
2.3.1	SST Patterns and Anomalies	15
2.3.2	Derivative Analysis	16
2.3.3	Angular Distance	17
2.4	Results	17
2.4.1	Ireland	18
2.4.1.1	SST Patterns and Anomalies	18
2.4.1.2	Derivative Analysis	18
2.4.1.3	Angular Distance	21
2.4.2	Hawai'i	23
2.4.2.1	SST Patterns and Anomalies	23
2.4.2.2	Derivative Analysis	23
2.4.2.3	Angular Distance	26
2.5	Discussion	28

2.5.1	Using a Multi-Approach to Incorporate Color Data for PSGD Detection	28
2.5.2	Pairing of Sea Surface Temperature and Color to Improve PSGD Detection	29
2.5.3	Refining PSGD Detection	30
2.5.4	Identification of Scene and Local SST Anomalies	33
2.5.5	Validation Using Data from the Case Studies	35
2.5.6	Methods Limitations	38
2.6	Conclusions	38
3.0	Spatio-Temporal Patterns in Remote Sensing Signatures of Submarine Groundwater Discharges in the Southern Irish Coast	39
3.1	Introduction	39
3.2	Materials and Methods	41
3.2.1	Observed and Expected Spectra in Irish Coastal Waters	41
3.2.1.1	Sea Surface Temperature and Color Observations from Land- sat 8	41
3.2.1.2	Expected Thermal Signature	42
3.2.1.3	Expected Color Signature	44
3.2.2	Detection of Potential Submarine Groundwater Discharge	44
3.2.3	Temporal Consistency of PSGD Plumes	46
3.2.4	Plume Buffering	48
3.3	Results and Discussion	48
3.3.1	Temporal Consistency of PSGD Plumes	48
3.3.2	Seasonality of PSGD	51
3.3.3	Spatial and Temporal Patterns of SST Anomalies	56
3.3.3.1	Spring to Summer	58
3.3.3.2	Summer to Autumn	59
3.3.4	Multi-criteria Approach	60
3.3.5	Methods Limitations	65
3.4	Conclusions	67

4.0 Mapping Submarine Groundwater Discharges Through Remote Sensing and Machine Learning	68
4.1 Introduction	68
4.2 Materials and Methods	69
4.2.1 Case Study	69
4.2.2 Sea Surface Temperature and Color	70
4.2.3 Data Preparation	71
4.2.4 Training Samples Generation	72
4.2.5 Convolutional Neural Network	72
4.3 Results and Discussion	74
4.3.1 CNN Classification	74
4.3.2 PSGD in Waterford Harbour	76
4.3.3 Features importance	79
4.3.4 Limitations and Future Research Directions	82
4.4 Conclusions	84
5.0 Conclusions	85
Appendix A. Supporting Information for Chapter 2	87
Appendix B. Supporting Information for Chapter 3	97
B.1 Introduction	97
B.2 Data	98
B.2.1 Remote Sensing	98
B.2.1.1 Table	98
B.2.1.2 Histograms	100
B.2.1.3 Maps	102
B.2.2 Groundwater Monitoring Stations	104
B.2.2.1 Table	104
B.2.2.2 Map	105
B.2.2.3 Graphs	106
B.3 Results	110
B.3.1 Derivative Analysis	110

B.3.1.1	Histograms	110
B.3.1.2	Maps	115
B.3.2	Angular Distance	125
B.3.2.1	Histograms	125
B.3.2.2	Maps	129
B.3.3	Derivative Analysis and Angular Distance	136
B.3.3.1	Histograms	136
B.3.3.2	Maps	140
B.3.4	Sea Surface Temperature	146
B.3.4.1	Maps	146
Bibliography	151

List of Tables

2.1	Landsat scenes used in this study.	13
2.2	Summary statistics of the SST intervals in the Irish case study.	19
2.3	Summary statistics of the SST intervals in the Hawaiian case study.	24
4.1	Confusion matrix.	74
A1	Angular distance maps in the Hawaiian case study.	92
A2	Number of PSGD pixels in the two lowest SST intervals.	93
B1	MODIS scenes used in this study.	98
B2	Groundwater temperature measured from monitoring wells in South Ireland. . .	104

List of Figures

1.1	Schematic depiction of groundwater flows from land to sea.	1
1.2	Temperature and color patterns attributed to submarine groundwater discharge.	2
1.3	Model-based map of fresh submarine groundwater discharge in the contiguous United States coast.	3
2.1	Maps of the study areas.	11
2.2	Reference spectrum of water with suspended sediments and chlorophyll content.	14
2.3	Results of the methods applied to imagery acquired over Ireland’s south coast. .	21
2.4	Distribution across SST intervals of the results from the Irish case study.	22
2.5	Results of the methods applied to imagery acquired over the Island of Hawai’i. .	26
2.6	Distribution across SST intervals of the results from the Hawaiian case study. .	27
2.7	Post processing of potential submarine groundwater discharge plumes.	33
2.8	Distribution of radon activity and salinity from field observations.	37
3.1	Mode of the expected thermal signature of PSGD plumes in each season.	45
3.2	PSGD detected by incorporating temporal consistency.	50
3.3	Seasonal PSGD detected by incorporating temporal consistency.	52
3.4	Seasonal PSGD detected regardless of temporal consistency.	54
3.5	Daily records of groundwater level near Waterford Harbour.	56
3.6	Mode of SST intervals in each season.	57
3.7	Number of PSGD pixels per scene resulting from each method independently. .	62
3.8	PSGD detected by combining derivative analysis, SST, and temporal consistency.	64
3.9	PSGD detected by combining angular distance, SST, and temporal consistency.	65
3.10	PSGD detected by combining SST interval with temporal consistency.	66
4.1	Pseudo-color representation of the surface reflectance in Waterford Harbour. . .	70
4.2	Sea surface temperature in Waterford Harbour.	71
4.3	Architecture of the convolutional neural network model.	73
4.4	Distribution of the probability of PSGD by CNN classification.	77

4.5	Pixels classified as PSGD by the CNN	77
4.6	SHAP feature importance as measured by the mean absolute Shapley values.	80
4.7	SHAP summary plot.	81
4.8	Distribution of the probability of PSGD by CNN classification with the thermal band included.	81
4.9	Pixels classified as PSGD by the CNN with the thermal band included.	82
A1	Distribution of the distance from PSGD plumes to the Irish coastline.	87
A2	Distribution of the distance from PSGD plumes to the Hawaiian coastline.	88
A3	Distribution of PSGD plumes area in the Irish case study.	89
A4	Distribution of PSGD plumes area in the Hawaiian case study.	90
A5	Angular distance maps in the Irish case study.	91
A6	Distribution of radon activity and salinity from field observations.	94
A7	Distribution of the distance from PSGD plumes to field samples.	95
A8	MODIS Level-3 chlorophyll concentration.	96
B1	Percent number of valid observations per pixel in each season.	100
B2	Cumulative percent number of valid observations per pixel in each season.	101
B3	Pixels with at least 5 valid observations.	102
B4	Pixels with at least 5 valid observations in each season.	103
B5	Subcatchments, monitoring boreholes and buoy, and previously identified PSGD in South Ireland.	105
B6	Comparison of the sea surface temperature observed from MODIS and from a monitoring buoy.	106
B7	Monthly aggregated measurements of groundwater temperature from monitoring boreholes in South Ireland.	107
B8	Comparison between groundwater temperature and sea surface temperature.	108
B9	Same as Figure B8, but only for the period 2013-14.	108
B10	Same as Figure B8, but only for the period 2016-18.	109
B11	Relative frequency of the number of times each pixel is flagged as PSGD by derivative analysis.	110

B12	Cumulative relative frequency of the number of times each pixel is flagged as PSGD by derivative analysis.	111
B13	Relative frequency of the number of times each pixel is flagged as PSGD by derivative analysis in each season.	112
B14	Same as Figure B13 but as a percent of the number of valid observations. . . .	113
B15	Relative frequency of the number of times each pixel is flagged as PSGD by derivative analysis and SST intervals in each season.	114
B16	PSGD detected by derivative analysis and temporal consistency.	115
B17	Same as Figure B16 but with more constrained temporal consistency.	116
B18	PSGD detected by derivative analysis, SST intervals, and temporal consistency.	117
B19	PSGD detected by derivative analysis, SST intervals, and temporal consistency, with at least 6 valid observations.	118
B20	Same as Figure B19 but with more constrained temporal consistency.	119
B21	Seasonal PSGD detected by derivative analysis and temporal consistency. . . .	120
B22	Seasonal PSGD detected by derivative analysis, SST intervals, and temporal consistency.	121
B23	Same as Figure B22 but with plume buffering.	122
B24	Same as Figure B23 but with at least 6 valid observations.	123
B25	Seasonal PSGD detected by derivative analysis and SST intervals.	124
B26	Relative frequency of the number of times each pixel is flagged as PSGD by angular distance.	125
B27	Cumulative relative frequency of the number of times each pixel is flagged as PSGD by angular distance.	126
B28	Relative frequency of the number of times each pixel is flagged as PSGD by angular distance in each season.	127
B29	Cumulative relative frequency of the number of times each pixel is flagged as PSGD by angular distance in each season.	128
B30	PSGD detected by angular distance and temporal consistency.	129
B31	Same as Figure B30 but with plume buffering.	130

B32	PSGD detected by angular distance, SST intervals, and temporal consistency, with at least 6 valid observations.	131
B33	Seasonal PSGD detected by angular distance and temporal consistency.	132
B34	Seasonal PSGD detected by angular distance, SST intervals, and temporal consistency.	133
B35	Same as Figure B34 but with at least 6 valid observations.	134
B36	Seasonal PSGD detected by angular distance and SST intervals.	135
B37	Relative frequency of the number of times each pixel is flagged as PSGD by derivative analysis and angular distance.	136
B38	Cumulative relative frequency of the number of times each pixel is flagged as PSGD by derivative analysis and angular distance.	137
B39	Relative frequency of the number of times each pixel is flagged as PSGD by derivative analysis and angular distance in each season.	138
B40	Cumulative relative frequency of the number of times each pixel is flagged as PSGD by derivative analysis and angular distance in each season.	139
B41	PSGD detected by derivative analysis and angular distance.	140
B42	PSGD detected by derivative analysis, angular distance, and temporal consistency, with at least 6 valid observations.	141
B43	PSGD detected by derivative analysis, angular distance at the 10th percentile, and temporal consistency.	142
B44	Seasonal PSGD detected by derivative analysis, angular distance, and temporal consistency, with at least 6 valid observations.	143
B45	PSGD detected by derivative analysis, angular distance, SST intervals, and temporal consistency, but with different order of intersection.	144
B46	Seasonal PSGD detected by derivative analysis, angular distance, SST intervals, and temporal consistency, but with different order of intersection.	145
B47	Mode of SST intervals.	146
B48	PSGD detected by SST intervals and temporal consistency.	147
B49	Seasonal PSGD detected by SST intervals and temporal consistency.	148
B50	Maximum expected thermal signature in each season.	149

B51 Median expected thermal signature in each season. 150

1.0 Introduction

Submarine groundwater discharge (SGD) is an important component of the water cycle that consists of flows of groundwater from coastal aquifers into the oceans (Figure 1.1). Recent studies have documented connections between groundwater and sea level rise, as most of the displaced groundwater eventually goes to the oceans, yet total freshwater discharge into the oceans remains uncertain. SGD often contains high nutrient concentrations and other dissolved solids, which may have ecological and economical impacts to coastal ecosystems. Furthermore, sustainable groundwater management recognizes SGD as a resource that may be used for human consumption, especially in coastal regions facing water shortages. However, locating these resources continues to be a challenge and an important current research topic, besides being an important part of the larger discussion of groundwater discharge. This research contributes to the mapping of SGD.




Figure 1.1: Schematic depiction (not to scale) of groundwater flows from land to sea. Arrows indicate flow direction. Green dashed lines depict possible paths of buoyant submarine groundwater discharge plumes up to the surface. Adapted from Burnett et al.[9]

SGD can be linked to patterns in sea surface temperature and color. Typically, groundwater temperatures exhibit less variability than sea surface temperature and tend to have temperature distinct from sea water. Therefore, coastal waters where SGD occurs may be affected by a thermally anomalous plume. Nutrient-enriched flows have been linked to changes in the color of coastal waters as a consequence of algae growth and the presence of colored dissolved organic matter. Given the assumption that groundwater carries those com-

ponents, SGD plumes also reveal a change in color. These are the two responses—thermal and optical—to SGD occurrence that are explored in this study (illustrated in Figure 1.2).

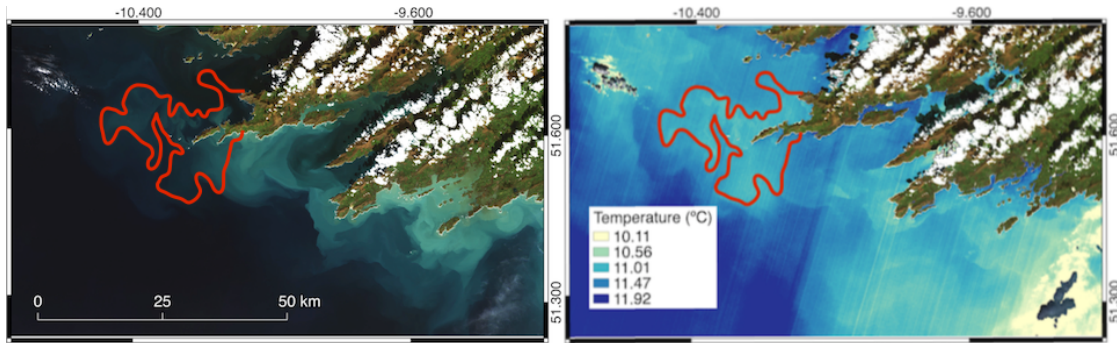


Figure 1.2: Patterns in seawater attributed to submarine groundwater discharge in Southwest Ireland, as captured by Landsat 8 on 10/17/2017. Left: Natural color representation (red, green, and blue bands). Right: Sea surface temperature ($^{\circ}\text{C}$), illustrating buoyant cold water plumes emanating from near shore coastal waters. There is a visible correspondence between greener (left) and colder plumes (right). The red line outlines one cold plume (right). The same outline is drawn in the color image (left) to highlight the similar patterns. The points with lower temperatures (yellow) are likely due to clouds and cloud shade.

Efforts to identify where and when SGD occurs in the globe have been restricted to modeling approaches (see example in Figure 1.3) and to observations over limited spatial domains. Observation-based studies that have a broader spatial coverage use satellite data, but only take advantage of a subset of the available data, and then rely on local secondary, harder to obtain—time consuming and expensive—information to further confirm the existence of SGD. Changes in sea surface temperature and color can be captured and assessed through images derived from satellite data. These data include the thermal infrared (temperature) and optical (color) wavelengths of electromagnetic radiation. Changes in sea surface temperature are proportional to changes in emitted thermal infrared radiation. Due to groundwater’s lower density, SGD floats to the surface, thus creating an observable plume (illustrated by the green dashed lines in Figure 1.1). The variations and patterns in the sea surface color can be assessed through analyses of the reflected radiation over the aforementioned wavelengths.

This dissertation is presented as the compilation of three core chapters, all on the development and application of methods to detect SGD occurrences through remote sensing. The main goal of this research was to demonstrate the improved ability of detecting SGD when combining observations of sea surface temperature and color. The methods developed in this dissertation improve upon previous remote sensing approaches, which have only utilized

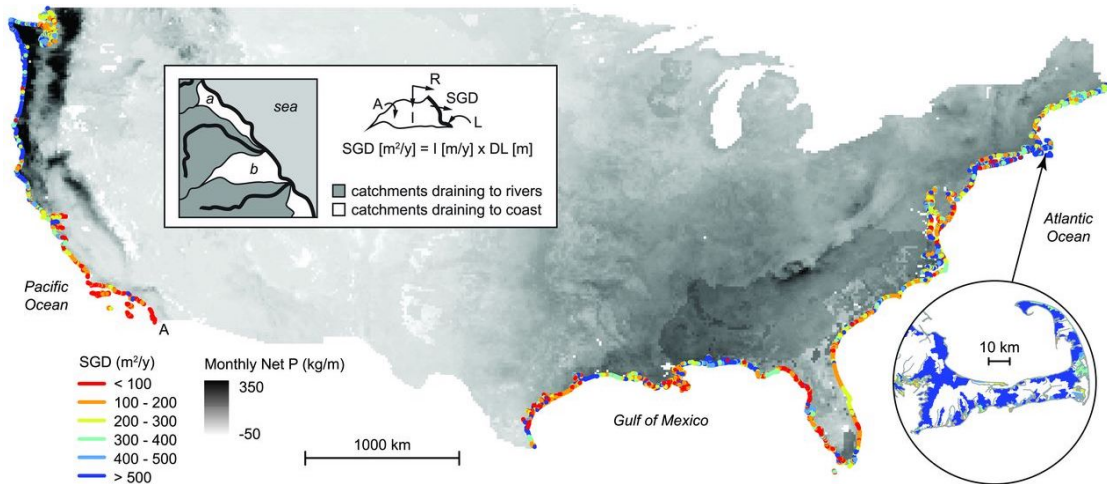


Figure 1.3: Map of fresh submarine groundwater discharge (SGD) rates along the contiguous United States coast, derived from a model-based approach. It illustrates the estimated ubiquity of SGD occurrences and the positive relation with precipitation (Net P), particularly on the West Coast. Adapted from Sawyer et al.[66]

SGD’s thermal response.

In Chapter 2, I show that thermal and color data can be combined to improve the reliability of remote sensing-based SGD detection, thus providing a higher level of scrutiny when determining SGD sites that warrant further investigation (e.g., field surveys). Results from my multi-method approach illustrate that combining thermal and optical signatures enables the identification of SGD sites that could be missed if only considering the thermal signal. Further, I compared the identified SGD plumes with in situ proven SGD occurrences, and verified that SGD detected from my methods contribute effectively to map areas where SGD is likely to occur. Additionally, I presented a simple, physically-based heuristic to reduce the number of false positives and refine SGD detection.

Chapters 3 and 4 rely on the relation between sea surface temperature and ocean color—the thermal and optical SGD signatures, respectively—established in Chapter 2, and build upon it to provide additional methods for SGD detection. In Chapter 3, the same detection methods of Chapter 2 were applied to build a time series of detected SGD in the same region. I explored this time series and applied a method based on temporal consistency, whereby time persistent signatures provide stronger evidence for prime sites for SGD occurrence. I also demonstrate a simple adaptive approach to the seasonally and spatially variable gradient

between groundwater and seawater temperatures. Accordingly, I analyzed seasonal variability of SGD, suggesting connections between groundwater level and the effectiveness of SGD signatures, likely linked to seasonal varying SGD flows. Results identified 15 potential SGD sites along the study area, a region where previous studies based on thermal patterns alone had only identified 4 sites, despite that model-based estimates have suggested a broader occurrence of SGD alongshore. Whereas field measurements provide the strongest evidence of SGD presence, these results demonstrate that my methods narrow the gap between the well supported field surveys and the model-based assessments, by improving the reliability of remote sensing-derived indicators.

In Chapter 4, I leverage the power of machine learning techniques to tackle the complex relation between temperature, color, and SGD flows. In Chapters 2 and 3, the detection criteria based on color relied on patterns extracted from one observed spectrum—in the visible and near infrared range—attributed to SGD influence. In the approach presented in Chapter 4, there is no dependence on a single target spectrum; instead, several target spectra were collected and the framework can be easily extended to amplify the target spectra. Additionally, this approach is less dependent on the thermal signature, as it is only utilized at the spectra collection stage. Then, the observed optical signatures can potentially be applied without dependency on thermal data. Therefore, this separation enables the utilization of remote sensing data from sensors without a thermal band, and eventually with higher spatial resolution in the optical bands. Furthermore, the reduced dependency on the thermal signature also makes this approach suitable for cases where the thermal gradient between groundwater and seawater is too low for a reliable detection (a limitation mentioned in Chapter 3). Ultimately, this approach has the potential to be more robust to variations in local and environmental conditions. Results match with ground-truthed SGD, thus demonstrate the ability of a neural network to capture SGD signatures embedded in remote sensing signals.

All the methods presented are applicable to remote sensing imagery from different regions of the globe and of varying coverage areas, thus contributing to the global mapping and monitoring of SGD.

Chapter 2 has been published in *Remote Sensing of Environment*: Caineta, J., Thomas,

B. F., & Bain, D. J. (2022). Submarine groundwater discharge detection through remote sensing: An application of Landsat 7 and 8 in Hawai'i and Ireland. *Remote Sensing of Environment*, 279, 113109. <https://doi.org/10.1016/j.rse.2022.113109>.

2.0 Submarine Groundwater Discharge Detection Through Remote Sensing: An Application of Landsat 7 and 8 in Hawai'i and Ireland

2.1 Introduction and Background

Submarine groundwater discharge (SGD) into the oceans is an important, poorly constrained component of the global water cycle [36, 50, 64], and locally SGD can be of the same scale as large rivers (e.g., $10^7 \text{ m}^3 \text{ d}^{-1}$) [11, 87]. These flows of groundwater from coastal aquifers represent roughly 6% of the total discharge of freshwater into the oceans [94, 10]. Despite the substantial magnitude of global SGD fluxes, recent studies that evaluate links between groundwater and oceans have not explicitly considered SGD [61, 85]. Notwithstanding, estimates of global runoff are larger when SGD is taken into account [62]. Therefore, constraining the role of SGD in the water cycle clarifies global water dynamics. This study presents a method to locate SGD occurrences over large areas, as a step towards improving the global characterization of SGD.

SGD influences multiple processes in coastal geology, geomorphology, geochemistry, biology, hydrology, oceanography, and ecology [30, 75], and recognition of SGD as an important flux from landscapes to oceans increased in the past two decades [9, 30, 52, 36, 80]. SGD can occur in either continental (e.g., Ireland) [92] or insular shelves (e.g., Hawai'i) [27]. Regardless of where it occurs, SGD can cause both negative and positive impacts [64]. The negative impacts are mostly related to changes in coastal ecology due to its nutrient enrichment, which, in excess, can contribute to eutrophication [17] and harmful algal blooms [38]. SGD is one of the main transport pathways of nutrient and dissolved terrestrial materials to the sea [79, 36], thus SGD can deliver contamination to the sea when coastal aquifers are polluted (e.g., excess nutrients, heavy metals, radionuclides and organic compounds) [9, 43]. These environmental impacts degrade economic systems in coastal zones [38]. SGD-driven nutrient inflows to the sea can also have positive impacts. SGD is responsible for a significant percentage, if not the majority, of nutrient inputs to key coastal ecosystems systems (e.g., salt marshes, estuaries, coral reefs) [33, 52, 64]. Whether positive or not, groundwater

discharges can drive changes in biodiversity in coastal ecosystems [69]. Other positive impacts refer to societal benefits: populations in coastal regions use SGD for a diverse range of activities, such as irrigation and fishery production [53, 80]. Most importantly, SGD may be used for human consumption, thus highlighting a potential target for a valuable extraction [59, 14]. This is of particular relevance in most coastal regions, which are densely populated, straining available water resources [21, 89]. One way to exploit SGD for human consumption at large scale is through desalination [59, 60], since processing brackish water provides an economical and operational advantage relative to saline water [4]. Mapping SGD occurrences can contribute to better physical representation of groundwater systems, by providing valuable data to numerical groundwater models [71, 6, 22]. In summary, SGD has fundamental implications for: a) the uncertainty in the contribution of SGD to the global water budget; b) impacts on coastal ecosystems; c) leakage of potential freshwater sources; and d) improvements in hydrological models.

Case studies from around the globe have identified the occurrence of SGD under different hydrogeologic and climatic conditions, using a variety of methods that can be divided into three categories: 1) in situ measurements; 2) modeling; and 3) remote sensing (RS) approaches. In situ measurement methods include seepage meters, piezometers, geophysical surveys, and natural geochemical tracers [9]. While in situ measurements provide ground-truthed data, those techniques are expensive and site specific, and thus are not adequate to capture SGD magnitude and variability over large spatial and temporal scales [9]. Modeling techniques include water balance [9], hydrograph separation [95], theoretical models [43], and numerical simulation [90, 83]. To date, modeling techniques are the methods that have been used to study SGD on large spatial scales (continental to global scales). [95] developed an estimate of SGD on a global scale based on hydrograph separation technique, although this technique only applies to coastal areas with developed stream networks and to zones of relatively shallow and mainly fresh aquifers [9]. [90] studied regional SGD using numerical simulations, suggesting that local studies only capture a portion of SGD processes, since flow systems contributing to SGD may span at least 20 km of the surrounding coast and continental slope. Using an inverse model combined with a global compilation of point observations of radium isotopes (^{228}Ra), [36] estimated SGD integrated over the global ocean (60°S to

70°N). Instead of referring to point discharges, this study sought to quantify the total fluxes, estimating that SGD is at least 3 to 4 times greater than the river water fluxes to the Atlantic and Indo-Pacific oceans. In a novel approach, [66] identified locations along the continental United States (CONUS) coast where SGD occurs, while also providing flux estimates. These authors used data assimilation systems and land cover information, however, land-derived flows were incorporated into SGD, and the temporal resolution followed the typical climate normals period (30 years). This hydrological modeling approach allows mapping of SGD occurrences over large areas (e.g., CONUS) and at high spatial resolution (1/8th degree) for a fixed period of time. Further discussions on the total freshwater discharge on a global scale focused on major spatial and temporal patterns [76], however, discharge was not partitioned into different sources of freshwater. Other global analyses of the coastal interface have focused on deriving general relations from analytical solutions, applicable to coastal areas around the globe [20]. While several SGD studies in the literature cover large spatial scales, they tackle SGD assessments from a modeling perspective. The observation-based characterization of SGD occurrences is limited to local scales. Therefore, a gap remains in mapping SGD occurrences across large areas (continental to global). [66] approached this gap using high resolution hydrologic data, however this approach did not incorporate direct or indirect observation of SGD.

RS studies for SGD have utilized various airborne [27, 77] and spaceborne sensors, including Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) [84, 83], Advanced very-high-resolution radiometer (AVHRR) [84], and different iterations of the Landsat program—5 [3], 7 [67, 92, 46], and 8 [49, 28]. They all used sea surface temperature (SST) gradients to reveal contrasting SGD plumes, and then proceeded to complement and validate their findings by acquiring in situ measurements. In general, these methods are applicable wherever there are temperature and density differences between coastal sea waters and SGD plumes, either at high or low latitudes, where groundwater is warmer or colder than sea water, respectively [27]. Both warm and cold SGD plumes float to the surface, since salinity is the main driver in determining density and thus buoyancy [58]. We underline two RS studies which have proposed SGD plumes detection criteria based on SST. One used a standardized anomaly index to highlight thermally anomalous plumes linked to

SGD [92]. This index is a linear transformation of SST, which enables the comparison of SST derived from images acquired at different dates. The other study utilized consistency in groundwater temperature across seasons [18] to identify plumes with lower temperature variability over time [45]. Both of these methods require the visual inspection of images to assess the existence and location of potential SGD occurrences, a limitation addressed in this study, by automatically isolating locations likely linked to SGD. RS images have the potential to provide estimates of SGD volumes, e.g., [31] document a positive relation between total SGD flux and the area of the mapped plume. This possibility offered by RS approaches is valuable, considering that unlike the easily gauged flows from major rivers, there is no simple means to gauge and analyze SGD fluxes to the sea [9].

Spaceborne-based RS combines both temporal and spatial coverage to aid holistic analysis and understanding of SGD. In this study, we explore the effect of SGD on water-leaving radiance (emitted and reflected) as observed from space. Changes in water-leaving radiance enable a broad spatial understanding and evaluation of potential SGD (PSGD) occurrences. The term PSGD is used here to differentiate between an RS inferred SGD occurrence and SGDs verified through ground-truthing. Specifically, in this study, we focus on the less commonly occurring freshwater component of SGD [9] and ignore the mixing and seawater components (recirculating flows). In many cases, SGD from recharged aquifers on land, driven by hydraulic gradients, only represents a minor portion of the total SGD flux [80]. We introduce a methodology which uses radiance data from the visible and near infrared (VNIR) and the thermal infrared (TIR) ranges of the electromagnetic spectrum. Groundwater and sea water usually have different temperatures, therefore, a change in SST can be measured where SGD occurs and reaches the water surface. This change in SST may be inferred through changes in the surface radiance emitted on the TIR range: warmer surfaces have higher radiance than colder ones. As mentioned above, SGD carries high nutrient concentrations that contribute to an increase in primary productivity [72], which then increase chlorophyll content [81, 72]. In turn, these effects change the color of coastal waters—generally referred to as ocean color (OC)—which then translate into changes in the VNIR surface reflectance [88]. RS-based observations of SST and OC can be confounded by several factors. For example, the SST signal can be obfuscated by low tides (warmer temperatures and mixed water pixels), up-

welling events (cold plumes), surface water, wind and ocean currents and eddies (possibly cold and increase mixing between SGD and seawater), and anthropogenic inputs (warming discharges). [28] consider the RS limitations for temperature-based detection of SGD, and categorize them in five groups: technical limitations, geological and hydrogeological characteristics, environmental and marine conditions, coastal geomorphology, and anthropogenic sources. Retrievals of OC can be influenced by almost the same factors, in addition to sun glint (which masks the signal). Specifically, nutrient-enrichment links to SGD can be hindered by other sources of nutrients such as upwelling events, surface and anthropogenic discharges.

We introduce a new approach which improves previous RS-based methods by combining both thermal and optical observations of the sea surface, thus combining SST and OC to systematically identify SGD occurrences, using historical RS archives. By leveraging two layers of observable evidence (SST and OC) across broad spatial and temporal coverages, this study fills a gap between the well supported field surveys and the model-based regional assessments of SGD. We attempt to address that gap through observation, taking advantage of medium resolution RS data available around the globe, thus contributing to advancing our knowledge of SGD occurrences over large areas (10^4 – 10^5 km²). The novel methods presented here are applicable to RS data archives regardless of their spatial extent. We demonstrate our methods with case studies from Ireland and Hawai'i, where SGD has been ground-truthed, and therefore also provide validation points. The specific objectives of this study are: a) to explore the links between SGD and effects measurable from RS data (changes in SST and OC); and b) to map PSGD occurrences on selected sites.

2.2 Data

2.2.1 Case Studies

Our methods were applied to two case studies from contrasting climates and with known and distinct spatial and temporal SGD scales ground-truthed through in situ sampling:

Southern Ireland [92] and the West coast of the Island of Hawai'i [27]. These case studies also contrast in their bathymetric profile—continental shelf and steeper, insular shelf, respectively—and both locations have fractured bedrock—typically linked to point-sourced SGD—and sedimentary bedrock—typically linked to diffuse SGD. These locations were selected from previous studies documenting SGD, which provide ground-truthed data used to validate and calibrate our methods.

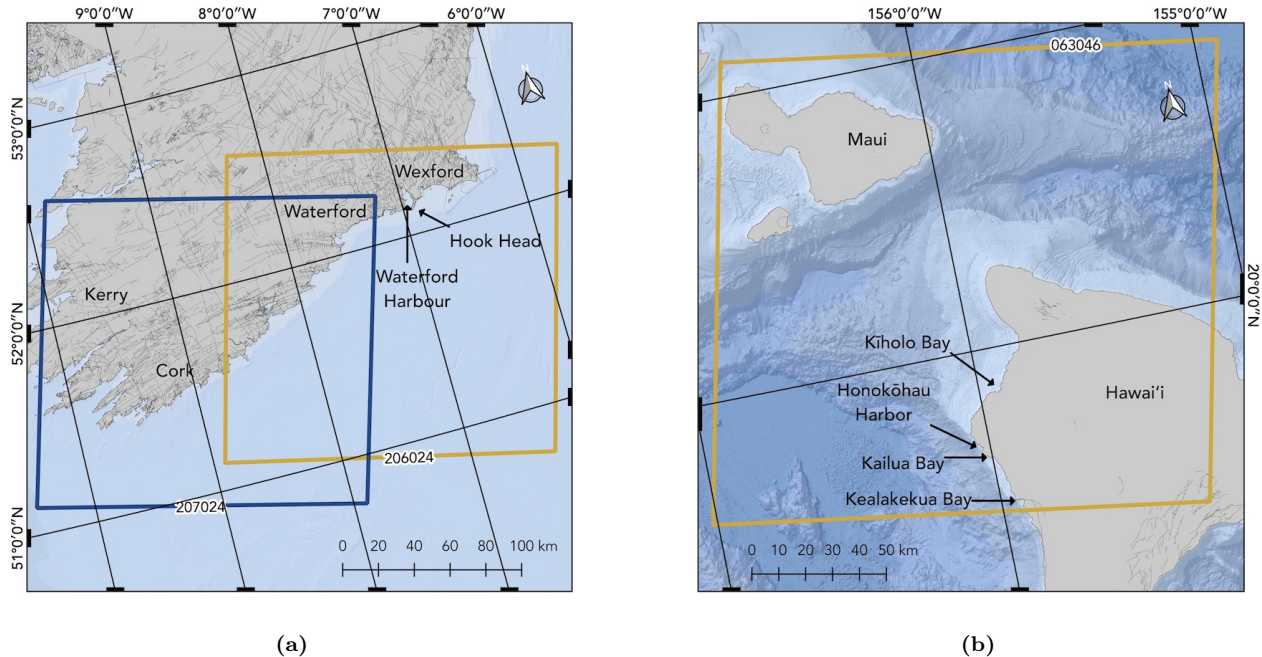


Figure 2.1: Location maps of the study areas in (a) Ireland and (b) the Island of Hawai'i. Arrows indicate sites where field surveys related to SGD were conducted. In Ireland, [92] measured ^{222}Rn isotope activity. In Hawai'i, in situ measurements of salinity were acquired by [27]. The blue and yellow boxes show the spatial extent covered by the Landsat scenes used in this study, according to the Landsat World Reference System: path 206 row 024 (blue, Landsat 8) and path 207 row 024 (yellow, Landsat 7) in Ireland; path 063 row 046 (both Landsat 7 and 8), in Hawai'i. Dashed lines show bedrock faults [23, 68].

2.2.1.1 Southern Ireland

The western and southern coastlines of the Republic of Ireland are encompassed by many islands, peninsulas headlands and bays [92]. Ireland has a temperate maritime climate, which is moderated by the warm influence of the Gulf stream, producing generally warm summers and mild winters [32].

Wilson and Rocha [92] studied PSGD occurrences along most of the Irish coastline, applying data from the Landsat 7/ETM+ thermal infrared band to detect temperature

anomalies in coastal waters. In one of the sites identified as a potential SGD location—Hook Head, Wexford Co., in the southeast coast—in situ concurrent measurements of ^{222}Rn isotope activity and salinity confirmed that detected temperature anomalies were in fact caused by SGD.

2.2.1.2 West coast, Island of Hawai'i

The Island of Hawai'i has the widest extremes of climate in the entire Hawaiian archipelago, ranging from humid tropics to arid, semiarid, temperate, and ice climates, mostly depending on orography [29]. In the western coastline, the climate is mostly dry [40]. The average groundwater temperature in coastal aquifers in the west coast is approximately 20 °C, which contrasts with the typical coastal ocean temperatures of 24–28 °C [27]. [27] surveyed the region using an airborne TIR sensor to identify 31 SGD locations along approximately 100 km of the west coastline. The aerial survey was combined with three in situ measurements: 1) vertical water column profiles of temperature and salinity; 2) collection of water samples to determine regional nutrient characteristics; and 3) continuous measures of ^{222}Rn and salinity. These analyses established temperature-salinity and nutrients-salinity relationships, therefore showing that SST can be used to estimate nutrient concentrations at the surface. They also determined that SGD flows into the ocean must be the only source of new nutrients to the nutrient-limited waters of the region.

2.2.2 Data Sets

2.2.2.1 VNIR and TIR Observations from Landsat 7 and 8

Two Landsat images were used for each case study location (see Table 2.1): one from Landsat 7 (L7), chosen for temporal proximity—not an overlap—with in situ measurements; and the other from Landsat 8 (L8), to seek more up-to-date SGD occurrences, and to showcase the application of the methods with more current data sets. VNIR data were obtained as surface reflectance from the level-2 products, and TIR data were obtained from the level-1 products and calibrated to brightness temperature. Both VNIR and TIR data are provided

at 30-m pixel resolution—TIR data are downscaled from 60 (L7) and 100 m (L8) resolution using cubic convolution [63, 37]—, which is preserved throughout our entire processing and analysis steps. All Landsat scenes used were from the Tier 1 of the Collection 1 archive and were sourced from USGS.

2.2.2.2 Reference Spectrum of Water With Suspended Sediments and Chlorophyll Content

Goodin et al. [24] conducted an experiment to assess changes in the water surface reflectance due to the presence of sediments and chlorophyll content. The experiment was conducted under a controlled environment, whereby spectra were captured from clear water, suspended sediments and algal chlorophyll, creating a reference reflection signature characteristic of chlorophyll production in coastal waters (referred to here as the Goodin spectrum). The authors showed that the first derivative of the reflectance spectrum of clear water is close to zero at all wavelengths (from visible to near infrared). They also verified that the effect of suspended sediments in the spectrum is almost eliminated with the second derivative, and that the presence of algal chlorophyll is still visible in the second derivative. Figure 2.2 shows the observed spectrum of water with suspended sediments and chlorophyll content (0th order, blue line), and the corresponding first (1st order, yellow line) and second derivatives (2nd order, pink line). In this figure, two segments of the second derivative are estimated (dashed line) from the 0th order data, using a Savitzky-Golay filter [65]. The referenced

Table 2.1: Details of the four Landsat scenes used in this study. All scenes were sourced from USGS.

Case study	Platform	Date	Landsat Product ID (level-1)
Ireland	L7	2010-06-02	LE07_L1TP_206024_20100602_20161214_01_T1
	L8	2016-06-01	LC08_L1TP_207024_20160601_20170324_01_T1
Hawai'i	L7	2007-04-29	LE07_L1TP_063046_20070429_20160923_01_T1
	L8	2019-04-06	LC08_L1TP_063046_20190406_20190422_01_T1

study only included a shorter segment of the second derivative (solid line).

For the current study, we explore this marked signature in the second derivative to identify where coastal waters may have been enriched with algal chlorophyll, and therefore are potentially under the effect of SGD. Here, we assume that the change in the reflectance due to chlorophyll is similar across different conditions (atmospheric conditions, incident light, different sources of chlorophyll, different water properties such as salinity and temperature).

Prior to further calculations, we resampled the Goodin spectrum to the same spectral resolutions of L7 (4 bands: blue, green, red, near infrared) and L8 (5 bands: coastal blue, blue, green, red, near infrared). This resampling process was performed using the predefined spectral response functions of both satellite sensors with the R package HSDAR [42]. Accordingly, the angular distances are calculated on a 4-dimensional space and on a 5-dimensional space, for L7 and L8, respectively.

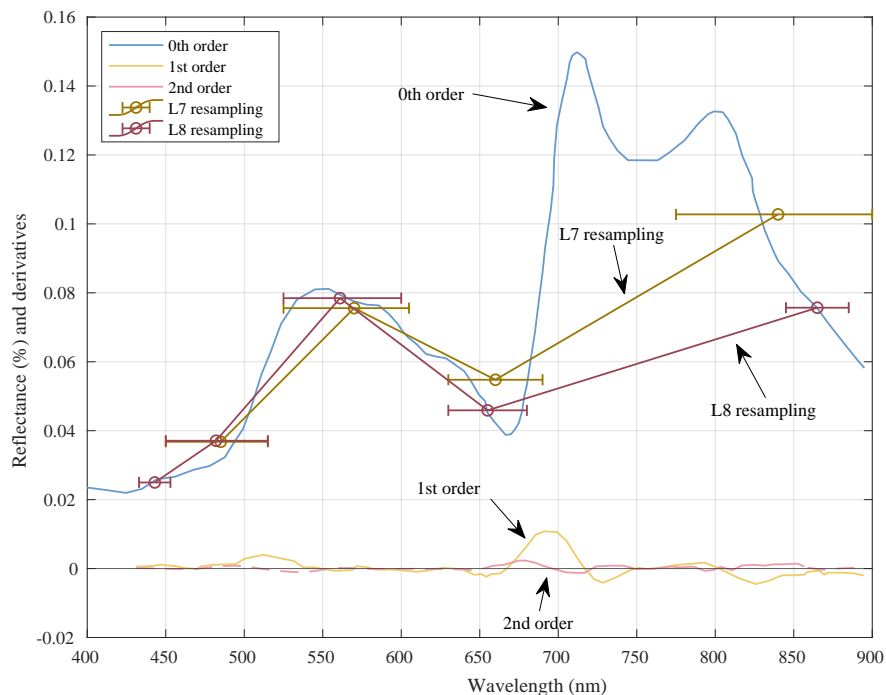


Figure 2.2: Reference spectrum of water with suspended sediments and chlorophyll content (0th order, blue line), and the corresponding first (1st order, yellow line) and second derivatives (2nd order, pink line; adapted from [24]). Parts of the second derivative were estimated (dashed line) using a Savitzky-Golay filter [65]. The reference spectrum is resampled to the Landsat 7 (dark yellow line) and Landsat 8 (purple line) resolutions. Circles point to the average wavelength in each band and the horizontal error bars show the corresponding wavelength ranges.

2.3 Methods

Our methods leverage SST and OC to determine PSGD locations by analyzing the observed VNIR and TIR spectra of coastal waters. TIR is used to infer SST anomalies. The VNIR spectra—here representing OC—are processed through two methods: 1) derivative analysis; and 2) angular distance. Parallel processing of both spectral analysis methods is akin to using co-related variables without assuming any dependence between them, following an analysis where those results are intersected.

2.3.1 SST Patterns and Anomalies

SST clustering isolates the low and high SST intervals and identifies local SST anomalies. Water pixels in the scene are grouped according to brightness temperature derived from TIR. Typically, RS approaches require inference of SST from brightness temperature, a procedure dependent on the RS platform [5]. However, this procedure is only required if absolute values of SST are needed; otherwise, brightness temperature can be used directly as a proxy to relative SST, which is the case here. SST intervals are created using a k -means variant, *Ckmeans.1d.dp* [86], a version optimized for univariate applications, and therefore equivalent to a Jenks natural breaks approach. We set the number of clusters, k , to 5, following the simple approach of having a group of low temperatures, a group of high temperatures, a group in the medium range of temperatures, and two other intermediary groups in between those three. The resulting groups are constituted by non-overlapping ranges of temperatures. In practice, SST clustering is used to create a predefined number (k) of SST intervals across a scene.

Pixels where SST is within the two lowest intervals are marked as SST anomalies and combined with color-related metrics—VNIR second derivatives (Section 2.3.2) and angular distance (Section 2.3.3)—to identify PSGD. These SST intervals are then mapped for visual analysis of SST spatial patterns. Additionally, we use SST intervals to analyze the distribution of the color-related metrics across the range of SST. In previous studies, the mean SST in an RS image had been used as a simple threshold to find pixels with temperature

anomalies [92], however, daily SST variations have a non-Gaussian distribution [73]—unless spatially averaged over multiple degrees [74]—and, therefore, the mean represents a suboptimal metric for the central tendency. Here, cluster analysis is used as a non-parametric approach to identify SST anomalies.

2.3.2 Derivative Analysis

Derivative analysis is used to highlight pixels where surface reflectance indicate SGD influence, by capturing a marked signature on VNIR spectra observed in the Goodin spectrum. We compute second derivatives in a pixel-wise manner, by applying a Savitzky-Golay filter [65] (SG) to the VNIR data. SG is a signal processing filter that smooths the spectrum by fitting a polynomial. A 2nd order polynomial is used for the spectral resolution in the L7 and L8 data sets. SG then computes the derivatives using a finite differences approach.

Considering the Goodin spectrum, and the limited spectral resolution available from both L7 and L8, it becomes less straightforward how to capture the differences in the spectrum caused by the chlorophyll content. As seen in Figure 2.2, the Goodin spectrum has a downward concavity (negative second derivative) around the green band range (500–600 nm), whereas around the red band range (625–700 nm) it has an upward concavity (positive second derivative). Accordingly, we expect that an inflection point exists between the two wavelength ranges, such that the second derivative changes sign from negative to positive. The inflection point can be found by comparing the sign of the second derivative on the green band (negative second derivative expected) and the red band (positive second derivative expected).

Output of the derivative analysis approach is defined by the intersection of the set of pixels selected based on color and the set of pixels selected based on temperature: pixels with SST within the two lowest intervals (see Section 2.3.1) are intersected with pixels where inflection occurs.

2.3.3 Angular Distance

The angular distance (AD) is used to find pixels whose spectrum is most similar to the Goodin spectrum. The angular (cosine) distance measures the angle between two vectors (two spectra) with any number of dimensions (bands). In the context of remote sensing, particularly hyperspectral remote sensing, this technique typically takes the name of spectral angle mapper [35] and is used to discriminate among endmembers. The angular distance is a measurement of orientation and not magnitude, thus it primarily depends on spectral shape and is tolerant to variations in illumination and albedo effects [35]. In more general terms, it represents correlation between pairs of spectra, including from multispectral sensors such as Landsat 8 OLI [93]. In our application, we compute the angular distance between the Goodin spectrum—resampled to the L7 and L8 resolutions (Figure 2.2)—and the spectrum of each pixel, where increased similarity is gauged by a smaller angle. PSGD detection is triggered for the pixels with the lowest AD values among all. Following a conservative approach, we specify a percentile threshold of 1% to define the set of PSGD pixels.

The output of the method is defined by the intersection of the set of pixels where AD values are below the 1st percentile threshold and the set of pixels where SST is within the two lowest intervals (see Section 2.3.1). The intersected pixel set, from both color-based AD and anomalous SST, is flagged as PSGD.

2.4 Results

The methods described above were applied to four RS scenes: one from L7 and one from L8 for each case study. The results are organized by case study and are described below.

2.4.1 Ireland

2.4.1.1 SST Patterns and Anomalies

The SST intervals resulting from the k -means clustering were mapped to visualize their spatial patterns. These maps can be interpreted as an ordinal representation of SST (using brightness temperature as a proxy), where the two lowest intervals are used to identify PSGD pixels.

The spatial representation of the SST intervals in south Ireland (Figure 2.3a) showed several coastal areas where plumes of colder waters emanated from near the shore. SST intervals based on the L7 scene over the southeast of Ireland (Figure 2.3a–1) revealed two coastal areas with larger plumes of locally colder waters (SST intervals with mean 11.30°C and 12.05°C, see Table 2.2a) near the coast in Cork and in Waterford/Wexford. Other smaller and locally colder plumes were revealed along the coastline. The pixels belonging to the coldest SST interval (mean 10.20°C) were mostly located within an offshore plume on the northeast of the scene, attributed to regional-scale current patterns. In the southwest coast, as seen in the L8 image (Figure 2.3a–2), the coldest waters (SST interval with mean 13.66°C, see Table 2.2b) were located in three areas along the coastline, two of them adjacent. Note that the lowest SST interval from L8 includes some pixels likely contaminated with cloud shadow, hence the significant difference between the minimum and mean values. Similarly, the highest values, both from L7 and L8, are likely due to land contamination in pixels not fully captured on the Quality Assessment band.

2.4.1.2 Derivative Analysis

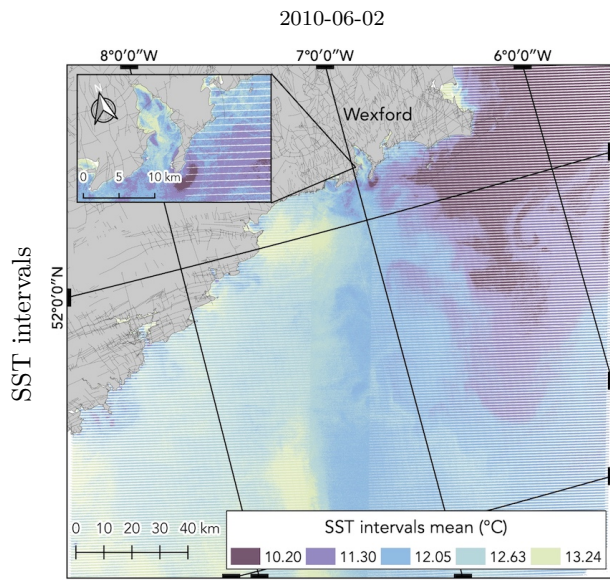
Derivative analysis of the VNIR spectrum—using Landsat 7 bands 1 to 4 and Landsat 8 bands 1 to 5 of the images captured in the southern Irish coast on 2010-06-02 and 2016-06-01, respectively—revealed that the areas where the sign of the 2nd derivative changed (yellow points in Figure 2.3b) enveloped almost the entire coastline captured in both images. The areas where SST was within the two coldest SST intervals (blue points in Figure 2.3b) were mostly located farther offshore, with other smaller areas located closer to the coast line.

There were fewer locations where the criteria outlined above—sign of the 2nd derivative changed and SST was below the mean—co-occurred (red points in Figure 2.3b), and were mostly located closer to the coastline.

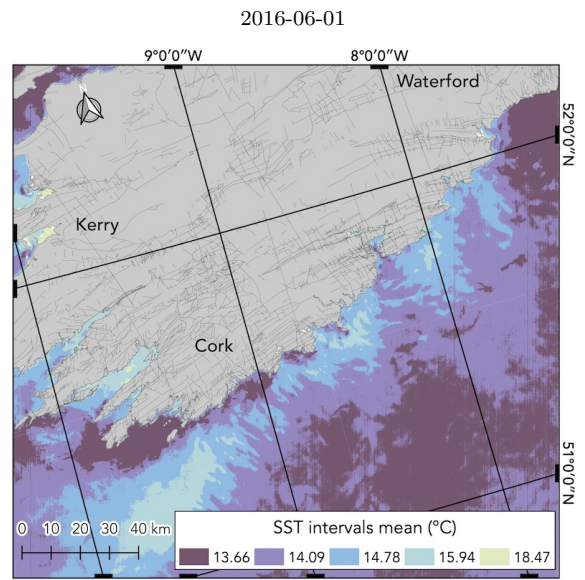
The distribution of the points where the sign of the derivative changes shows different behavior between the two scenes. In the L7 scene, the two coldest SST intervals are also the ones where there are more points where the sign of the derivative changes, which is observed in more than 50% of the pixels within each interval, comparatively to less than 30% in the warmer intervals (Figure 2.4a-1). However, in the L8 scene, the opposite behavior is observed: the two coldest SST intervals are the ones with a lower fraction of pixels where the sign of the derivative changes (less than 20%, see Figure 2.4a-2), therefore highlighting the importance of considering color and temperature simultaneously. Considering both scenes over Ireland, we also note that it is mostly the sign of the 2nd derivative in the green band (negative) that acts as the discriminating factor of the derivative analysis approach, since the sign in the red band is almost always negative across all SST intervals in both scenes, except in the warmest interval in the L8 scene (Figure 2.4a).

Table 2.2: Summary statistics (size, mean, minimum, maximum and standard deviation) of the SST intervals derived from k -means clustering applied to (a) Landsat 7 and (b) Landsat 8 imagery, captured in coastal south Ireland on 2010-06-02 and 2016-06-01, respectively. Brightness temperature is used as a proxy to SST.

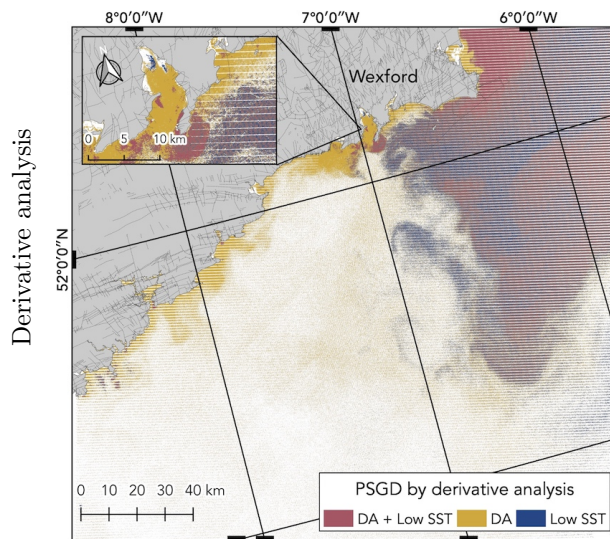
(a) 2010-06-02						(b) 2016-06-01					
k	Size (%)	Mean	Min.	Max.	σ	k	Size (%)	Mean	Min.	Max.	σ
1	11.1	10.20	5.44	10.61	0.37	1	31.3	13.66	11.66	13.87	0.25
2	13.5	11.30	10.92	11.55	0.26	2	45.0	14.09	13.88	14.43	0.15
3	32.0	12.05	11.86	12.17	0.15	3	17.2	14.78	14.43	15.36	0.26
4	32.7	12.63	12.48	12.79	0.15	4	6.0	15.94	15.36	17.20	0.45
5	10.8	13.24	13.10	23.69	0.32	5	0.5	18.47	17.20	26.62	1.23




(a-1)



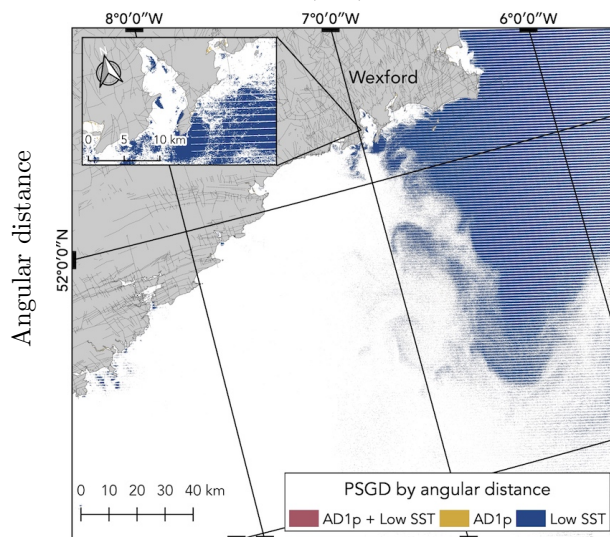
(a-2)




(b-1)



(b-2)



(c-1)



(c-2)

(Caption on next page.)

Figure 2.3: Results of the methods applied to imagery acquired over Ireland’s south coast on (1) 2010-06-02 from Landsat 7 and on (2) 2016-06-01 from Landsat 8: (a) sea surface temperature (SST; actual brightness temperature) intervals from k -means clustering; (b) potential submarine groundwater discharge (PSGD) pixels from derivative analysis of the visible spectra, showing locations where the sign of the 2nd derivative changed (DA; yellow), where SST was within the two coldest SST intervals (Low SST; blue), and where the two criteria co-occurred (DA + Low SST; red); and (c) PSGD pixels from angular distance (AD) applied to reveal the 1% points (AD1p; yellow) with the most similar spectrum to that of water with algal chlorophyll, where SST was within the two coldest SST intervals (Low SST; blue), and where the two criteria co-occurred (AD1p + Low SST; red). Inset shows details around Waterford Harbour. White pixels represent no data points, masked out clouds, or points unrelated to PSGD. Dashed lines show bedrock faults [23].

2.4.1.3 Angular Distance

The AD approach was used to find the pixels with higher spectral similarity to the reference spectrum of water with algal chlorophyll (yellow pixels in Figure 2.3c). When applied to the L7 image over the southeastern coast of Ireland (Figure 2.3c-1), the AD method revealed small plumes—mostly single-pixel plumes (900 m^2 ; see the distribution of the plumes area in Figure A3b-1—which were scattered along the coastline. In this scene, the 1% threshold for the AD is 0.38 rad. Several of the extracted pixels were located farther offshore in the east half of the scene, however, they were mostly isolated pixels, i.e., not adjacent to one another. In the AD map derived from the L8 image over the southwest coast (Figure 2.3c-2), the pixels with the smallest 1% AD values (less than 0.46 rad) were also dispersed along the coastline, including larger plumes (including up to a few tens of pixels) located in estuaries and bays. The majority of the extracted pixels were located in the northeast of the scene, spanning a densely populated area of approximately 215 km^2 adjacent to the coastline, while only a few (less than 20) pixels were extracted from offshore. Conversely, in the corresponding area of the L7 results to that dense region, PSGD pixels were extracted from along the coastline.

In both scenes, the AD values were differently distributed across SST intervals. From the L7 scene, the AD values within the two coldest intervals were markedly different from the two warmest intervals, with the distribution within the middle interval being closer to the coldest intervals (Figure 2.4b-1). Specifically, the distributions within the two coldest intervals and the middle interval have similar modes, whereas within the two warmest intervals the modes are significantly higher—and therefore more dissimilar to the target spectrum—which is also

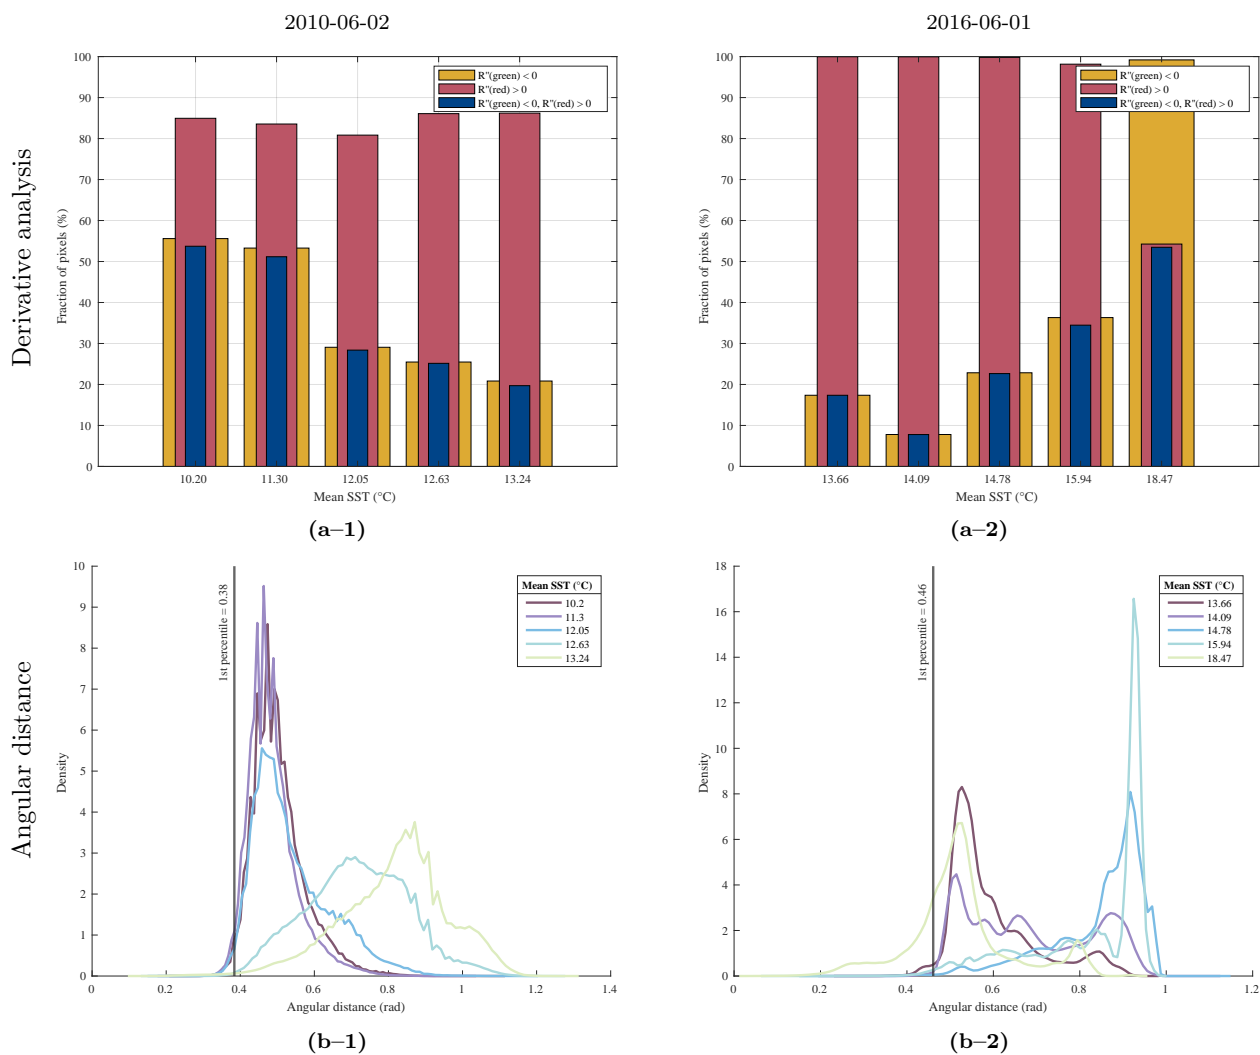


Figure 2.4: Distribution across SST intervals of the results of the methods applied to imagery acquired over Ireland's south coast on (1) 2010-06-02 from Landsat 7 and on (2) 2016-06-01 from Landsat 8: (a) derivative analysis and (b) angular distance (AD). Distribution of the derivative analysis results (a) shows the fraction of pixels where the 2nd derivative of the visible spectrum reflectance (R'') is negative over the green band (yellow bars), those where it is positive over the red band (red bars), and those where both of these conditions are verified (blue bars). Distribution of the AD results (b) shows the probability density estimate for the AD values within each SST interval, including the value of the first percentile used as threshold (vertical black line). SST intervals are represented by their mean values.

reflected in the lower number of pixels below the 1% threshold (0.38 rad). AD values from the L8 scene differed: while the two coldest intervals have similar modes, they are similar to the mode within the warmest interval, yet considerably lower than the modes within the other warmer intervals (Figure 2.4b-2). In fact, in this scene, the warmest interval is the one with the most pixels below the 1% threshold (0.46 rad).

2.4.2 Hawai'i

2.4.2.1 SST Patterns and Anomalies

In the Hawai'i case study, for both L7 and L8 scenes (Figure 2.5a), the large majority of the pixels in the two lowest SST intervals (mean 19.59°C and 21.07°C, respectively) were located in the northern region of the scene, both near the north coast and offshore of the Big Island, as well as around Maui—where there are also several known SGD occurrences [54], however, we only discuss the case of the west coast of the Island of Hawai'i here. The case studies used as ground-truthed SGD are all located in the west coast of Hawai'i, which is dominated by comparatively warmer waters, as it is less affected by trade-winds [40]. In both scenes, the areas surrounding the entire west coastline were dominated by the highest SST interval (mean 21.20°C and 22.65°C for L7 and L8, respectively). The warmest pixels were located over shallower waters and low tides, likely due to misclassified wet, hot sand pixels as water pixels. The L7 scene was captured about 2 h after the low tide and the L8 scene was captured during the low tide peak, according to observed water levels (NOAA Tides & Currents, Kawaihae, HI: site 1617433). However, a few areas nearby the coastline had locally colder pixels, which are within the 3rd and 4th SST intervals, with no neighboring pixels on lower SST intervals. Note that the 3rd and 4th SST intervals from the L7 scene (Table 2.3a) are entirely comprised by one unique value, likely due to L7's 8-bit quantization.

2.4.2.2 Derivative Analysis

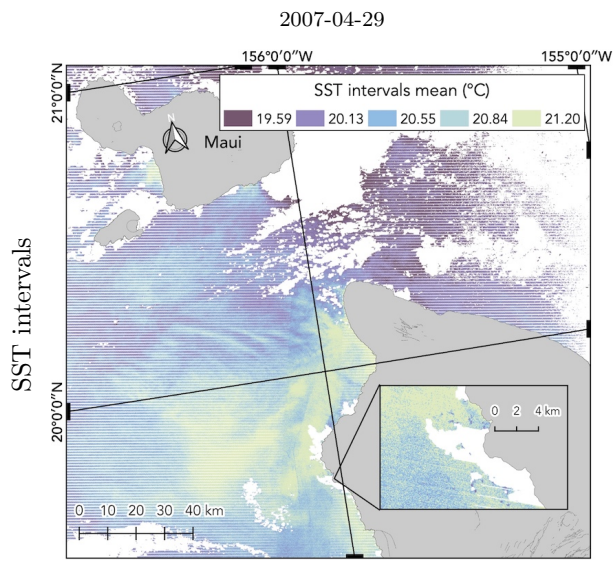
In the Hawai'i case study, the derivative analysis produced similar results for both scenes (Figure 2.5b): the sign of the 2nd derivative changed (yellow pixels) along most of the coastline, notably covering a large area near the west coastline, including over bays with shallow waters; the pixels with SST within the two coldest intervals (blue pixels) were predominantly located across the northern region; and, therefore, the pixels where both the derivative sign changed and the SST within the two coldest intervals (red pixels) were also mostly located along the north coastline. In contrast, the derivative analysis of the VNIR spectra did not work differently for the west and north coasts; although, because the

extracted PSGD pixels (red pixels in Figure 2.5b) are also a function of SST, they are almost exclusively located along the north coastline. The results from the two scenes differ mostly in terms of what happened farther from the coastline: from the L7 scene, several isolated pixels were extracted throughout the ocean on the west of the island, whereas from the L8 image a considerably lower number of isolated pixels were extracted, and those were located mostly over a smaller region off the north coast. Additionally, from the L7 scene, there is a region near the northwest coast with a denser cloud of extracted pixels (yellow pixels). In this region, there is a significant number of pixels with locally lower SST (blue pixels), even if only a small number of them are juxtaposed (red pixels) with locations where the derivative analysis is indicative of PSGD.

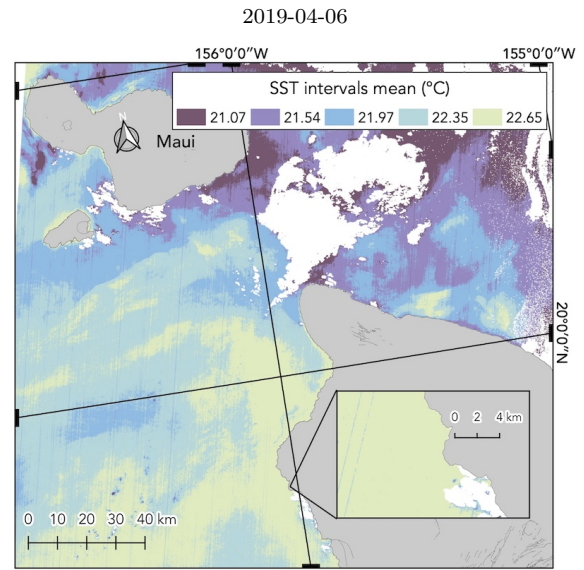
In terms of the distribution of the fraction of pixels where the sign of the derivative changes (Figure 2.6a), a striking feature is the much smaller number of pixels extracted as PSGD, when compared to the Ireland scenes, as also noted in the resulting maps (Figure 2.5b). Additionally, the fraction of pixels where the 2nd derivative over the green band is negative is the constraining factor, as the fraction where the derivative over the red is positive is considerably higher across all SST intervals, namely, from the L8 scene, this con-

Table 2.3: Summary statistics (size, mean, minimum, maximum and standard deviation) of the SST intervals derived from k -means clustering applied to (a) Landsat 7 and (b) Landsat 8 imagery, captured in the west coast of the Island of Hawai'i on 2007-04-29 and 2019-04-06, respectively. Brightness temperature is used as a proxy to SST. The 3rd and 4th SST intervals from the L7 scene are entirely comprised by one unique value, likely due to L7's 8-bit quantization. In addition, rounding to two decimal points obscures the differences between the maximum and the minimum SST values within the 3rd and 4th SST intervals from the L8 scene.

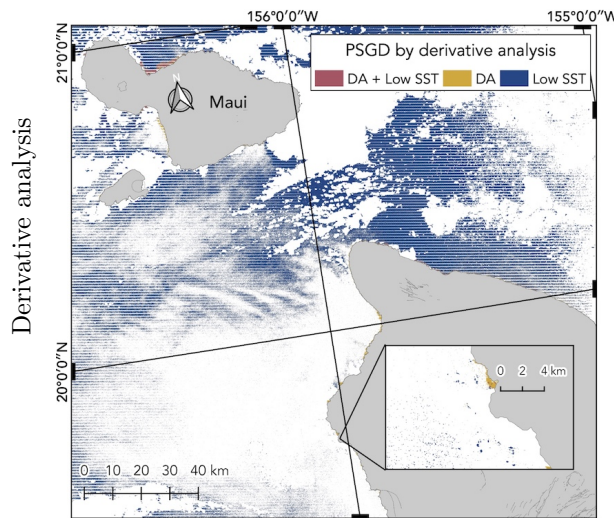
(a) 2007-04-29						(b) 2019-04-06					
k	Size (%)	Mean	Min.	Max.	σ	k	Size (%)	Mean	Min.	Max.	σ
1	15.5	19.59	14.02	19.68	0.19	1	11.3	21.07	15.97	21.30	0.23
2	31.1	20.13	19.97	20.26	0.14	2	19.4	21.54	21.31	21.75	0.13
3	22.7	20.55	20.55	20.55	0.00	3	18.3	21.96	21.75	22.16	0.12
4	18.1	20.84	20.84	20.84	0.00	4	31.4	22.36	22.16	22.50	0.09
5	12.6	21.20	21.13	29.46	0.17	5	19.7	22.65	22.50	35.34	0.21




(a-1)



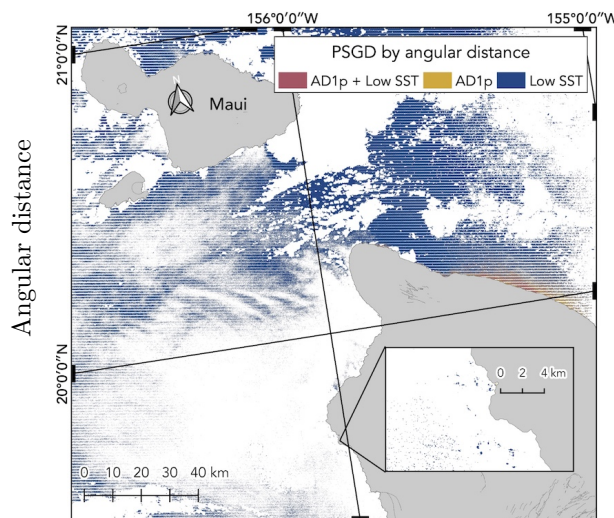
(a-2)




(b-1)



(b-2)



(c-1)



(c-2)

(Caption on next page.)

Figure 2.5: Results of the methods applied to imagery acquired over the Island of Hawai‘i on (1) 2007-04-29 from Landsat 7 and on (2) 2019-04-06 from Landsat 8: (a) sea surface temperature (SST; actual brightness temperature) intervals from k -means clustering; (b) potential submarine groundwater discharge (PSGD) pixels from derivative analysis of the visible spectra, showing locations where the sign of the 2nd derivative changed (DA; yellow), where SST was within the two coldest SST intervals (Low SST; blue), and where the two criteria co-occurred (DA + Low SST; red); and (c) PSGD pixels from angular distance (AD) applied to reveal the 1% points (AD1p; yellow) with the most similar spectrum to that of water with algal chlorophyll, where SST was within the two coldest SST intervals (Low SST; blue), and where the two criteria co-occurred (AD1p + Low SST; red). Inset shows details around Kailua Bay. White pixels represent no data points, masked out clouds, or points unrelated to PSGD. Dashed lines show bedrock faults [68].

dition is verified practically everywhere (Figure 2.6a–2). There are no major differences in the fraction of pixels flagged by the derivative analysis across SST intervals, although the second coldest interval has the higher fraction of pixels from both scenes.

2.4.2.3 Angular Distance

The general patterns of the AD results over both scenes in Hawai‘i were similar: very few pixels were extracted along the west coastline, either isolated or in small plumes—mostly single-pixel plumes (900 m^2 ; see the distribution of the plumes area in Figure A4b–1—, while the majority of them was located in the north coast, including near the coastline as well as farther offshore. Specifically, from the L7 image (Figure 2.5c–1), the pixels with the lowest 1% AD values that were located offshore of the northeast coast were mostly isolated and sparse, but less sparse and grouped together near the coastline. The few pixels that were extracted in the west coast were all adjacent to the coastline, isolated from one another and spread out over an extent of 60 m of the coastline, except for a two-pixel plume located on the Kealakekua bay. With the L8 scene (Figure 2.5c–2), the pixels with the lowest 1% AD values were mostly located over five regions: 1) an area of 7.2 km^2 in the Kawaihae bay; 2) an area of 9 km^2 in the Honokohau harbor; 3) two stretches in the south of the west coastline, where the extracted pixels are close to each other, although they do not form a multi-pixel plume; 4) a 16 km-long stretch near the north coastline; and 5) a large area offshore of the north coast. There were several other single-pixel plumes adjacent to the coastline.

In both scenes from the Hawai‘i case study, it is also noted that the AD values within the two coldest SST intervals are differently distributed than within the warmer intervals

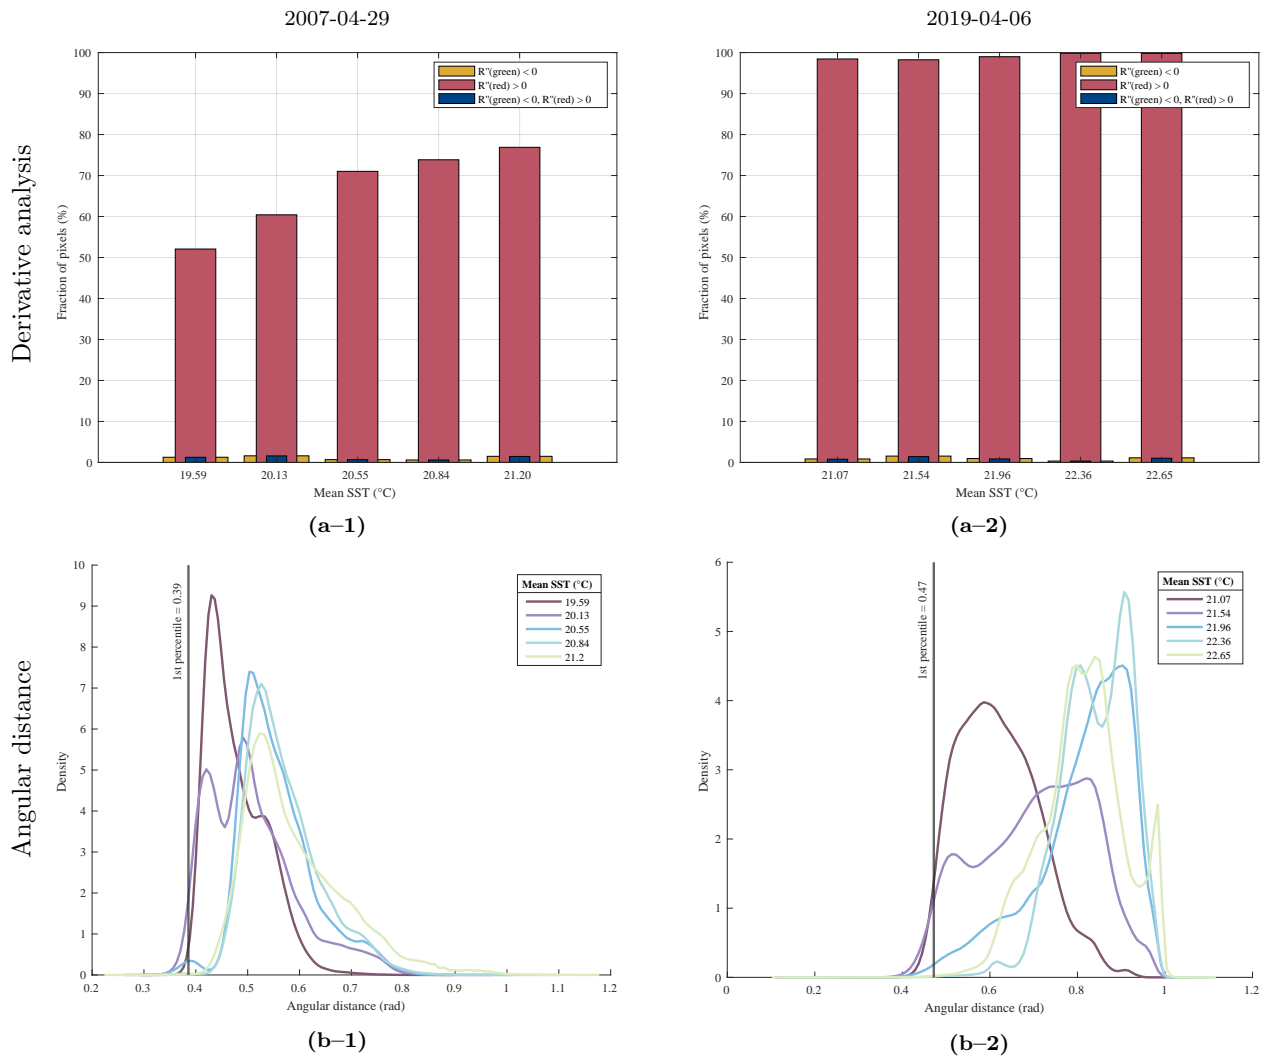


Figure 2.6: Distribution across SST intervals of the results of the methods applied to imagery acquired over the Island of Hawai'i on (1) 2007-04-29 from Landsat 7 and on (2) 2019-04-06 from Landsat 8: (a) derivative analysis and (b) angular distance (AD). Distribution of the derivative analysis results (a) shows the fraction of pixels where the 2nd derivative of the visible spectrum reflectance (R'') is negative over the green band (yellow bars), those where it is positive over the red band (red bars), and those where both of these conditions are verified (blue bars). Distribution of the AD results (b) shows the probability density estimate for the AD values within each SST interval, including the value of the first percentile used as threshold (vertical black line). SST intervals are represented by their mean values.

(Figure 2.6b), in particular within the coldest interval, which has the lowest mode. The distribution of AD values within the second coldest interval is more dispersed across the range of values, however, it still presents a lower mode and a higher number of pixels below the 1% threshold for both L7 and L8 scenes (0.39 rad and 0.47 rad, respectively).

2.5 Discussion

2.5.1 Using a Multi-Approach to Incorporate Color Data for PSGD Detection

We used different detection approaches as independent, parallel lines of evidence to support our findings of PSGD occurrences—derivative analysis and angular distance (AD). The developed algorithm digests RS images and applies the independent methods to extract PSGD pixels by exploring the relations between reflected and emitted water-leaving radiance (VNIR/color and TIR/temperature, respectively). Results from our case studies showed that the same method, when applied to different scenes, can yield differences in terms of: pixel connectivity (whether plumes are formed of individual or multiple adjacent pixels; see the distribution of plumes areas in Figures A3 and A4); proximity to known SGD sites (see Section 2.5.5); and percentage of extracted pixels located outside of expected areas (i.e., too far offshore; see the distribution of the distance from plumes to the coastline in Figures A1 and A2). The method based on derivative analysis, when applied to the Irish scenes (Figure 2.3b), extracted PSGD pixels within several large plumes, including some areas farther offshore. In contrast, the pixels extracted from Hawai'i (Figure 2.5b) were for the most part located near the coastline. Additionally, the criteria defined for different methods can lead to a different set of extracted PSGD pixels—in number and location—from the same scene. For all processed scenes, the method based on AD yielded far more restrictive results—i.e., fewer and less connected extracted pixels—than derivative analysis, likely due to the AD's specified 1st percentile threshold (Section 2.3.3; Figure A5 and Table A1 show other percentiles). Results from different methods to detect PSGD complement each other, forming an augmented set of PSGD pixels. This multi-method approach relies on the assumption that where extracted pixels are spatially coherent across methods PSGD is more likely to occur.

2.5.2 Pairing of Sea Surface Temperature and Color to Improve PSGD Detection

The separation of SGD from other possible drivers of SST anomalies within a single RS scene stands as a current challenge. SST anomalies can be driven by factors not related to SGD, such as upwelling, wind and ocean currents, and precipitation patterns. Despite these multiple drivers, SST is one of the most common SGD predictors in the literature, and the only one in RS studies [80]. Typically, in field applications, SST measurements are coupled with measurements of salinity, radon, and other natural tracers for SGD [9, 80]. However, these tracers cannot be reliably measured through spaceborne RS. Instead, in this study, surface reflectance in the VNIR range (color) is used to separate SGD-driven SST anomalies. In all four scenes, we observe several locations along the coastlines that were shortlisted as PSGD—i.e., marked for extraction—either based on SST anomalies or based on OC criteria (derivative analysis or angular distance), but not both simultaneously. The scenes over Ireland show examples of nearshore areas, bays and river mouths with pixels shortlisted as PSGD by OC criteria, but otherwise excluded for their warmer temperatures (yellow pixels in Figure 2.3), representing cases of chlorophyll presence likely linked to surface waters. Conversely, OC criteria excluded thermal anomalies located further offshore and limited the extent of some of the cold plumes emanating from nearshore (blue pixels in Figure 2.3), thus representing cases of SST anomalies possibly driven by factors other than SGD. Among all four scenes, cross-checking results from derivative analysis with the two lowest SST intervals reduced the number of shortlisted pixels by 40–60%, and by 15–55% in the case of angular distance (Table A2 shows the number and percent of shortlisted pixels in all scenes).

The case study of Ireland also illustrates the importance of combining SST with OC to identify PSGD locations that could be missed if only considering thermal anomalies. For example, there are a few PSGD plumes located southwest of Waterford Harbour, along the coastline—e.g., in the Dungarvan Harbour and in the Youghal Harbour—, whose corresponding cold plumes might seem insufficiently anomalous to be identified as potential sites for SGD, however, those cold plumes stand out when cross-checked with OC criteria and are flagged as PSGD. The Youghal Harbour example has the particularity of having a bedrock

fracture mapped along the middle of the estuary (dashed lines in Figure 2.3), around which we have identified PSGD plumes by both OC criteria. Previously, [92] have highlighted 35 PSGD sites—along the southern and western Irish coastlines—, based on visual inspection of thermal anomalies, however SST anomalies in the locations mentioned above were omitted as PSGD. These examples also illustrate that color data can be combined with SST to improve the assessment of what locations would better warrant field surveys.

As illustrated above, SST and color data can be combined to refine PSGD detection in a commutative manner. Particularly, results of the derivative analysis showed that the majority of the coastal waters surrounding the southwestern Irish coastline (L8 scene) had spectra characteristic of chlorophyll-enriched waters (yellow and red pixels in Figure 2.3b–2). Independent same-day observations from MODIS (Level-3 Daily Chlorophyll Concentration product, see Figure A8) corroborate with the spatial patterns observed in the derivative analysis results. This spatial coherence between the result from derivative analysis and the MODIS chlorophyll product supports that the derivative analysis method—which targets the spectrum of waters with algal chlorophyll—is an indirect method to detect high chlorophyll concentrations at higher spatial resolution than other well established remote sensing products. However, SGD is not the only—nor, necessarily, the major—driver of chlorophyll enrichment, so this result does not suggest that SGD might be occurring all along the Irish coast. In fact, by combining SST data with color data (through the derivative analysis approach) we show that only 59% (see Figure 2.4a and Table A2) of the pixels with spectra characteristic of chlorophyll enrichment are then extracted as PSGD (red pixels in Figure 2.3b–2). SST data constitute a valuable means to detect PSGD, and are commonly used for that end. Ocean color data are also useful for SGD studies, though not commonly utilized. More importantly, our study shows that combining SST and color results in a more reliable PSGD detection.

2.5.3 Refining PSGD Detection

PSGD delineations resulting from each method can be refined in a (optional) post-processing step. We can filter PSGD pixels (plumes) that are likely unrealistic or due

to noise and artifacts in RS images. The two methods presented yield results that can be visualized and analyzed utilizing the same heuristic as follows. First, recall that both methods work on a pixel-wise manner, therefore each pixel can be extracted as PSGD—i.e., selected—or not, thus creating a binary variable. Second, we can use the location of the extracted pixels, relative to the coastline and to each other, to further refine their selection (example in Figure 2.7). The interpretation of extracted pixels relies on the following spatial factors: 1) pixel connectivity—how extracted pixels form connected components, i.e., sets of pixels that are spatially continuous; 2) plume density—the number of plumes (defined as connected components) within a predefined area; and 3) distance to the coast—as plumes farther offshore are unlikely to result from SGD. Together, these factors constitute ways of further assessing the spatial nature of PSGD pixels.

Pixel connectivity, as usual in image processing literature, refers to how pixels are connected to each other. If two or more pixels are adjacent, they form a connected component. We use 8-adjacency, which includes pixels in the diagonal neighborhood. We refer to connected components as PSGD plumes, including single-pixel components, and where the plume area can be determined by the corresponding number of pixels (see the distribution of plumes area in Figures A3 and A4). Despite the spatial resolution of the RS scenes used, PSGD plumes may be represented by a large or by a limited number of pixels. SGD can rise to the surface with a small footprint (e.g., less than 10 m^2) [31], but most spaceborne sensors are not reliably sensitive to such occurrences. For example, both the L7 and 8 products used here have pixel areas of 900 m^2 . Therefore, a single extracted pixel might still represent a substantial PSGD plume. Additionally, since mechanisms driving SGD patterns at the sea surface are not fully understood nor easily predicted, PSGD occurrences identified with a single or a low pixel count cannot be automatically dismissed, despite having an inherent higher uncertainty than pixels that represent a larger PSGD plume [97]. Conversely, several adjacent pixels might represent a very large plume. Particularly, a unrealistically large PSGD plume could happen due to noise and artifacts in RS images, and, therefore, can be excluded from the set of the identified PSGD plumes. From our derivative analysis and AD results, we observed that defects in RS images commonly translate to a high number of extracted pixels in close proximity to each other, but mostly nonadjacent to other pixels. Therefore, we use

plume density to remove areas of dense non-adjacent pixel extractions as follows. We count plumes defined by connected components. Then, we divide RS images in square blocks that are 40 pixels (1200 m) wide. Plume density is the number of plumes in each block. If there are more than 5 plumes within a block, then all plumes within that block are marked for removal (example in Figure 2.7a–2). Since small plumes are of primary interest, we use the number of pixels that constitute each plume (connected components) to assess plume area. We assume that a large SGD plume can lead to a multi-pixel PSGD plume, as well as to other nearby and smaller plumes. To account for the relation between large and small plumes, each block is divided into quadrants (20 pixels wide), a simple partition approach. Then, if all plumes within a quadrant have an area of 10 pixels (9000 m² for L7 and L8) or less, plumes within that quadrant are effectively removed and are no longer considered PSGD plumes (example in Figure 2.7b–1). Small plumes will remain as PSGD if a large plume (composed by more than 10 pixels) is near. The size and shape of blocks, the maximum number of plumes within each block, the partition of blocks in 4, and the minimum plume area, can all be adjusted as warranted by conditions in the region of interest, and which we defined as working assumptions. A hydrogeological setting characterized by high hydraulic transmissivity, such as in the case of karsts areas in south Ireland, is expected to produce large SGD rates [28], which in turn would be linked to large plumes. However, no conclusive difference was observed in terms of plume area across scenes from both case studies (see more details in the supplemental), therefore, no additional hydrogeological information was incorporated at this step.

Globally, SGD is known to occur mostly within 55 km of the coast [50]. We use this information to exclude PSGD plumes located more than 55 km away from the coast. This considerably reduces the number of defect-originated plumes, for example, among all PSGD plumes in Ireland (L7), 38.4% are located beyond the 55 km mark (see the distribution of the distance from plumes to the coastline in Figures A1 and A2). We combine distance to the coast, pixel connectivity and plume density to produce a revised set of PSGD plumes. Note, however, that these post-processing steps are not required to use the methods presented here, which are simple yet effective; instead, these refining steps optimize the methods results, removing known issues arising from data noise and the limits of spaceborne observations.

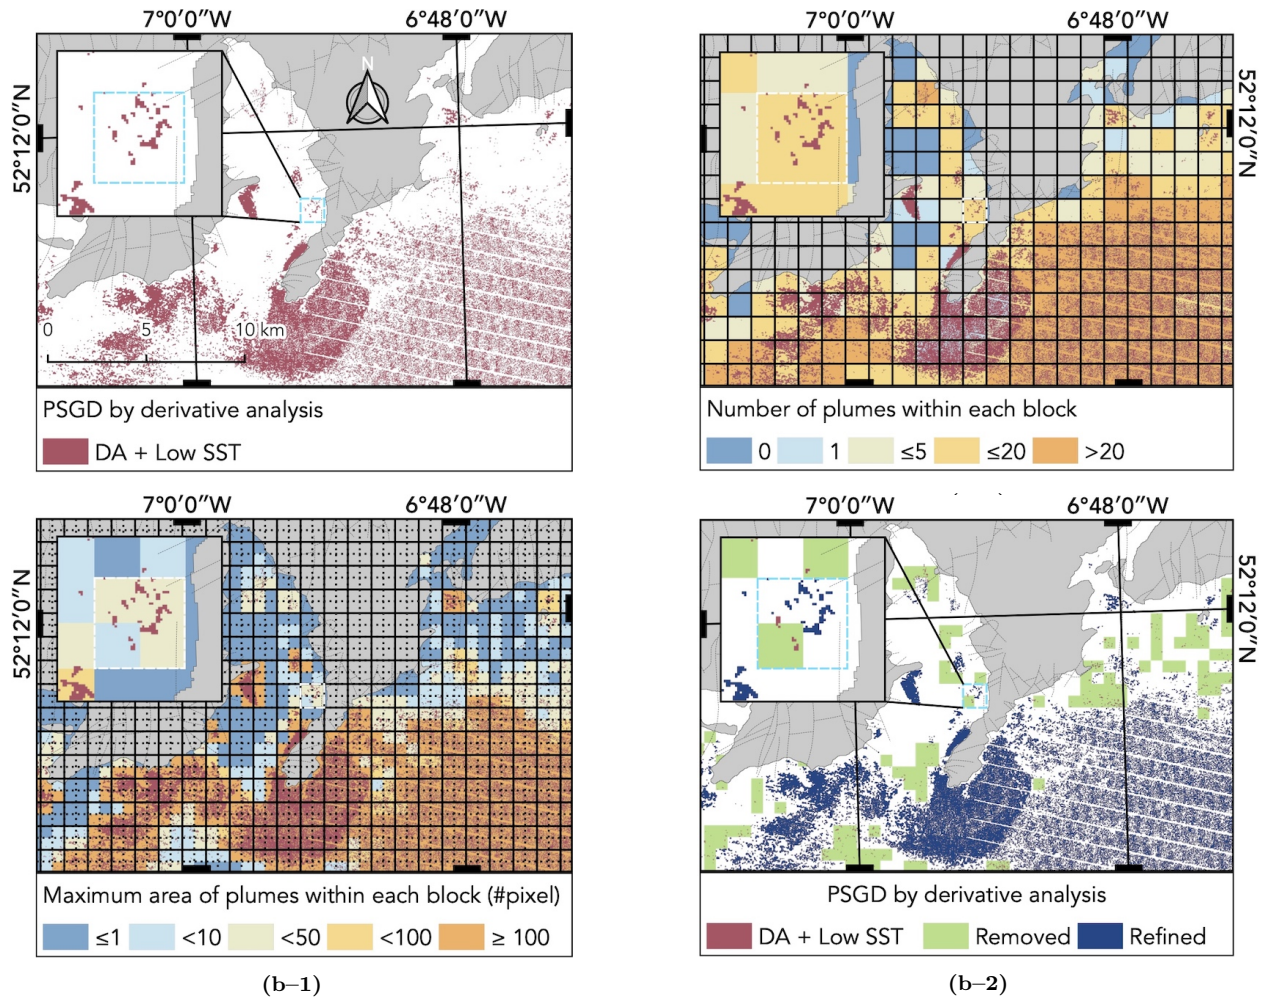


Figure 2.7: Post processing of potential submarine groundwater discharge (PSGD) plumes resulting from applying the derivative analysis method to imagery acquired over Waterford Harbour, south Ireland (see the inset in Figure 2.3b-1), on 2010-06-02 from Landsat 7. a) PSGD pixels from derivative analysis of the visible spectra, showing locations where the sign of the 2nd derivative changed and SST was within the two coldest SST intervals (red pixels). b) Plume density, given by the number of plumes (connected components) within distinct 40-pixel wide blocks (solid grid). Plumes are shortlisted for removal where there are more than 5 plumes. c) Maximum area (number of pixels) of the plumes within quadrants (20-pixel wide, dotted grid) of the same blocks (solid grid). Plumes are shortlisted for removal where the maximum area is at least 10 pixels. d) Plumes that were shortlisted both by plume density and maximum area are removed (green shades), resulting in a refined PSGD data set (blue). The inset shows details within the 40-pixel wide block outlined by the dashed rectangle.

2.5.4 Identification of Scene and Local SST Anomalies

The process of identifying SST anomalies depends on whether it is applied to an entire scene or to a subregion of a scene (local anomalies). In this study, the methods are designed to be platform agnostic and applicable to RS images regardless of spatial coverage and of

the definition of a subregion. For the case studies analyzed, we processed entire scenes to demonstrate the potential for analyzing PSGD on a regional scale, and for implementation as an operational tool without user identification of scene subregions. Both the derivative analysis and the AD methods work on a pixel-wise manner, making them independent of scene area. On the contrary, the definition of the SST intervals depends on the area of the scene, since SST range and variability are a function of area [78] (among other factors). However, as discussed below, local SST anomalies may exist outside of the coldest SST interval. For this reason, we adopted the maximum in the two lowest SST intervals in the scene (Table 2.2 and ??) as a generally applicable threshold for PSGD identification.

Local and scene SST anomalies can be expressed differently in the SST intervals. In Ireland, SST intervals derived from the L8 scene (Figure 2.3a–2) revealed two offshore, very large bodies of cold water (approximately located in the center of the scene) that belonged to the coldest interval. Another large anomalous plume in the same scene was located in the northeast region. Those three bodies of cold water are SST anomalies unlikely related to SGD, since they occur or extend beyond 55 km from the coastline, and at depths ranging 100–200 m, where significant mixing with sea water would have most likely occurred and obfuscated any eventual thermal signature in the sea surface [47, 28]. Local SST anomalies near the coastline fell in the same SST interval as those three bodies of cold water. However, those smaller, near-shore SST anomalies were also extracted as PSGD through the derivative analysis method (red pixels in Figure 2.3b–2), whereas the larger bodies of cold water were not. Comparatively, local and scene SST anomalies may fall in different SST intervals. For example, in Hawai'i, the largest SST anomaly in both scenes (Figure 2.5a), located offshore of the north region of the island, comprised the entirety of the lowest SST interval. This caused the smaller, local SST anomalies on the western shore of the island to fall in a higher SST interval (2nd and 3rd intervals; see Table 2.3). Smaller point-sourced SGD fluxes also tend to produce plumes that are colder than surrounding waters, however they may be as cold or less cold than other, unrelated thermal anomalies (e.g., moored ships and shoreline vegetation) in the same bay [31]. This example also shows that limiting PSGD extraction to pixels within the lowest SST interval may be too constraining, whereas using the two lowest intervals in the scene is a more adaptive approach. Our methods allow the distinction

among SST patterns of different spatial scales, thus enabling the identification of both local and scene-based SST anomalies without defining subregions.

2.5.5 Validation Using Data from the Case Studies

We compare the identified PSGD plumes with observations from the case studies presented here. Previous studies have taken field measurements of natural tracers for SGD: radon (^{222}Rn) activity in Ireland [92] and salinity in Hawai'i [27] (Figure A6 shows the distribution of these datasets). In this study, survey data are thought to represent tendencies for localized SGD, and reapplied here to document PSGD locales. The generalization of the location of the measurements from past studies comes as a consequence of the challenges of temporal matching RS images and the field surveys [28]. For that reason, we incorporated all in situ measurements into a single data set for each case study, regardless of their acquisition dates, and compared them against the corresponding L7 images, which are relatively close temporal matches.

We applied a simple approach for the validation assessment, defining quantitative measures that are applicable regardless of the case study, namely across different sampling strategies—including temporal and spatial distribution of samples. First, we calculate the distance between PSGD plumes (centroids) and sampled locations. We use this information to select PSGD plumes that are within 1 km of the samples, since it is desirable to only include the plumes that are potentially related to the surveyed areas. Considering this reasonably permissive cut-off of selecting plumes, we then compute the distance between plumes and samples and the corresponding sample values. That allows us to assess whether or not plumes are closer to samples that are strong indicators of SGD, i.e., samples with values of high radon activity or low salinity. In this analysis, we used the PSGD plumes refined as described in Section 2.5.3.

In Ireland, L7 data were acquired on 2010-06-02, whereas field surveys were conducted in the 5th and 11th of 2010-08, gathering a total of 55 data points that covered approximately 100 km² of one bay (Waterford Harbour). Measurements of radon activity were retrieved at 1 m depth [92]. Considering derivative analysis, our results found 682 plumes (including

single-pixel plumes) within 1 km of the samples (see the distribution of the distance between plumes and samples and their corresponding value in Figure A7). Among this subset of plumes closer to sampled locations, the median of linked radon activity is 5 Bq m^{-3} . Note that these measurements reflect radon activity in excess to ambient levels within the region [92]. The majority of these plumes are small (up to 2 pixels). There is one large plume of approximately 17 km^2 which extends beyond the 1 km around sampled locations. All of the 15 nearest PSGD plumes from AD are comprised of 1 pixel, and the linked samples have a median value of 4 Bq m^{-3} . Despite that PSGD plumes do not seem to be exclusively concentrated in the vicinity of high radon activity points, we consider notable that several plumes are located close to this area known for SGD. Radon tracers for SGD are known to include circulated seawater [55], which may differ from meteoric SGD in terms of temperature and nutrients [77]. On the contrary, here, PSGD plumes capture the thermal and nutrient-induced signatures from meteoric SGD. This difference could explain the several AD plumes for which the nearest radon samples were kilometers away.

In Hawai'i, L7 data were acquired on 2007-04-29, whereas we used data from field surveys that were conducted over 10 different dates between 2005-08 and 2007-05, gathering a total of 67 data points that covered approximately a 100 km coastline stretch over four bays (Kailua, Kealahou, Kiholo, and Honokohau Harbor, see Figure 2.1). Measurements of salinity were retrieved up to 15 cm depth [27]. In this case study, a small number of plumes were located within 1 km of the samples, 10 from derivative analysis and 2 from AD (Figure A7 shows the distribution of the distance between PSGD plumes and field samples). The reduced number of plumes near the samples are likely a consequence of the observed cloudiness in the RS image over the sampled bays (note the white patches in Figure 2.5a), in addition to the temporal gap between the in situ sampling and the L7 overpass. Both methods produced relatively small plumes in the region of the samples (half of the 12 plumes have up to two pixels, and the largest one has 17 pixels). 6 of the plumes resulting from derivative analysis and both of those derived from AD have a reduced value of salinity linked to them (≈ 30 psu).

In relative terms, PSGD plumes near sampled locations in Ireland tend to be primarily linked to samples that are within the lower range of excess radon activity (Figure 2.8).

However, the distribution of the values assigned to PSGD plumes resulting from derivative analysis (red bars in Figure 2.8a) generally match with field samples (blue bars in Figure 2.8a).

The fewer PSGD plumes near measurements from Hawai'i are relatively harder to compare in terms of distribution of the linked samples. Nevertheless, they also tend to be linked to strong indicators (reduced salinity samples). We conclude that PSGD plumes from both methods contribute effectively to map areas of where SGD is likely to occur.

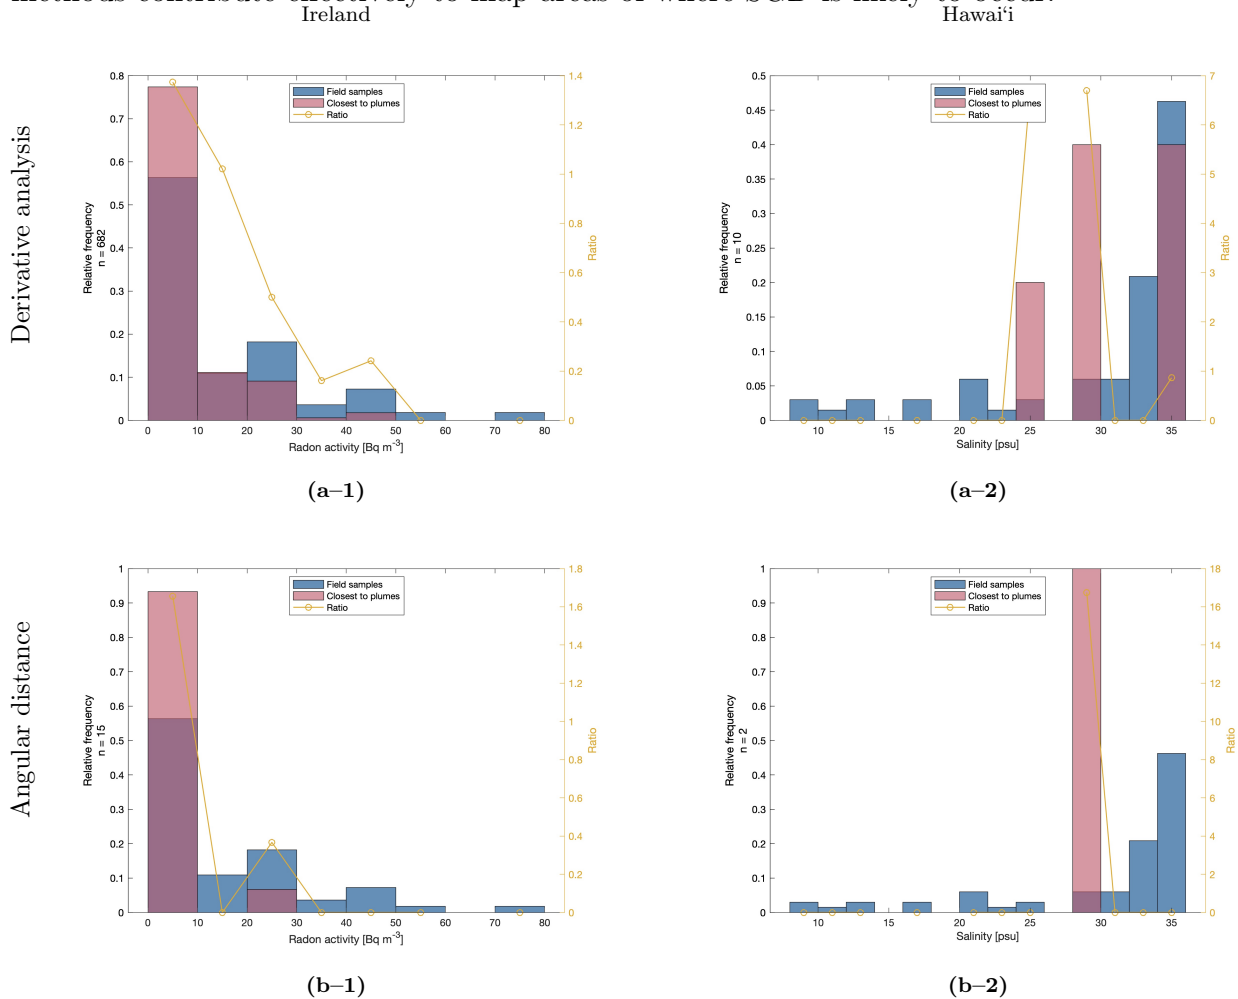


Figure 2.8: Distribution of radon activity (Bq m^{-3}) and salinity (psu) from field observations in (1) Waterford Harbour, Ireland [92], and in (2) four bays of the Island of Hawai'i [27], respectively. The complete distributions of the values from the field samples (blue bars) are compared to the distribution of the values assigned to the corresponding nearest PSGD plume (up to 1 km away; purple bars) from (a) Derivative analysis; and (b) Angular distance, both applied to the Landsat 7 scenes. Yellow lines and axes show the ratio of the relative frequencies of the field values to the relative frequencies of the plume values. The label in the vertical axis show the number of plumes considered (n).

2.5.6 Methods Limitations

With this study, we provide an improvement over the current techniques for SGD detection by striving to accurately capture PSGD. However, there may be SGD occurrences not detected through our methods. Those missing occurrences are likely due to standard RS limitations [28]. Color (VNIR) interpretations from derivative analysis and angular distance rely on a single spectrum of observed water-leaving radiance in a lab setting. Ideally, one would use a spectra library with multiple samples, including from field surveys. In addition, the SGD sites in Ireland and Hawai'i are specific locations measured at specific moments in time. In the assessment of our methods, particularly for the most recent scenes (using Landsat 8), we assume that SGD is perennial in those locations. Remote sensing observations from the sea surface are inherently tied with hydrodynamics, such as currents and upwelling events, as well as other common challenges like sun glint and bottom reflectance (in shallow waters). While we recognize that our results are susceptible to interference from those issues, we believe that the necessary additional work, which is outside of the scope of the current study, should not pose a major scientific obstacle to wider application of our methods.

2.6 Conclusions

In this study, RS data are leveraged to improve our ability to locate PSGD occurrences. Typically, only thermal data are used in RS applications for SGD, relying then on in situ, harder to obtain data to further assess the locations shortlisted by temperature analyses. We demonstrate that RS-observed thermal and color data can be combined to enable a more unambiguous PSGD detection, and therefore provide a higher level of scrutiny when determining the PSGD locations that warrant further investigation. Additionally, the methods presented are applicable to RS scenes from different regions of the globe and of varying covering areas, thus contributing to the global mapping and monitoring of PSGD.

3.0 Spatio-Temporal Patterns in Remote Sensing Signatures of Submarine Groundwater Discharges in the Southern Irish Coast

3.1 Introduction

The role of submarine groundwater discharge (SGD) in the global water cycle continues to be recognized [50, 64], yet we still lack frameworks for global SGD monitoring and assessment. SGD flows are also relevant at estuarine scales, as they can be of the same magnitude as large rivers (e.g., $10^7 \text{ m}^3 \text{ d}^{-1}$) [11, 87]. Therefore, we are in need of long term, cost effective solutions to determine where, when, and how much SGD occurs.

This study evaluates the consistency of potential SGD (PSGD) signatures—i.e., repeated evidence—across time and leverages long term, frequent remote sensing imagery to detect seasonality of SGD. Most remote sensing-based methods to identify SGD only use a single snapshot in time over a target location. Our method captures important variability that single snapshot/static approaches cannot capture, additionally to being less sensitive to noise and confounding factors. Limited temporal analysis leaves knowledge gaps with respect to coastal aquifer budgets. Major SGD occurrences may have significant seasonal variability that impacts coastal aquifer budgets [16]. Our seasonal analysis of PSGD contributes to filling this gap. With a multi-scene approach, a temporally consistent signal would suggest a higher potential SGD (i.e., higher certainty of true positives), and less likelihood of it being a false positive. A false negative is still possible, though. Either because the SST or color signal could not be sufficient to trigger detection (which would be the same issue on the single snapshot approach), or because it was not observed enough times to satisfy the temporal consistency. (consistency threshold)

Using the algorithms developed by [12], we aim to document seasonal SGD as a function of PSGD signatures, including thermal anomalies—as observed through emitted thermal infrared (TIR)—and patterns in the visible and near infrared (VNIR) spectrum. We create a time series of detected PSGD by mapping the PSGD signatures over multiple remote sensing scenes. Thermal anomalies are inferred through sea surface temperature (SST) intervals, and

spectral patterns are found through derivative analysis and angular distance. We showcase this application over south Ireland (path 206, row 24) and leverage the Landsat 8 catalog to improve our assessment of PSGD. TIR is used based on the expected temperature differences between groundwater and seawater. VNIR is used based on the fact that nutrient-rich SGD can cause changes in the ecosystems, namely through the contribution to an increase in primary productivity (growth), which may increase algae growth and chlorophyll content. These are then linked to changes in the surface reflectance in the optical spectrum, hence color.

Typically, remote sensing applications for SGD detection exploit scene-level SST gradients driven by groundwater influence. Mallast et al. [46] follow a somewhat different approach. They use the reasoning that groundwater temperature is almost constant throughout the seasons, whereas surface water temperature is more variable, changing in function of air temperature and solar radiation. Accordingly, a location where groundwater is discharging would observe an attenuated variation across seasons, comparatively to non-affected areas, which would observe a wider variation. The decreased SST variability was assessed by applying pixel-wise metrics (e.g., median, range) across the entire data set. Mallast et al. [46] also pioneered the analysis of the temporal variability of SGD, linking plume area variations on a biennial scale to measured discharge rates. Our approach extends the temporal analysis to the inter-seasonal scale, as we utilize the temporal consistency of PSGD signatures across seasons and across the 9-year study period as a whole. First, we identify PSGD by scene, and then compare results across seasons, resulting in a temporally finer scale (seasonal). This approach extends what can be captured by summarizing an entire data set by applying a pixel-wise metric. Applying a metric is a way to leverage signatures of perennial SGD events to better identify them, although intermittent events are less likely to be captured that way. Therefore, it is useful to evaluate PSGD time series through both approaches (with and without the seasonal aggregation). Temporal consistency is applied pixel-wise. Because PSGD can reach the sea surface at different specific locations over time, our approach targets the general identification of PSGD sites, and is not intended to obtain precise estimates of the surface area of plumes. However, a change in surface area (number of PSGD pixels) over seasons in a given location can correspond to a difference in SGD flow.

This work leverages remote sensing data to capture spatial and temporal PSGD patterns derived from parallel algorithms based on SST and color, and to map high PSGD sites. We identify regions where PSGD persist through time, as well as those that can occur intermittently. However, a long and temporally homogeneous data set is required to determine whether PSGD is perennial or intermittent. In particular, PSGD may appear intermittent as a consequence of biased scene acquisition times, which can favor some seasons over the others. We also tackle two questions arising from seasonal differences in PSGD occurrences: 1) how do they arise, suggesting a link to groundwater levels and seasonality of groundwater recharge [26], and, 2) how does the inversion of thermal gradients and the attenuated signature as a consequence of a reduced flow impact our ability to detect PSGD.

3.2 Materials and Methods

3.2.1 Observed and Expected Spectra in Irish Coastal Waters

The western and southern coastlines of the Republic of Ireland are encompassed by many islands, peninsulas headlands and bays [92]. The Irish coastline has been the target of past studies looking to identify SGD occurrences [92, 91]. Wilson and Rocha [92] utilized Landsat 7/ETM+ thermal infrared data to map 35 potential sites of SGD, providing ground-truthed observation in one of those sites (Hook Head, Wexford Co.) that confirmed the occurrence of SGD in that area (around the headland and in the embayment next to it, the Waterford Harbour). We have developed our PSGD detection methods with a study on the southern Irish coast [12], using the ground-truthed site to calibrate and assess the efficacy of our methods. Whereas in our previous study we only utilized two scenes (one from Landsat 7 and one from Landsat 8), here we take advantage of using multiple scenes across time.

3.2.1.1 Sea Surface Temperature and Color Observations from Landsat 8

We collected 473 Landsat 8 overpasses over the southern Irish coast (path 206, row 24) between 2013-05-30 and 2021-09-28. For the purpose of computational efficiency, we

restricted scene selection as a function of cloud coverage: 62 scenes out of the 473 were selected (see the list of product IDs in Table B1) based on cloud coverage (less than 50% within 5 km of the coastline; clouds derived from the Level-2 quality band), aerosol content (low or medium), and data quality (Tier 1).

Color and sea surface temperature (SST) data were obtained as surface reflectance and surface temperature, respectively, both from the Level-2 products of the Landsat Collection 2 catalog and were sourced from USGS.

3.2.1.2 Expected Thermal Signature

We explored temperature and density differences between coastal sea waters and groundwater to reveal thermally contrasting SGD plumes. Groundwater, and consequently SGD plumes, can be either warmer or colder than sea water, whereby buoyancy is mainly driven by density[58], thus fresh or brackish SGD plumes float to the surface regardless of temperature. In the south Ireland, groundwater temperature is not always colder than seawater, as both show seasonal variability. Groundwater temperature has a lower variability than SST, in Ireland generally within the range 9.5–11 °C [1], whereas SST in the Celtic Sea can be in the range 8–18 °C (see SST time series from a buoy in the Celtic Sea in Figure B6).

Our methods rely on the identification of SST anomalies in coastal waters, which therefore can be colder or warmer than the surrounding seawater. To determine whether SGD-induced SST anomalies are expected to be colder or warmer than seawater, we applied a systematic comparison between groundwater temperature, as measured in quality monitoring wells, and seawater (surface) temperature, as derived from satellite TIR observations. We use daily groundwater temperature data from 20 boreholes in the Irish National Water Monitoring Stations, sourced from the Irish Environmental Protection Agency (EPA; Resource identifier: EF.IE.EPA.MON.Waterstations; see the location of the boreholes in Figure B5). Groundwater temperature in this subset is mostly within the range previously identified—apart from outliers—therefore we assume temperature measurements are representative of the tapped aquifer, despite the known occurrence of geothermal springs in south Ireland (e.g., near Cork, Tramore, and Wexford [2]). We resample the daily records to 8-day

percentiles (10th, p_{10} , and 90th, p_{90}) and reindex them to match the dates of acquisition of the Landsat 8 scenes, assigning the nearest (in time) EPA 8-day record to each scene, with a 30-day tolerance. These EPA records have periods of time without valid measurements, spanning several months at times (see the groundwater temperature time series in Figure B8). Whenever there is no 8-day record within 30 days of each scene acquisition time, that scene is excluded from our analysis, as we cannot reliably infer the expected thermal signature from SGD (colder or warmer than SST). Accordingly, the 62 scenes subset is further reduced to 29 scenes: 10 in autumn, 7 in spring, 10 in summer, and 2 in winter (see Table B1).

For the purpose of comparing groundwater temperature with SST, we use a Level-3 MODIS product of 8-day SST composite (AQUA_MODIS, L3m.8D.SST.sst.4km, Masked, SMI, NASA GSFC OBPG, R2019.0, Global, 0.041 66°, Science Quality), SST_{MODIS} , sourced from the PO.DAAC at JPL, NASA. This product is used instead of the Landsat 8 surface temperature for two reasons: first, this Level-3 product is calibrated to represent absolute SST, therefore it is more accurately comparable to the EPA groundwater temperature measurements (this was verified by comparing SST_{MODIS} to ground-truthed SST from a buoy, see Figure B6); second, it has a coarser spatial resolution—here also aggregated on a relatively coarse temporal resolution—and therefore provides a representation of the local average SST—from the point of view of the 30 m Landsat 8 pixels—and a buffer against possible SGD-induced SST anomalies. We matched each of the 29 Landsat 8 scenes with one MODIS SST scene, which are then reprojected to the same projection, resolution (through nearest neighbor downscaling), and spatial extent of the Landsat 8 scenes. Then, for each scene, we compare each SST_{MODIS} pixel to the groundwater temperature percentiles, such that if $SST_{\text{MODIS}} > p_{90}$, then SGD is expected to produce a colder plume in the area corresponding to that pixel (green pixels in Figure 3.1), and if $SST_{\text{MODIS}} < p_{10}$, then SGD is expected to produce a warmer plume in the same area (yellow pixels in Figure 3.1). If $p_{10} \leq SST_{\text{MODIS}} \leq p_{90}$, then we assume the temperature gradient between groundwater and seawater is insufficient to reliably detect the corresponding SST anomaly, and therefore that pixel is excluded from being flagged as PSGD (blue pixels in Figure 3.1, which also represent missing data in SST_{MODIS}). These steps produce maps—one for each Landsat 8 scene—of

the expected SGD thermal signature. However, because they are derived from the coarser spatial resolution of SST_{MODIS} , nearshore areas tend to have missing data, as pixels over land are masked out in SST_{MODIS} . To minimize the data loss that would be induced from excluding the nearshore areas—which are prime locations for SGD occurrence and detection—, we make the assumption that the expected SGD thermal signature in those coastal pixels is the same as the nearest SST_{MODIS} pixel, up to 3 km (100 Landsat 8 pixels). This approximation is done through nearest neighbor interpolation (which introduces artifacts in the spatial continuity, as revealed by irregular shapes in Figure 3.1, for example, in the cold/warm interface in spring).

3.2.1.3 Expected Color Signature

As in Chapter 2, here we use a reference reflectance signature characteristic of chlorophyll production in coastal waters, derived from an experiment conducted under a controlled environment, whereby spectra were captured from clear water, suspended sediments and algal chlorophyll, thus assessing changes in the water surface reflectance due to the presence of sediments and chlorophyll content [24]. The authors showed that the presence of algal chlorophyll produces a feature in the second derivative of the reflectance spectrum, whereas clear water and suspended sediments alone would not manifest in the second derivative (near zero). In our application, we explore the signature introduced in the second derivative to identify coastal waters that may have been enriched with algal chlorophyll, and therefore are potentially under the effect of SGD. We assume the induced signature is verified across different conditions (atmospheric conditions, incident light, different sources of chlorophyll, and different water properties such as salinity and temperature).

3.2.2 Detection of Potential Submarine Groundwater Discharge

We leverage SGD linkages to patterns in SST and OC to determine SGD locations by analyzing the observed optical and thermal infrared spectra of coastal waters. These data are used to infer thermal anomalies and to find spectral patterns attributed to SGD. Both approaches applied here are described in greater detail in [12].

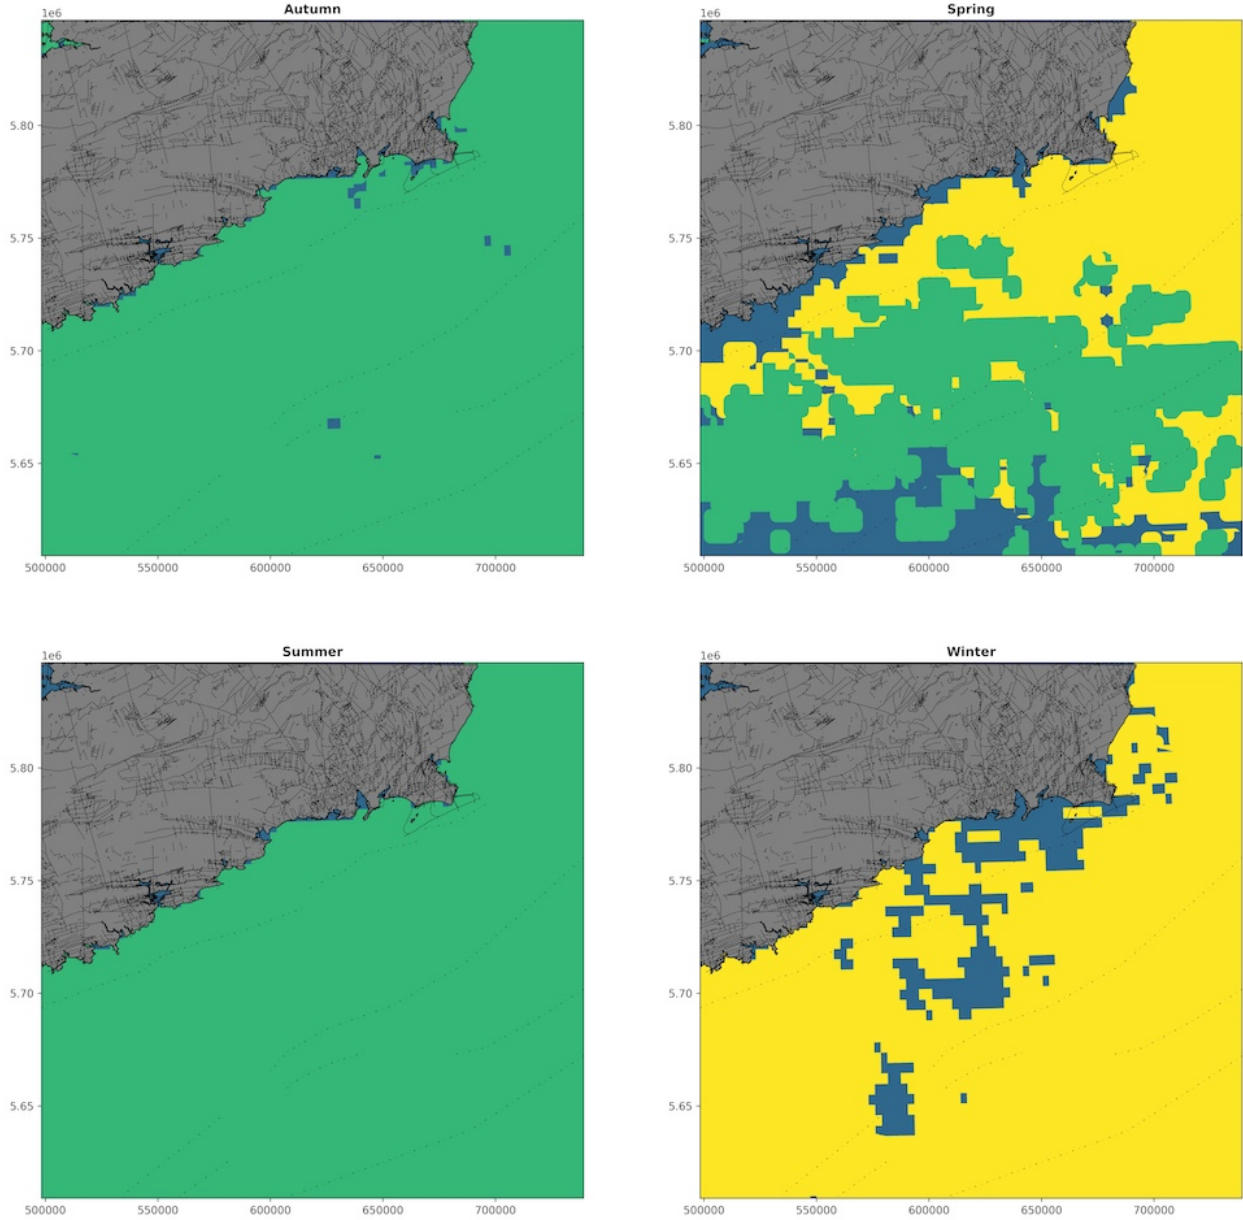


Figure 3.1: Mode of the expected thermal signature of PSGD plumes in each season: cold plumes (green), warm plumes (yellow). The expected thermal signature is determined by comparing groundwater temperature to SST observations from MODIS. The thermal signature is undetermined where there are missing data in MODIS scenes, or where the gradient between groundwater temperature and SST is diminished (blue pixels). Note that the mode computation includes undetermined signature, therefore it does not imply that PSGD plumes cannot be detected in any scene.

First, SST anomalies are identified using a clustering method, a k -means variant optimized for univariate applications [86]. For each scene, clustering classifies the SST values into five temperature intervals (1 to 5, with relation of order preserved). Then, SST intervals

are compared to the expected SGD thermal signature, such that if a pixel is in the coldest (warmest) SST interval and SGD is expected to be colder (warmer) than seawater, then it is flagged as an SST anomaly. We note that in our previous study, we had applied a less stringent criterion to identify SST anomalies, as we considered the two coldest SST intervals, instead of only the coldest one. Second, color patterns linked to SGD are identified using derivative analysis (DA) and angular distance (AD). Here, derivative analysis is used to flag pixels whose spectra have an inflection point between the green and red bands—a spectral signature linked to SGD. DA results is a binary map of the pixels where the spectral signature—inflection, i.e., change of sign of the second derivative—is observed. AD is used as a measure of similarity between the observed visible spectra in each scene and the expected spectrum for water with suspended sediments and chlorophyll content [24], as a proxy to SGD influence, resulting in a binary map of the pixels whose spectra are the most similar—the top 5%—to the expected spectrum. Here, we have also adjusted the sensitivity of AD in comparison with our previous study, where we had applied a more conservative percentile threshold (1%). We define PSGD pixels as those simultaneously flagged as SST anomalies and flagged by either DA or AD. Here, we take one additional step and refine PSGD pixels as those flagged by both DA and AD (each set already crosschecked with SST anomalies).

These methods are applied scene-by-scene, resulting in a time series of images showing PSGD pixels. This time series is then analyzed in two ways: 1) full time series; and 2) seasonally aggregated time series. In both cases, we apply a temporal consistency criterion to further refine the identified PSGD occurrences.

3.2.3 Temporal Consistency of PSGD Plumes

When PSGD is detected from a single image, false positives are highly likely, as noise and artifacts from remote sensing data propagates through the applied methods. Here, we take advantage of repeated observations to evaluate temporal consistency and diminish the influence of those confounding factors. Therefore, in both full and seasonal time series analysis, we further refine the set of PSGD pixels using the number of times a pixel is flagged

as PSGD, evaluated as a percentage of the number of valid observations in the same pixel. Specifically, we assess pixels that are flagged as PSGD more than 50%, 60%, 70%, 80%, and 90% of the valid observations in the time series.

In the full time series analysis, the collection of 29 scenes is summarized in one single image as a result of applying the temporal consistency criterion. We sum the number of times each pixel is flagged as PSGD—either by DA, AD, or the intersection of the two (also being crosschecked with SST intervals)—and divide by the number of valid observations in each pixel out of the 29 selected scenes (pixels may have different number of valid observations depending on cloud coverage), resulting in the PSGD incidence percentage for each pixel.

An identical calculation is applied to the collection of scenes when divided by season, where seasons are defined by equinoxes and solstices in each year, resulting in four seasonal time series: autumn, spring, summer, and winter. For each season, the set of PSGD pixels is refined by retaining only the pixels that are flagged as PSGD more than 50% of the observations in the time series—e.g., in a given season, a PSGD pixel is retained if it was observed 10 times during that season, and flagged as PSGD at least 6 times. We also assess this threshold percent at 60%, 70%, 80%, and 90% of the valid observations in the time series.

Because the temporal consistency criterion is a function of the number of valid observations, we combine it with the requirement of a minimum number observation to validate the computed percentage: each pixel should have at least five observations. Accordingly, a PSGD pixel that satisfies the temporal consistency criterion is flagged as PSGD at least three times.

A note about the visualization of PSGD plumes in our figures: because small plumes (e.g., 1–5 pixels) might not be easily visible, we add circles around the plumes. The radius of these circles is defined by the minimum enclosing circle radius plus 40 pixels, thus each circle has a minimum radius of 41 pixels. There is also a limitation in the way we computed the number of PSGD plumes (connected components): plumes at different incidence percentages always considered distinct plumes, even if they are adjacent.

3.2.4 Plume Buffering

PSGD plumes are moved by currents as they rise to the surface, and there is no guarantee they emerge at the same exact surface location over time. For PSGD plumes with surface area smaller than the spatial resolution of the remote sensing imagery—30 m for Landsat 8—, the PSGD plume may still consistently surface in the same pixel location. However, it is still possible that the same PSGD flow resurface at different pixel locations, particularly for seepage occurring at greater depths, i.e., with an increased vertical travel distance.

To account for plume movement, we apply a spatial buffering of PSGD pixels on a scene-by-scene basis, and only then apply the temporal consistency criterion. Note that buffering only takes place after crosschecking DA or AD with SST intervals; otherwise there is an additional increase in noise and false positives. For the refined PSGD derived from crosschecking DA with AD, the crosschecking takes place after the plume buffering. We set the buffering to a 5-pixel radius disk (150 m at Landsat 8 resolution). We chose 5 pixels as the likely maximum plume area: from the two scenes over the southern Irish coast we previously studied[12], we observed that plumes with an area up to 5 pixels accounted for at least 80% of the PSGD plumes (derivative analysis) and up to almost 100% (angular distance; see Figure A3 in Appendix A). This plume buffering acts like a naïve plume tracking approach, that enables accounting for the fact that SGD plumes are not tied to the same exact location and can move like a loosely anchored buoy, and although it may introduce some false positives, it also amplifies features and has the advantage of being applicable regardless of the time interval between scenes. The minimum number of valid observations applied with the temporal consistency also helps to mitigate noise introduced by buffering.

3.3 Results and Discussion

3.3.1 Temporal Consistency of PSGD Plumes

Our analysis of spatial and temporal consistency of PSGD plumes identified 15 areas along the southern Irish coast with potential for SGD, accounting for a total of 127 mapped

plumes (Figure 3.2).

PSGD plumes identified here correspond to 6.4% of the pixels where DA and AD plumes intersect, i.e., only that subset of plumes satisfies the temporal consistency criterion of being flagged as PSGD in more than 50% of the observations (see the distribution of the number of flagged PSGD pixels in Figures S35–38). This subset consists of 457 pixels (grouped in 127 plumes), which satisfy all the detection criteria simultaneously: 1) SST anomaly—i.e., they belong to the SST interval correspondent to the relative groundwater temperature (colder or warmer than seawater); 2) minimum number of valid observations—i.e., there are at least 5 observations in each pixel’s time series; and 3) temporal consistency—i.e., they are flagged by DA and AD in more than 50% of the observations (which when combined with the minimum number of observations criterion, means that they are flagged as PSGD at least 3 times). These PSGD plumes are mostly located nearshore, highlighting 8 coastal areas where SGD is more likely to occur (West to East): Youghal Harbour (Youghal Bay), Dungarvan Harbour, Copper Coast, Tramore Bay, Waterford Harbour, Bannow Bay, Saltee Islands, and Wexford Harbour. In Dungarvan Harbour, more specifically south of the Ring Peninsula, 3 plumes with higher percent incidence (60%, 70% and 80%) provide stronger support for PSGD in this location. PSGD plumes in all other locations have an incidence of up to 60%. It is worthwhile to note three plumes identified in the coastline of the Hook peninsula (Waterford Harbour), an area where SGD occurrence has been confirmed through in situ measurements of natural tracers [92]. The southern Irish coastline is rich in fractures, faults and karsts (see dashed lines in Figure 3.2) which are a common pathway for SGD, and are likely linked to the majority of the identified plumes that fall on or nearby them.

These results are conditioned by data availability, both temporally and spatially. The scene selection based on cloud coverage and on matching with groundwater records likely introduced a bias in the available data. The set of 29 Landsat 8 scenes is heterogeneously distributed across seasons, in particular, only two scenes were acquired during the winter, precluding seasonal analysis (at least 5 scenes needed). In addition to the different number of scenes captured across seasons, when scenes are aggregated per season, they provide different spatial coverage over the study area. Our approach applies a criterion to require a minimum of 5 valid observations per pixel, which is hard to obtain even with 10 scenes (as in




Figure 3.2: PSGD detected by crosschecking derivative analysis, angular distance, and SST intervals, and with temporal consistency of more than 50% of the valid observations. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents five categories of temporal consistency: > 50% (lighter orange) to > 90% (darker orange). Pixels with less than 5 valid observations are represented in gray. White pixels have at least 5 valid observation but no detectable PSGD. Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.

the autumn and summer seasons). For example, the autumn season has considerably fewer pixels that satisfy the minimum number of valid observations relative to the spring season, which only has 7 scenes. Southern Ireland has frequent rainy days year-round, although the

autumn is often the season with highest amount of rainfall and with the larger number of wet days (more than 1 mm), which could explain the reduced number of valid observations as a consequence of cloud coverage. Additionally, when working with a limited number of scenes, it becomes more likely that the spatial distribution of valid pixels changes across seasons. For example, the observable areas (minimum of 5 valid observations) in the summer and spring seasons are not the same: they overlap only on 27% of the observable summer pixels (Summer has more 25% of observable pixels than spring).

3.3.2 Seasonality of PSGD

The subset of 29 scenes from the Landsat 8 catalog revealed to be insufficient for a complete analysis of PSGD seasonal patterns, due the limited areas where the criterion of 5 minimum number of observations is satisfied. For that reason, no plumes were detected in the autumn and winter seasons (Figure 3.3). In the summer, 5 small plumes (41 pixels total) are detected, all close to each other in an offshore site. This site is located approximately 8 km south of Brownstown Head, a location previously identified as a potential site for SGD based on the analysis of thermal anomalies [92]. Although this PSGD detection could be attributed to a possible SGD occurrence near Brownstown Head, at 50 m b.s.l. it is more likely a false positive [28]. In the spring, all 26 plumes (190 pixels) detected are located along the coastline, covering 8 locations, two of which with plumes of 70% incidence (south of the Ring Peninsula and west of the Carnsore Point). All of these 8 sites were also identified through the non-seasonal analysis, although not necessarily by the same exact pixels; e.g., the plume detected near the Carnsore Point was on the east side of the headland.

As mentioned above, in the summer we have 10 scenes, more than the 7 in the spring, and with a different spatial coverage, which seems to limit the observable area near the southern Irish coastline (note the nearshore gray pixels). However, when ignoring the minimum number of observations criterion, we see PSGD plumes detected in several sites along the coastline across all seasons Figure 3.4. But in this case the detected plumes lack the support otherwise warranted by temporal consistency. Additionally, the combination of all scenes in each season yield thousands of false positive plumes, which tend to be located in the border

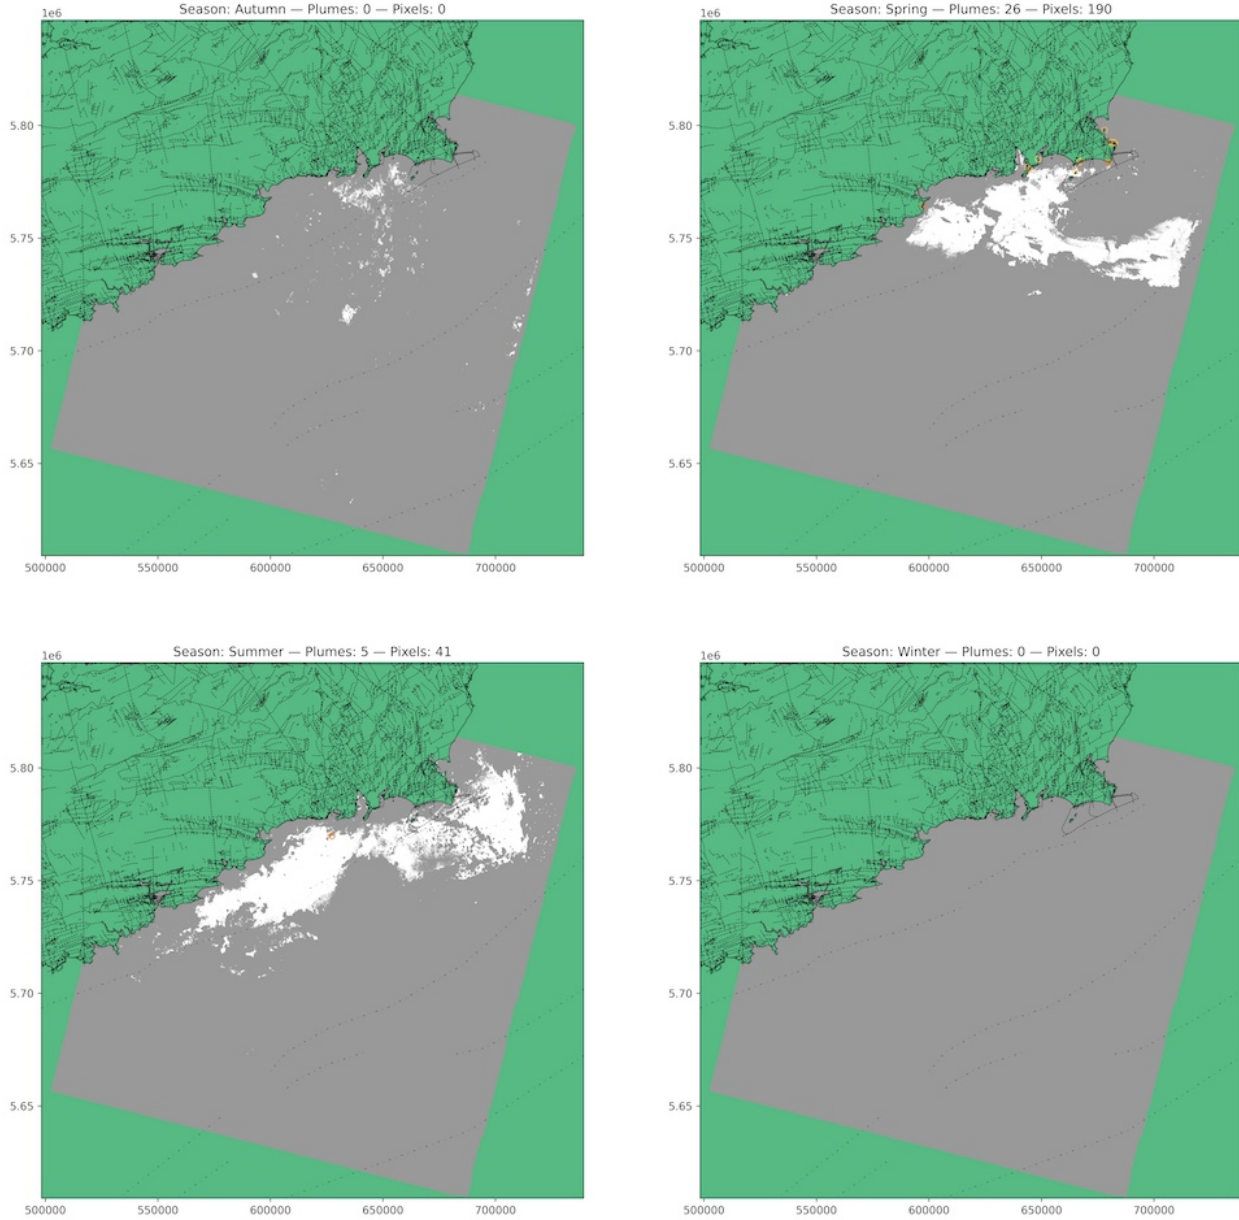


Figure 3.3: Seasonal PSGD detected by crosschecking derivative analysis, angular distance, and SST intervals, and with temporal consistency of more than 50% of the valid observations. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents five categories of temporal consistency: > 50% (lighter orange) to > 90% (darker orange). Pixels with less than 5 valid observations are represented in gray. White pixels have at least 5 valid observation but no detectable PSGD (no PSGD detected in autumn and winter). Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.

of masked out clouds (included in the gray pixels). We hypothesize that this confounding factor is explained by the lower apparent SST due to cloud shadow—mimicking the lowest SST interval, hence we do not see this phenomenon during the winter, when PSGD would

likely be associated to the highest SST interval—and the effect that cloud shadow have on the reflectance, particularly in the near infrared range [57], thus influencing the ocean color approaches—derivative analysis and angular distance, perhaps in a way that approximates the reflectance signal to the target spectrum. This side effect is particularly observed in the autumn and summer, whereas in the spring there are relatively fewer false positives.

Nevertheless, these results can provide some insight on the seasonality of PSGD in the southern Irish coast. Generally, there are more plumes detected per site during the spring and winter seasons, suggesting a higher flow in these seasons. The locations with PSGD plumes detected in all seasons (cyan outline in Figure 3.4) are likely sites with perennial SGD. In those sites, there are less plumes in the summer, and in some cases also less in the autumn, again, possibly due to a decreased flow. In the sites with plumes detected in all seasons except in the summer (blue outline in Figure 3.4), it is not implied that SGD is intermittent, but that, at a minimum, PSGD is harder to detect in the summer, possibly corresponding to similar conditions as the sites with plumes detected in all seasons and with less plumes in the summer. The same applies to the site with plumes only detected in the spring and winter (purple outline in Figure 3.4), SGD could be intermittent there, but it is more likely that it just has a weaker signal, possibly due to a reduced flow. The plumes detected in sites signaled in all seasons except in the winter (orange outline in Figure 3.4) are likely due to the lack of observations (gray pixels in Figure 3.4) in those areas during the winter (which is represented here by only two scenes). Similarly, plumes detected only in the spring and summer (yellow outline in Figure 3.4) also correspond to areas not well observed in the winter, and which could possibly correspond to a reduced discharge in the autumn.

We note one particular observation from the results without the temporal consistency criterion: fewer, small plumes are detected in Waterford Harbour in the summer, comparatively to the other seasons. Note that when our approach results in no or few plumes detected in a certain location and season, it implies PSGD is harder to detect in that location and season, possibly due to a reduced discharge. Recall that SGD has been verified in situ in this site, specifically, during the summer (August 5th and 10th 2010)[92]. Nevertheless, the reduced number of PSGD plumes in the summer could be explained by a decreased discharge

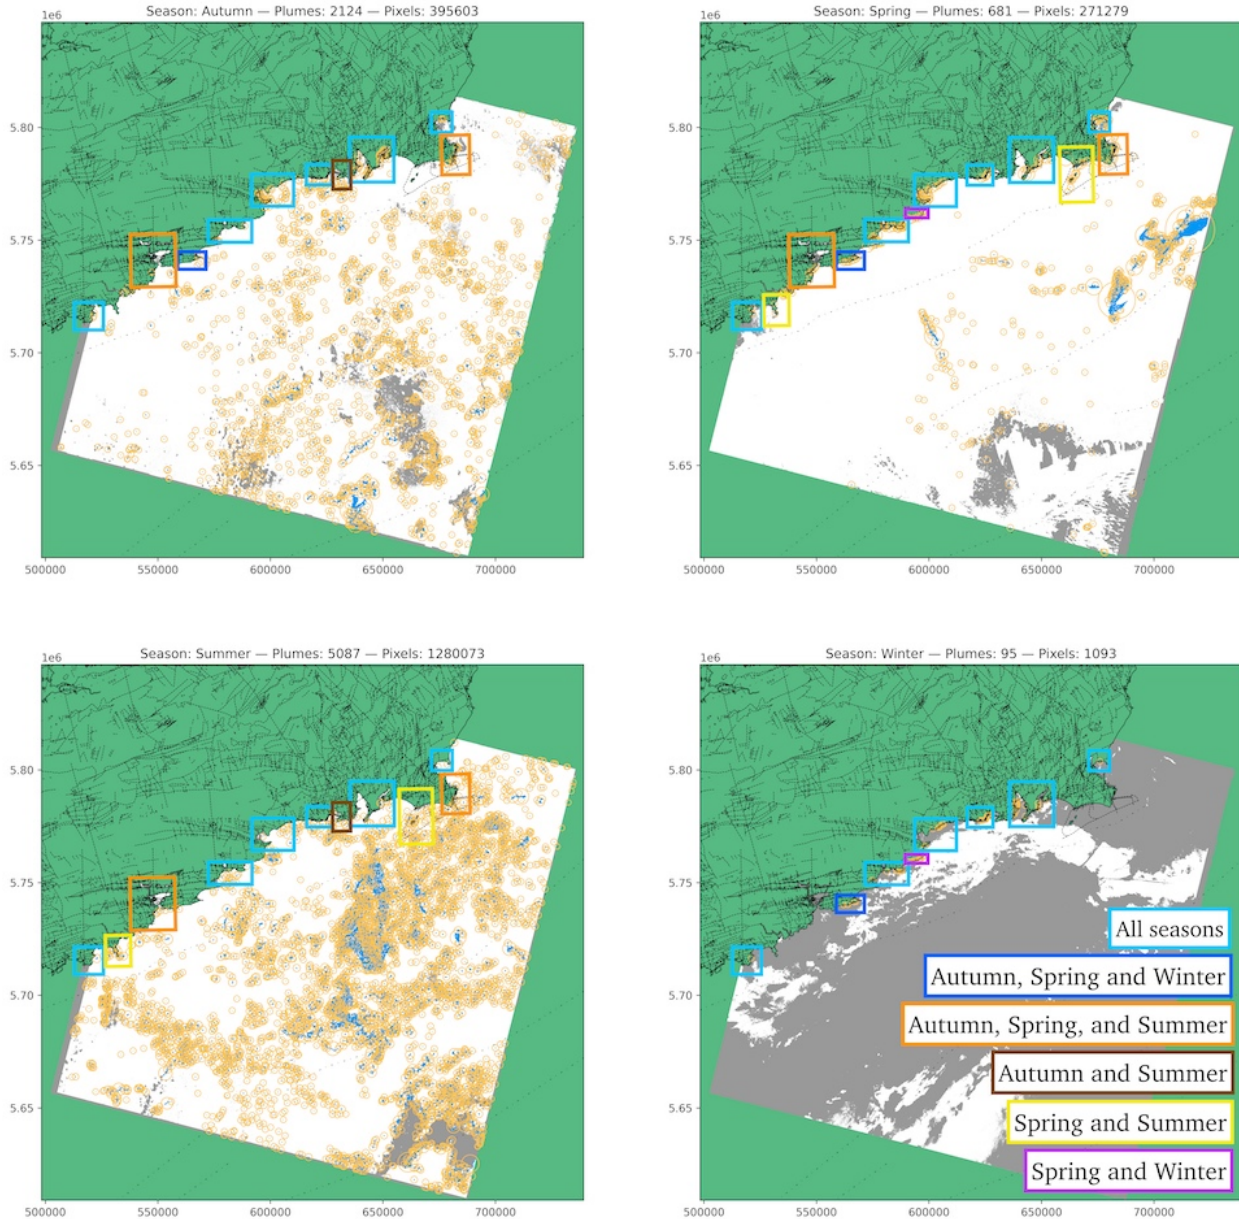


Figure 3.4: Seasonal PSGD detected by crosschecking derivative analysis, angular distance, and SST intervals, but regardless of temporal consistency. PSGD plumes (blue pixels) are encircled (orange) to assist visualization. No data points are represented in gray. White pixels have no detectable PSGD. Colored outlines identify seasonal patterns in PSGD sites. Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.

in this season. Groundwater levels in a nearby coastal aquifer show seasonal fluctuations, as observed through a monitoring well near Waterford Harbour (Figure 3.5). Generally, the groundwater level at this well reaches its yearly maximum in the winter, then decreases during the spring and summer, reaching the yearly minimum in the summer or in the first

half of the autumn. With some working assumptions, we can verify the relation between groundwater levels at the monitoring well and the number of PSGD pixels in Waterford Harbour throughout the study period. Figure 3.5 shows that, with a few exceptions (indicated by purple circles), a decreasing trend in the groundwater level is usually followed by a relatively lower number of PSGD pixels within the Waterford Harbour (orange lines and dots). We consider this a remarkable observation, suggesting that the likelihood of SGD occurrences, or SGD flow in general, may be linked to groundwater levels observed from select wells. In Figure 3.5, the number of PSGD pixels is derived from DA without crosschecking with AD, as DA is a more permissive approach, thus more likely to capture PSGD signature when under low flow conditions—although it also becomes more prone to noise when not crosschecked with AD. This analysis assumes that SGD in Waterford Harbour is driven by hydraulic head, and that the monitoring well is representative of some upper head condition, with a down gradient towards the discharge locations. It also requires the assumption that the non-linear relationship between groundwater storage and discharge borrowed from recession theory holds true [8, 19]. With these, it should be expected SGD to occur—or its flow to increase—during periods of high driving head. It should be noted that a change in sea level—which we considered constant in this simplistic analysis—would change the lower head condition, and therefore it could influence the discharge relation.

More broadly, we consider that there is added value in analyzing the results without the temporal consistency criterion, particularly when dealing with a reduced number of scenes for a seasonal analysis. Where available, data from groundwater monitoring wells can be combined to seek to explain possible differences in PSGD plumes across seasons and possibly establish a predictor for SGD flow. Additionally, our methods also contribute to a broader detection of potential SGD sites SGD. For example, with a single-scene approach, it is more likely to miss potential sites for SGD, such as in cases where SGD flow is variable across seasons. Based on a single-scene approach that uses thermal anomalies, [92] listed four potential sites for SGD within our study area: Old Head of Kinsale, Brownstown Head, Swines Head, and Hook Head. The scene used was acquired in late spring (from a Landsat 7 overpass on 2010-06-02). Near Old Head of Kinsale, we found PSGD in spring and summer (yellow outline in Figure 3.4). Near Brownstown Head and Swines Head, in the autumn and

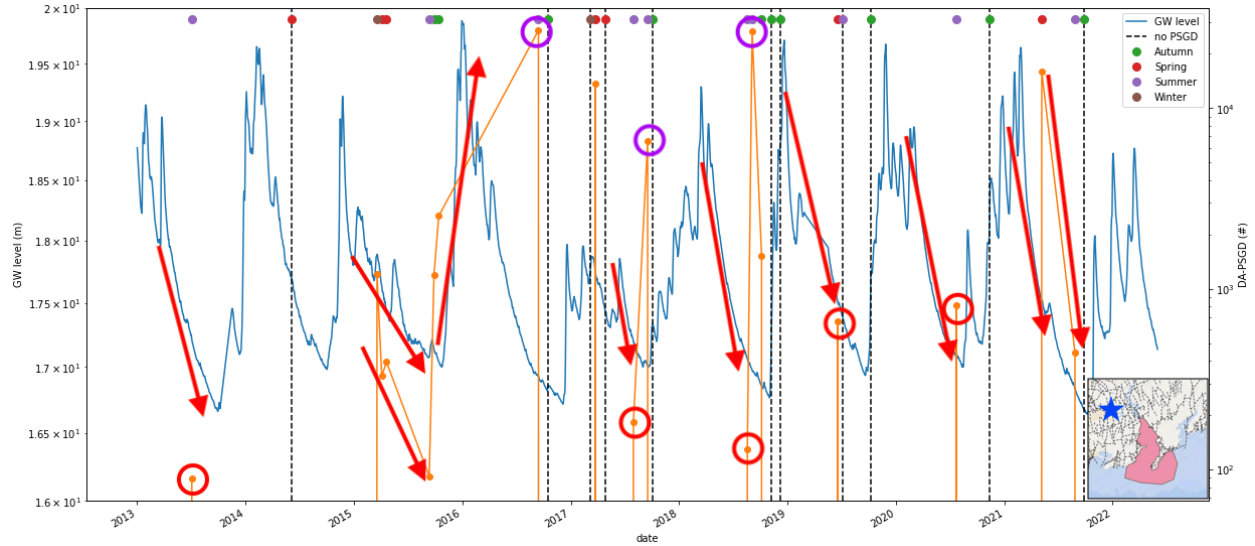


Figure 3.5: Daily records of groundwater level (solid blue line) at one monitoring well (blue star in the inset) near Waterford Harbour (pink shade in the inset) showing seasonal variation. The Slieveroe well (StationID GWIE_SE_G_14915000018) is a 100 m deep borehole situated in Ordovician Volcanics, and is located 8 km–23 km from the Waterford Harbour. The number of pixels flagged by derivative analysis crosschecked with SST anomalies (DA-PSGD, orange solid line and dots) within Waterford Harbour shows similar trends to groundwater levels (indicated by red arrows), apart from a few exceptions (encircled in purple). Colored dots in the top indicate the season of the corresponding Landsat 8 scene (green: autumn, red: spring, purple: summer, brown: winter). Vertical dashed lines indicate scenes without PSGD detected within Waterford Harbour.

summer (brown outline); our study period starts in 2013, being possible that the onset of some hydrological conditions that favored SGD had shifted a few weeks between late spring and early summer. This explanation is supported by the fact these sites had relatively smaller thermal anomalies, e.g., with approximately half the magnitude of the thermal anomalies verified near Hook Head [92], which we found PSGD plumes in all seasons (cyan outline), thus this site also satisfies the temporal consistency criterion (as seen in Figure 3.2). Our results expand the list of PSGD sites, even with the more constraining temporal consistency criterion, therefore providing a more comprehensive and certain—when considering the multi-criteria approach—outlook of PSGD occurrences.

3.3.3 Spatial and Temporal Patterns of SST Anomalies

In our methods, SST is separated in intervals through a non-linear transformation, applied at the scene level. This transformation creates five ordered categories to which pixels

are assigned: 1 to 5, coldest to warmest SST. Because this is performed at the scene level, SST intervals are comparable across seasons. Therefore, we can compare the spatial distribution of the coldest and warmest intervals in each season, regardless of the corresponding absolute SST values. Figure 3.6 shows the pixel-wise mode of the SST interval per season, and we can see how they changed through from one season to another (except to and from the winter, due to low data availability).

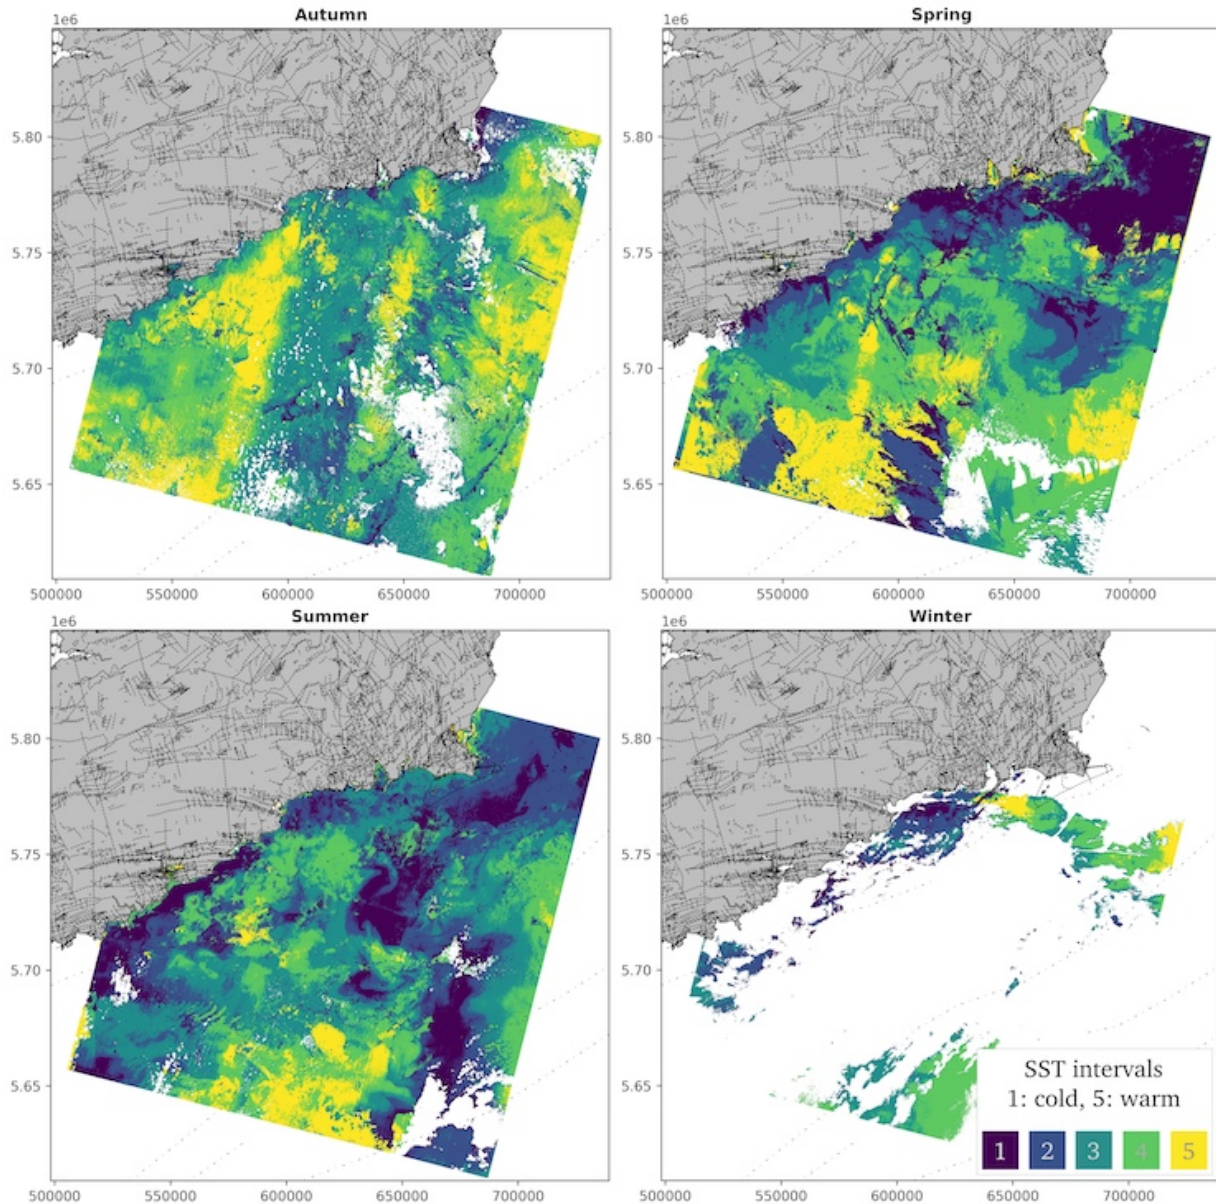


Figure 3.6: Mode of SST intervals from k -means clustering in each season, where interval 1 is the coldest (dark blue) and interval 5 is the warmest (yellow). The pixel-wise mode shows the most commonly assigned interval for each pixel. White pixels represent no data points or masked out clouds. Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.

3.3.3.1 Spring to Summer

From spring to summer, we see a change in the large, cold water mass (lowest SST interval) in the northeast corner of the scene—corresponding to the area where the Irish Sea and the Celtic Sea meet—, where the large, cold plume tends to fall within less cold SST intervals. In the spring, we tend to look for warm plumes in the northeast (Figure 3.1), but not in the summer, showcasing the benefit of determining the target SST interval (coldest or warmest) per pixel and instead of using the same for the entire scene. This is particularly relevant in this case, because the boundary between the Irish Sea and the Celtic Sea is known to promote primary growth[70], thus enabling a confounding factor to SGD signatures. In the summer there are still cold plumes in this region, which could explain the PSGD plumes flagged by DA—crosschecked with SST—far offshore in the northeast (Figure 3.8).

We also note the shift in SST intervals within the Waterford Harbour. In the spring, coastal waters within the embayment tend to be assigned to the two warmest intervals, whereas in the summer there they are also assigned to the colder intervals. Pixels in the warmest interval during the spring are more prone to be conflicted with surface water signatures, because in this area during the spring there is a tendency to look for warm plumes. Note, however, that it is hard to assess this tendency from the mode (Figure 3.6) alone, because there is no minimum number of observations imposed. In fact, most of the areas assigned to the warmest interval—within the Harbour during the spring—are excluded when considering a minimum of 5 valid observations (note the gray pixels in Figure 3.3). This assessment about the warm PSGD plumes in the Waterford Harbour, and in particular near Hook Head, seems to contradict the plumes highlighted in previous studies on a spring (2010-06-02) scene, where PSGD was expected to be linked to cold plumes[92, 12]. We argue these results are not necessarily contradictory, as the scene in those past studies was from a late spring date, where likely PSGD plumes were expected to be cold; likewise, we processed two scenes in late spring, 2020-06-02 and 2019-06-19, where only cold plumes were targeted (see Table B1).

Surface water inputs to the coastal region seem to be captured through a shift in SST intervals across seasons. In the spring, SST from within river mouths and estuaries tend

to fall in the two warmest intervals (4 and 5), but those plumes do not extend much far outside of the corresponding embayments. Then, in summer, while SST in the same areas remains in the warmer intervals, it tends to be less in the warmest interval (5), and more in the two next warmest intervals (3 and 4). Additionally, there seems to be an increased influence of surface waters in SST along the coastline, possibly reflecting a seasonal change in surface water discharge. These changes seem to be more easily seen in the western half of the coastline (Kinsale Harbour, Cork Harbour and Youghal Harbour), where we note a striking separation between a surface water-influenced zone, and a colder (SST interval 1) plume parallel to the coastline. This example showcases the importance of SST to separate SGD from surface waters.

3.3.3.2 Summer to Autumn

From summer to autumn, the cold large plume from the Irish Sea is further diminished, now present only near the Wexford Harbour. Overall, the Irish Sea and the Celtic Sea warm up during the summer and that is reflected in the trend in early autumn. Recall that the majority of the autumn scenes were captured earlier in the season, so the separation from the summer is more nuanced. Accordingly, autumn—at least the first half—could be a prime season to detect SGD, at least in terms of the reduced confounding factors with cold water, although, this season might as well be characterized by lower, harder to detect SGD flows (see Figure 3.5 and the discussion in Section 3.3.2). We note the presence of relatively cold plumes emanating from Hood Head, as well as from within the Waterford Harbour, Brownstown Head and Swines Head. However, they may also fall within the second coldest interval, thus being rejected as PSGD plumes in our current parameters (only pixels in the coldest/warmest SST interval gets shortlisted as PSGD). Again, this could be explained by a decreased SGD flow, hence the attenuated thermal influence of SGD in coastal waters, causing a weaker signature. In general, this could be true for the entire study area, hence we have considerably fewer or none PSGD plumes detected in autumn. This observation suggests that the decision of using one or two (coldest/warmest) SST intervals could depend on season, even within the same study area (scene). In fact, it could be

incorporated in the analysis of the expected SGD thermal signature, and use the comparison between groundwater temperature and SST_{MODIS} to determine the number of SST intervals to consider, thus better exploring the availability—and perhaps the number—of the multiple SST intervals. On the other hand, in autumn, the long, cold plume along the western half of the coastline is still visible but less pronounced (divided in SST intervals 3 and 4), and surface water pixels tend to be a colder SST interval than in summer, which would indicate that it is best to continue to limit PSGD detection to the coldest SST interval, to avoid having surface waters as a confounding factor. Recall that these remarks (mode of the SST interval) are not taking into account the number of valid observations, so there is a bias here in possibly favoring some specific scenes. Nevertheless, we think there is an opportunity for improvement in PSGD detection by better determining the SST intervals classified as anomalies.

3.3.4 Multi-criteria Approach

The combination of different detection criteria is one of the features of this approach, providing independent, parallel lines of evidence to support detected PSGD occurrences[12]. Previously, we have applied three detection criteria based on spectral signatures as observed from remote sensing data: SST intervals (thermal signature), derivative analysis and angular distance (color signature). Here, we add the temporal consistency criterion—observing all the other three criteria in the same pixel consistently over time—and we extend our previous analysis of the effects of crosschecking the different criteria.

With relatively few observations in the time series (see the distribution of the number of valid observations per season in Figures S1 and S2), crosschecking DA and AD is an important step to reduce noise and better assert PSGD consistency over time. Results from DA (crosschecked with SST intervals) would yield 25 times more plumes (and 47 times more pixels, thus larger plumes) than when crosschecking with AD. Also from DA alone, additional coastal locations would be highlighted as prime sites for PSGD (e.g., Old Head of Kinsale, Kinsale Harbour, Ballycotton Bay) and more plumes would add support to the previously listed locations, however, that would also come with the cost of more noise (several plumes

detected offshore). AD is the most stringent criterion, which tends to produce fewer and smaller plumes than DA, even with an increased percentile threshold. When this threshold is set to the 10th percentile, it would result in 35% more detected plumes (171 plumes, 663 pixels) after crosschecking with DA, however, a more lax percentile threshold does not necessarily translate in an increase in PSGD sites (only 1 plume outside of the previously listed sites), instead, it would mostly increase plumes count and size in the existing sites, as well as an increase in noise (additional likely false positives located far offshore). For a multi-scene application, we consider the 5% threshold to offer a better trade off when considering the increase in uncertainty associated with a more lenient threshold. Recall that plumes from both DA and AD are first crosschecked with SST anomalies—their corresponding SST has to belong to either the coldest or warmest SST interval, depending on season—and only then DA and AD plumes are crosschecked. Crosschecking relies on the assumption that PSGD is more likely to occur where plumes from each criteria intersect. Generally, when looking in terms of the number of flagged pixels, AD is the most constraining criterion, followed closely by SST (Figure 3.7).

The three detection criteria follow similar trends through the study period in terms of the number of flagged pixels, which tends to increase in the spring and summer, and decrease in the autumn and winter. This is particularly true for DA and AD, which have an almost perfect linear correlation ($r \approx 1.0$), while both have a weaker linear correlation with SST ($r = 0.48$ for both). The seasonal trends could be explained by different light conditions across seasons, since primary production—a proxy we explore to link optical RS to PSGD—decreases with increased light attenuation[41], an effect on primary production whose magnitude depends on latitude. At south Ireland latitude, primary production is generally the lowest in the winter, then autumn, spring, and highest in the summer[82]. However, from our results, it is not straightforward to link light conditions to the seasonal variation in PSGD detection, due the reduced, unequal number of scenes in the winter and the corresponding varying cloud coverage across seasons.

While SST and AD tend to constrain resulting pixels in a similar manner, the spatial and temporal patterns of SST are crucial for an effective detection; results directly from DA or AD, without crosschecking with SST, are harder to interpret and less reliable. Plumes




Figure 3.7: Number of PSGD pixels per scene resulting from derivative analysis (DA, blue line), angular distance (AD, orange line), and SST intervals (SST, green line), prior to any crosschecking. Colored dots in the top indicate the season of the corresponding Landsat 8 scene (green: autumn, red: spring, purple: summer, brown: winter).

from only the DA criterion tend to be relatively larger—when compared to DA combined with the other criteria—and noisy, with several plumes located far offshore, which can cover a significant portion of the scene. We observed this limitation of DA in our previous study, which was based on single-scene analysis. Accordingly, it is more likely that such numerous and large plumes would have temporal consistency, particularly on a limited number of scenes: DA itself results in hundreds of thousands of plumes (505855 plumes, 1495945 pixels; see resulting map in Figure B16) that are detected in more than half of the observations. Yet the power of this criterion is still visible: out of those plumes, only 0.14% (7244 plumes, 18253 pixels) have a temporal consistency of more than 90%, and those plumes are mostly located along the coastline (see Figure B17). The same is verified in the seasonal analysis (0.19% in spring and 0.10% in summer; see Figure B21). Finally, when crosschecking with the SST criterion, the number of plumes that satisfy the temporal consistency is further reduced, even when applying plume buffering (3227 plumes, 21497 pixels; see Figure 3.8).

The more stringent AD criterion, without crosschecking with SST, also produces plumes along the coastline, with considerably fewer likely false positives offshore (see Figure B30) (2753 plumes, 4118 pixels). The same is verified in the seasonal analysis, including a few

plumes in the Autumn (27), a season when DA failed to detect nearshore plumes (see Figure B33). However, without plume buffering, AD cross checked with SST becomes excessively constraining, resulting in less and smaller plumes, which fail to show any temporal consistency (only 2 plumes, 3 pixels). With buffering, but without crosschecking with SST, several plumes are identified along the coastline, with also many located offshore (36655 plumes, 367740 pixels; see Figure B31), although with relatively lower temporal consistency, and which are again reduced when crosschecking with SST (264 plumes, 978 pixels; see Figure 3.9).

The temporal consistency criterion can also be applied to SST intervals, producing PSGD plumes only derived from the thermal signal. Each pixel whose SST interval matches the target interval—determined according to the comparison against groundwater temperature (Figure 3.1)—is then shortlisted as PSGD. Out of the three spectral criteria, SST is the one that by itself results in the most constrained set of temporally consistent plumes (963 plumes, 7593 pixels; see Figure B48), even less than AD, but with plumes on average 5 times larger (hence, more pixels than AD). Plumes derived only from SST highlight several nearshore sites, without a significant number of offshore false positives, where the majority of those seem to be due to the confounding factor from the Irish Sea mentioned earlier (Section 3.3.3.1). When plumes are buffered, more areas of the coastline are covered with shortlisted plumes (11049 plumes, 130203 pixels; see Figure 3.10), but not all are included in the final set of PSGD plumes (Figure 3.2).

Combining the three criteria improves the reliability of our results, although it is also useful to assess each criterion individually. AD seems to be most constraining criterion when considering plume buffering, which perhaps could be attenuated with a larger buffering specifically for AD (we applied the same buffer radius for all criteria). For example, in the Kingsale Harbour and Cork Harbour areas, all three criteria shortlisted several plumes, but those sites are not present when AD is crosschecked with SST. That could be explained by a temporal mismatch (i.e., AD plumes are not identified in the same scene as SST plumes) or by a spatial mismatch (i.e., AD and SST plumes may be close but do not overlap), in which case a larger buffer could increase the sensitivity to a weaker PSGD signature—although with the cost of added noise and uncertainty. There is certainly a trade-off between detecting

weaker PSGD signatures and an increase in false positives (noise and confounding factors), but in certain cases that could be the only way to detect a weaker signal. As illustrated above with the case of the Waterford Harbour (Section 3.3.2), the more lenient DA criterion can be useful to detect weaker signals.




Figure 3.8: PSGD detected by derivative analysis crosschecked with SST intervals and with temporal consistency of more than 50% of the valid observations. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents five categories of temporal consistency: > 50% (lighter orange) to > 90% (darker orange). Pixels with less than 5 valid observations are represented in gray.




Figure 3.9: PSGD detected by angular distance (5th percentile threshold) crosschecked with SST intervals and with temporal consistency of more than 50% of the valid observations. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents five categories of temporal consistency: $> 50\%$ (lighter orange) to $> 90\%$ (darker orange). Pixels with less than 5 valid observations are represented in gray.

3.3.5 Methods Limitations

In this study, we improve our previous methods for SGD detection by exploring the temporal consistency of plumes. A direct limitation of this approach is the ability to only

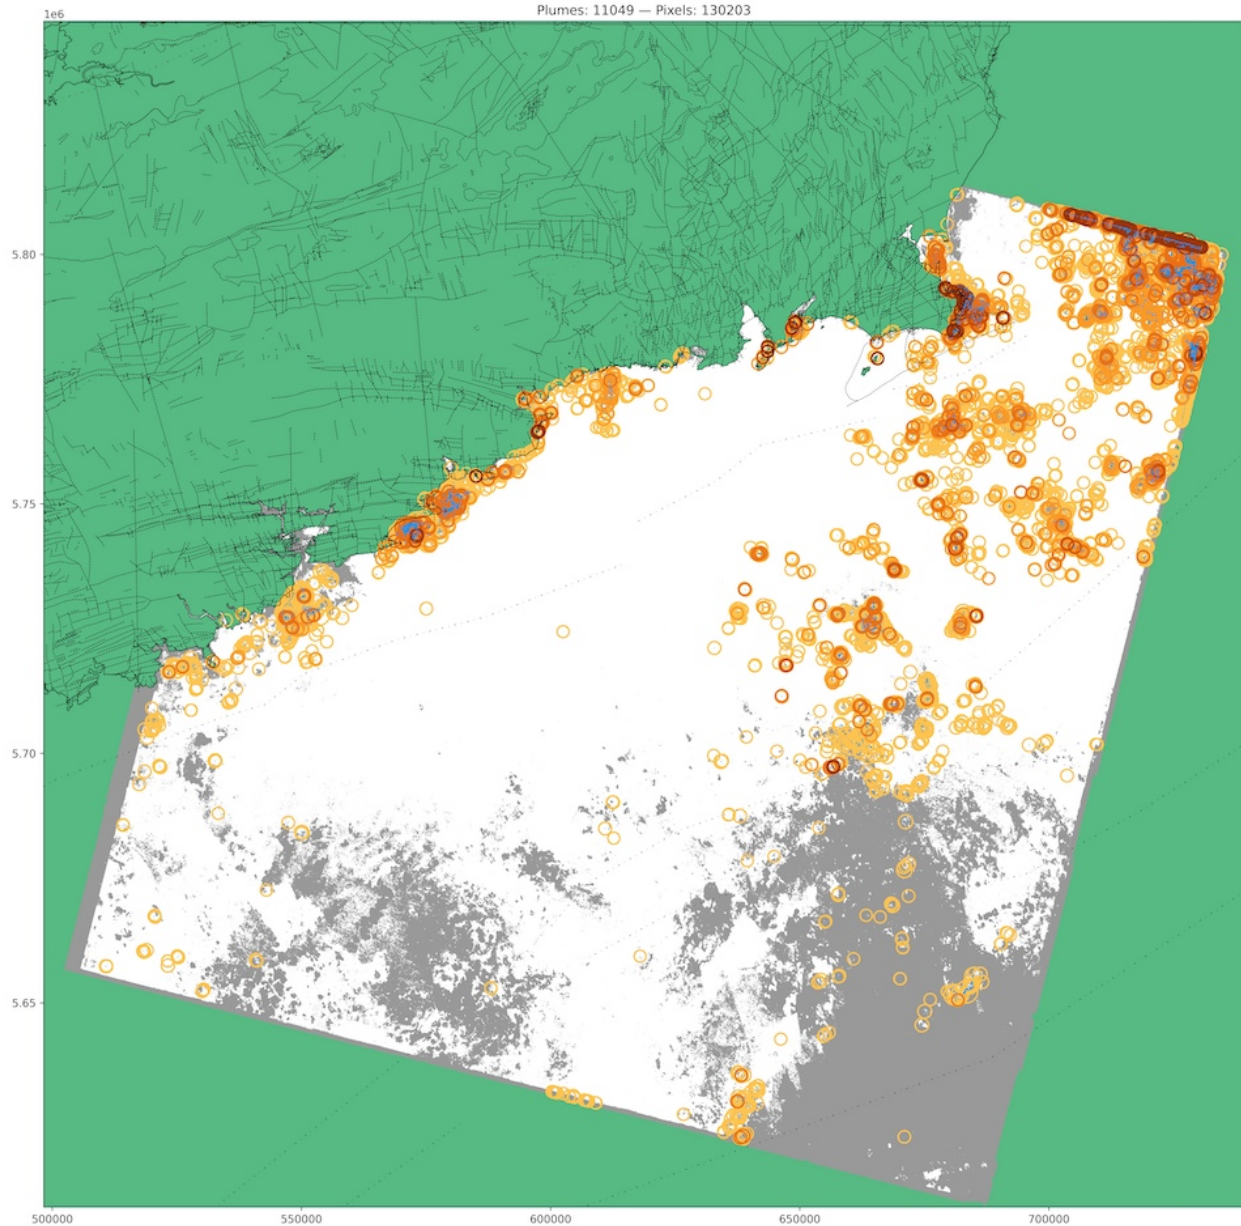


Figure 3.10: PSGD detected by SST intervals and with temporal consistency of more than 50% of the valid observations. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents five categories of temporal consistency: > 50% (lighter orange) to > 90% (darker orange). Pixels with less than 5 valid observations are represented in gray.

capture perennial occurrences. We attempt at capturing intermittent PSGD by aggregating observations by season, but that comes with the challenge of a reduced number of valid observations, particularly in autumn and winter, when cloud coverage is typically higher. Additionally, a more refined temporal analysis would be required to capture PSGD that

may occur more sporadically, or on a temporal cycle out of phase or larger than season.

In addition to the limitations inherited from our previous study (e.g., using a single spectrum for the expected spectral signature), the other limitations here are tied to working assumptions and simplifications: the scene selection criteria highly reduced the number of scenes analyzed, limiting the scope of the seasonal assessment; the dependency of groundwater records further limits the number of usable scenes, such that an alternative approach or data set could be used—e.g, infer groundwater temperature based on land surface temperature as in [7], thus minimizing the dependence on the more spatially and temporally limited in situ data; the nearest neighbor interpolation to fill the gaps induced by MODIS’ coarse resolution might be introducing artifacts and bias determining SST anomalies, as it is possible that some coastal pixel and the nearest MODIS pixel have significantly different SST, whereby the expected SGD thermal signature could differ; the expected thermal signature is a function of groundwater temperature from measurements in several wells, ignoring groundwater flow direction, catchments delineation, and the likely discharge location, thus possibly missing localized differences along the coastline. Despite these limitation, we believe that the work needed to address the limitations is clear and unlikely to pose major scientific obstacles, and that our results still provide an important improvement over the existing methods and a meaningful extension to the knowledge of SGD in south Ireland.

3.4 Conclusions

In this study, we expand our previous work on using optical and thermal RS data to locate PSGD occurrences. We show strengthened support to the identified PSGD sites by combining remote sensing imagery acquired across multiple years, thus showcasing the applicability of the long, global Landsat catalog to the global mapping and monitoring of PSGD. We demonstrate that time persistent signatures of PSGD provide stronger evidence for the prime sites for SGD occurrence. We also analyze seasonal variability of PSGD, suggesting connections between groundwater level and the effectiveness of SGD signatures, likely linked to seasonal varying SGD flows.

4.0 Mapping Submarine Groundwater Discharges Through Remote Sensing and Machine Learning

4.1 Introduction

An accurate budget for the global water cycle requires next-generation observation systems that leverage the increasing spatial and temporal resolutions of satellite platforms [62]. Submarine groundwater discharge (SGD) is an important component of the water cycle [50, 64], which is inherently challenging to account for and incorporate into conceptualizations of the groundwater-oceans interface, thus it is commonly neglected in those assessments [61, 85]. Notwithstanding, SGD flows can be of the same magnitude as large rivers (e.g., 10^7 $\text{m}^3 \text{d}^{-1}$) [11, 87], and they are estimated to represent 6% of the total discharge of freshwater into the oceans [94, 10]. With a systematic, globally applicable framework for the identification of locations where SGD occurs, we can constrain the role of SGD in the water cycle and clarify global water dynamics.

Remote sensing data have been used to identify potential SGD (PSGD) occurrences [92, 28], however, only signatures in the thermal infrared (TIR) wavelength region were linked to SGD. Groundwater temperature is a well-known geophysical tracer for SGD [9], therefore TIR data can be utilized to assess sea surface temperature (SST) wherever there are thermal gradients between groundwater and seawater [27]. However, there are several confounding factors that can cause changes in SST, such as low tides, upwelling events, surface water and anthropogenic inputs [28], implying that SGD assessments based solely on SST observations could result in significant uncertainty. Tackling this limitation was the main goal of our previous work, where we introduced the combination of ocean color (OC)—as observed by remote sensing of the visible and near infrared (VNIR) spectrum—with SST to improve the identification of PSGD [12]. SGD is one of the main transport pathways of nutrient and dissolved terrestrial materials to the sea [79], which can boost primary productivity [72] and increase chlorophyll content [81, 72]. Therefore, SGD induces signatures in OC, which can be captured in VNIR surface reflectance [88].

This study is based on the established link between SGD and the combined SST and OC (Chapters 2 and 3). Here, we propose a different approach to retrieve RS signatures and identify PSGD. We utilize a powerful and widely applied machine learning technique—convolutional neural network (CNN)—applied to Landsat 8 data to detect PSGD. We use spectra from each pixel—which can be seen as 1-D variables with 5 points, considering the first five bands in the VNIR range—as inputs on a CNN for classification. The output is tailored to produce a map of the study area indicating SGD or non-SGD (binary). The broad end product is a classification framework that digests a RS image and outputs the location of probable SGD occurrences. Additionally, that output can be combined in a similar application as in Chapter 3, i.e., we can apply the temporal consistency to plumes detected through this machine learning approach.

The potential for deep learning in the geosciences has been increasingly recognized, including for remote sensing applications [96, 39]. Optical and thermal signatures observed from remote sensing data have been leveraged to detect recent lava flows through a neural network [15]. Although lava flows and SGD are distinct natural phenomena, they both can be linked to changes in surface temperature as well as in surface reflectance. Therefore, the applicability of this type of machine learning to a task similar to SGD detection has been demonstrated. Certainly, the two phenomena have many differences and each case comes with different challenges. Here, we propose a new approach for SGD detection, in a novel application of neural networks and widely available remote sensing data, scalable for global mapping of potential sites of SGD.

4.2 Materials and Methods

4.2.1 Case Study

Similar to work presented in the first two core chapters, the selected location to demonstrate the PSGD detection method presented in this work is within the southern Irish coast. Specifically, we now focus on the Waterford Harbour region (Figure 4.1), a prime site for

SGD that has been ground-truthed [92].

In this region, SGD occurrence was confirmed through the collection of 55 water samples measuring a radon isotope activity [92], namely around Hook Head, where thermal anomalies had initially suggested the presence of SGD. In our work, we utilize the location of those samples as an indicator of true SGD in manual classification of plumes. Additionally, two other locations within this region have been suggested as potential SGD sites through the analysis of thermal anomalies [92]: Brownstown Head and Swines Head.

In our previous work, we have shown links between the location of radon isotopes and mapped PSGD plumes based on the combination of color and SST data (Chapter 2). Those links were evaluated over time and provided insights about the seasonality of PSGD within the area (Chapter 3). Here, we identify PSGD plumes through a different approach to locate the optical and thermal signatures, which corroborate with previously identified potential discharge sites.




Figure 4.1: Landsat 8 surface reflectance in the study area as observed in 2019-07-05, represented in pseudo-color—composite using bands 4 (red), 3 (green), and 1 (coastal blue)—with brightness and contrast adjusted to enhance features. Red polygons outline manually selected PSGD plumes. Inset shows the location of Waterford Harbour in Ireland. Dashed lines show bedrock faults [23].

4.2.2 Sea Surface Temperature and Color

This case uses data from one Landsat 8 (L8) overpass over the study area (path 206, row 24) in 2019-07-05. Color (bands 1–5) and sea surface temperature (band 10) data were

obtained as surface reflectance and surface temperature, respectively, both from a Level-2 product of the Landsat Collection 2 catalog (LC08_L2SP_206024_20190705_20200827_02_T1) sourced from USGS.

4.2.3 Data Preparation

The input image used here was manipulated to adapt it for utilization in the machine learning methods applied here. First, pixels corresponding to PSGD locations are manually classified, forming a labeled data set. Second, the classified data set is divided in several smaller images, which form the training and test data set.

The labeling process leverages the spatial coherence between cold plumes—represented as SST intervals resulting from applying k-means clustering in the L8-derived SST (see Chapter 2)—and in situ measurements of radon activity (55 samples in Waterford Harbour) to manually retrieve spectra from pixels that are believed to be groundwater discharge. The key assumption is that the manually retrieved cold plumes—and their spectra—are ground-truthed SGD, even if they are only occurring in the general vicinity of elevated radon activities, as per samples retrieved several years earlier (2010-08). A second plume was included, despite being in a location not covered by the in situ sampling (Brownstown Head), but which has been classified as a potential site for SGD [92], and in which a cold plume was present. Accordingly, every pixel within the manually outlined plumes (Figure 4.2) are labeled as SGD, and every other pixel is labeled as non-SGD. The set of pixels within the outlined plumes define the target spectra.

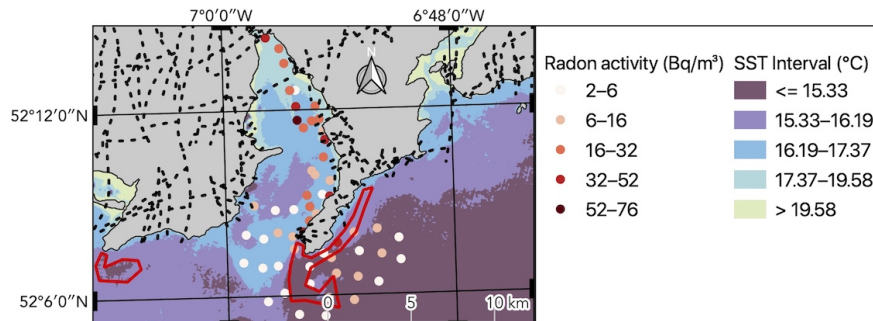


Figure 4.2: Landsat 8 surface temperature in the study area as observed in 2019-07-05, represented in SST intervals ($^{\circ}\text{C}$) derived from k-means clustering. Circles show the location and magnitude of 55 samples of radon activity (Bq m^{-3}), collected in the 5th and 11th of 2010-08 [92]. Red polygons outline manually selected PSGD plumes. Dashed lines show bedrock faults [23].

4.2.4 Training Samples Generation

The manually classified pixels are the basis for the generation of the training samples. A 7x7 moving window, with stride 5x5, traverses the L8 scene, creating one sample at each step (total 21240 samples). The center pixel of the each 7x7 window defines the label assigned to the corresponding generated sample, i.e., if the center pixel is within the previously classified plumes, then the sample takes the SGD label, otherwise, it takes the non-SGD label. This processing step was performed using a QGIS plugin specific for this purpose (<https://github.com/PratyushTripathy/QGIS-Plugin-Produce-Training-Samples-For-Deep-Learning>).

Out of the generated training samples, only 2.3% (494) were of the SGD class. Therefore, it is recommended that the number of samples in each class be balanced. For that purpose, we downsampled—randomly resampled without reposition—the non-SGD samples to match the size of the SGD subset. The balanced data set is then divided in three subsets: training (70% of the samples), validation (15%), and test (15%).

The preparation of training samples includes a normalization step. Although not a requirement, it improves the learning rate of neural networks. A z-scores normalization (mean removal and variance scaling), computed feature-wise, is typically the recommended method [13]. Test data are not included in the normalization step, but are normalized using the same parameters as the training set.

4.2.5 Convolutional Neural Network

CNN is a deep learning technique that has a large learning capacity, which can be controlled by varying the depth and breadth of the network. CNN's make strong and mostly correct assumptions about the nature of *images* (namely, stationarity of statistics and locality of pixel dependencies) [34]. CNN takes an image and applies a function to it to transform it to class scores (considering the problem of classification). The function consists of several rounds of convolutions where the filter entries are parameters, and a few matrix multiplications, where the elements of the matrices are parameters. A typical CNN might have 100 million parameters. CNNs usually have a standard structure: stacked convolutional layers are followed by one or more fully-connected layers.

Here, a simple CNN model is implemented in Tensorflow [48], with an architecture consisting of only two convolutional layers (see Figure 4.3). This simplified model establishes a foundation for further iterations, as well as providing an early assessment for proof of concept.

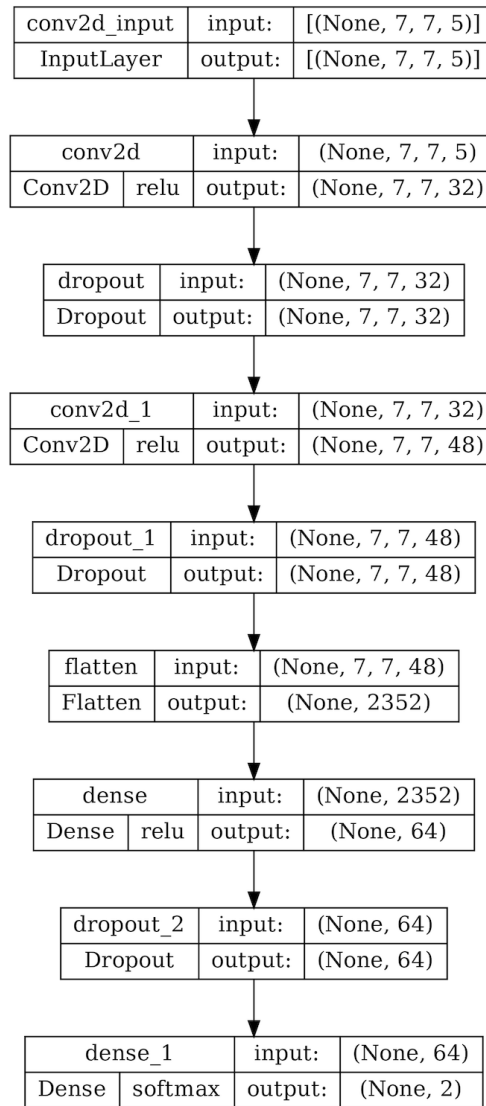


Figure 4.3: Architecture of the implemented convolutional neural network. Left column shows the layer type (top row) and the layer name and the corresponding activations, when present (bottom row). Right column show the input (top row) and output (bottom row) shapes. The first dimension represents the number of samples, where **None** means it accepts any number.

The CNN is used twice. First, the network is trained on the training set. The labels (PSGD/non-PSGD) are provided, as well as the validation set, which is used to assess the efficiency of the learning through each epoch and to minimize a loss function (here, softmax loss). One epoch corresponds to running all of the training spectra once through

the network. After running for a given number of epochs, the best model of the network, that was generated in the end of one of the epochs, is selected as the one that represents the trained network. Here, we experimented running during 30 epochs, however, the model quickly overfitted and we proceeded with the result of the first epoch.

In the second running of the CNN, the test spectra are passed through the network (here the notion of epoch is not used, because there is only one pass). The CNN works as a features extractor, so that by the end of this process we have a set of features that represent each spectrum. The last layer of the CNN is configured to output values according to two classes, like the input: PGDG and non-PSGD. These features in the last layer can be directly related to the probability that the spectrum belongs to each class, so it can also act as a classifier, where the highest value of the three values assigns a class to each spectrum.

4.3 Results and Discussion

4.3.1 CNN Classification

The test set was evaluated against the trained CNN, yielding the confusion matrix shown below (Table 4.1) and the following accuracy metrics: P-Score: 0.632, R-Score: 1.000, F1-Score: 0.775.

Table 4.1: Confusion matrix.

True \ Predicted class	non-SGD	SGD
non-SGD	66	85
SGD	0	146

In the confusion matrix, by definition, entry (i, j) is the number of observations actually in group (i), but predicted to be in group (j). In other words, each row corresponds to the true class/label, and each column to the predicted class/label. The 146 points belonging to

the SGD class are correctly identified as such; that is why recall (R-Score) is 1 (see below). However, 85 samples were inaccurately identified as SGD, hence precision (P-Score) is lower.

The precision (P-Score) is defined as

$$\text{Precision} = \frac{\text{TruePositives}}{\text{PredictedPositives}} = \frac{\text{TruePositives}}{\text{TruePositives} + \text{FalsePositives}} \quad (1)$$

which in our case translates to

$$\text{P - Score} = \frac{146}{146 + 85} = 0.632 \quad (2)$$

Achieving a higher precision is desirable, as that means more true positives, and that when a pixel is classified as PSGD, that is more likely to be true.

Recall (R-Score) is defined as

$$\text{Recall} = \frac{\text{TruePositives}}{\text{ActualPositives}} = \frac{\text{TruePositives}}{\text{TruePositives} + \text{FalseNegatives}} \quad (3)$$

hence, we have

$$\text{R - Score} = \frac{146}{146 + 0} = 1.0 \quad (4)$$

Recall indicates what fraction of all pixels that are originally PSGD are detected as PSGD, i.e., a high recall means when a pixel is actually PSGD it gets correctly classified.

Instead of only relying on the overall accuracy, precision and recall warrant a better sense of how an algorithm is actually doing, especially in the presence of a highly skewed data set (not the case here, because we balanced the data set). For example, if we predict non-SGD all the time and get 99.5% accuracy, then precision and recall will both be 0, because there would be no true positives, implying the classifier is not a good classifier. When precision and recall are both high, that is an indication that the algorithm is doing very well. That would be an ideal case, which is not observed here: our recall is high but precision is inferior. If we would increase the threshold for detection (e.g., $p > 0.7$), we would achieve higher precision and lower recall. On the other hand, if we decrease the threshold for detection (e.g., $p < 0.3$), we achieve lower precision and higher recall. The F1-Score can help assess whether a higher

precision or higher recall is better. F1-Score is a weighted average of precision and recall, defined as

$$\text{F1 - Score} = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (5)$$

therefore, we have

$$\text{F1 - Score} = \frac{2 \times 0.632 \times 1.0}{0.632 + 1.0} = 0.775 \quad (6)$$

All true PSGD was classified as such (high recall, 100%), however, it included noisy, wrongly classified pixels as PSGd (lower precision, 63.2%). In a refined experiment, it should be possible to achieve a higher F1-Score at the expense of a lower recall.

One important caveat should be noted. Given the simplistic approach to collect the training samples, where SGD pixels were manually selected and everything was considered non-SGD, we introduced the possibility of false negatives directly in the training set. Accordingly, it is possible that other SGD-affected pixels are present in the image. Therefore, some of the 85 false positives could, in fact, be true positives. This implies that the calculated precision should be interpreted as a minimum, e.g., it is at least 63.2%.

4.3.2 PSGD in Waterford Harbour

These initial, first-iteration results establish a basis for refinement. We show the result of a very simple CNN with only two convolutional layers and two classes (SGD, no SGD). The CNN input was defined as only two plumes manually selected over the VNIR range (5 bands). Nevertheless, we consider these initial results solid to warrant further development of the method, with several, well defined opportunities for improvement.

Figure 4.4 shows the distribution of the probabilities of belonging to the SGD class. Figure 4.5 shows the pixels where the probability of belonging to the SGD class is above 64%. As seen in the histogram, the number of PSGD pixels falls out quickly after 64%.

Compared to previous findings in this region [92, 12], which have identified a major SGD discharge point at the Hook Head, we note that the general location of the main plume off the Hook is there (at $p > 64.5\%$, but vanishes partially at $p > 65\%$). In part, this was expected, since that plume was manually classified for input, however it was not given in full for the CNN to learn, since part of it was extracted as a test set.




Figure 4.4: Distribution of the probabilities of each pixel being PSGD (as given by CNN classification).

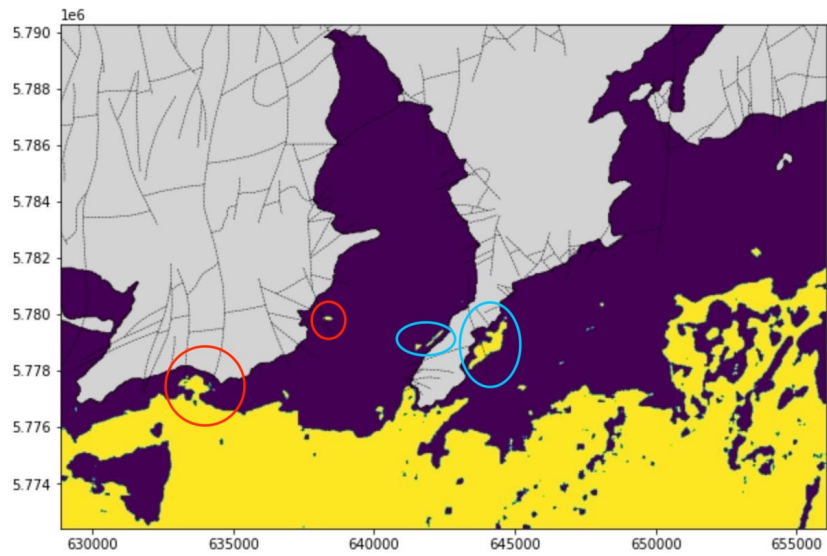


Figure 4.5: Pixels classified as PSGD (yellow) by the CNN ($p > 64\%$). Grey pixels represent land and blue pixels were classified as non-SGD.

We highlight 4 other plumes that stand out from the large noisy offshore plume (all at $p > 64.5\%$ only). The first two (circled in cyan) match with previously sampled areas of high radon activity [92]. They were not part of the manually classified input, despite having a thermal anomaly (not as evident/strong as off the Hook Head). We find it remarkable because the classifier picked up a difference between those plumes and surrounding waters. In particular, the one west of the Hook, despite the difference in SST with surrounding waters, the difference in reflectance is not as evident. The other two plumes (circled in red) are in areas that were not sampled for radon, but that were also identified previously as PSGD through a different approach [12]—the first two plumes (cyan) were also identified as PSGD there. The smaller plume (out of those encircled in red), within the Waterford Harbour, is close to the convergence of a few bedrock fractures, which could explain the occurrence of SGD in that location. As in results from the previous approach [12], it stands out isolated from surrounding waters. Although here with the particularity that the thermal signature was not used to detect the plume. The other one (encircled in red), off Swines Head, besides matching our previous results [12], it is in a location also previously identified as a potential site for SGD based on thermal anomalies [92]. It is connected to the large offshore plume (connected component), but almost fully surrounded by coastal waters with a different spectral signature. Brownstown Head, also a previously identified potential site, does not show a plume as distinct and close to it, but it is possible some of the pixels in the large—and likely noisy—plume could be attributed to SGD originating from this location.

This method improves upon our previous single-snapshot-in-time approach (Chapter 2) as it is less dependent on the thermal signature—while still utilizing it—, therefore more transferable to images acquired by sensors without a thermal band, which typically have a higher spatial resolution (e.g., Sentinel-2, but also platforms based on cubesats such as from Planet). The thermal band is still needed to provide the input, which could be based on a joint overpass, but then, after training, it could be transferable to other scenes without the overpass. The reduced dependency on the thermal signature also makes this approach suitable for cases—locations or seasons—where the thermal gradient between groundwater and seawater is too low for a reliable detection (example in Chapter 3). Additionally, this approach is an alternative to using just one target spectrum (as in Chapters 2 and 3), which

is particular to a limited and controlled environment. Therefore, this approach has the potential to be more robust to variations in local and environmental conditions.

When compared to approaches only based on thermal anomalies [92], the advantage of this approach is the same as the one presented in Chapter 2: combining color and SST improves our ability to detect SGD—i.e., we can detect SGD signatures that may be not as evident solely based on SST, for example, the cyan circles that match with high radon activity but do not have a strong thermal anomaly—and gives more confidence on plumes that are defined only based on thermal (e.g., Swines Head).

4.3.3 Features importance

In the context of machine learning, sensitivity analysis of the inputs is important step for evaluating the impact of each input feature in the model output. In machine learning applications, this sensitivity analysis is known as features importance. Computing features importance is fundamental for the interpretability of machine learning models. Interpretability is a measure of our ability to apply physical reasoning to understand the results: the higher the interpretability of a machine learning model, the easier it is for us to understand why certain predictions have been made [51]. There are global and local methods to assess models interpretability: a global method informs about how a model makes a decision as a whole, and local methods explain individual predictions. SHapley Additive exPlanations (SHAP) is a method derived from game theory to explain individual predictions by computing the contribution of each feature to the prediction [44]. Essentially, SHAP provides a framework to understand how the features in the input data are related to the model output.

In our particular application, we use SHAP to assess which bands (i.e., features) are globally—as in global interpretability—more important to explain the spatial distribution of SGD (model output). Note that we added the thermal band (preprocessed as SST) for the features importance assessment. Features with larger absolute Shapley values are more important. The average of the absolute Shapley values per feature across the entire data set yields the global importance (Figure 4.6).

The summary plot (Figure 4.7) shows the relationship between the value of a feature and




Figure 4.6: SHAP feature importance measured as the mean absolute Shapley values, which indicate the magnitude of the impact on the model. Features are sorted by decreasing importance. Sea surface temperature (SST, as inferred from the thermal infrared band, TIR) was the most important feature, changing the predicted absolute SGD probability on average by 0.7 percentage points. The remaining features, ordered by importance, are the green (G), blue (B), near infrared (NIR), red (R), and ultra-blue (uB) bands.

the impact on the prediction. As expected, low values (blue dots) of SST tend to contribute to the classification as SGD (positive SHAP values). Recall that the training input was manually selected specifically based on low SST values (in addition to color patterns and proximity and high radon activity samples). It is interesting to note that the green band (G) is the second most important feature—and the most important when SST is excluded. Whereas SGD is linked to color patterns driven by algal chlorophyll content, if reflectance in the green band is high (red dots), it leads to a pixel being classified as non-SGD (negative SHAP values).

In this particular case study, this seems to be a consequence of the greener plumes apparently originating from surface waters discharge through the river mouth within the Waterford Harbour (as seen in Figure 4.1). Blue reflectance (B) has the highest positive impact on SGD predictions—although only in a reduced number of pixels (note the low jitter in the y-axis towards the high SHAP values)—, including higher values of reflectance, which could represent a lighter blue color. High red (R) and near infrared (NIR) reflectance can be linked to predicted SGD—as also seen in the expected spectrum shown in Figure 2.2)—, however, high values in those bands can be similarly present in non-SGD pixels. Ultra-blue reflectance (uB; also known as the coastal band in Landsat 8) shows as the least important feature according to SHAP, although high reflectance values are mostly linked to non-SGD




Figure 4.7: SHAP summary plot. Each point represents a Shapley value for a feature and a prediction. The position on the y-axis is determined by the feature, ordered by decreasing importance, and on the x-axis by the Shapley value. The color represents the value of the feature from low to high. Overlapping points are jittered in y-axis direction, representing the distribution of the Shapley values per feature. Low values of sea surface temperature (SST, as inferred from the thermal infrared band, TIR) increase the probability of a point being classified as PSGD. The remaining features, ordered by importance, are the green (G), blue (B), near infrared (NIR), red (R), and ultra-blue (uB) bands.

predictions, likely as a consequence of bottom reflectance (a confounding factor).




Figure 4.8: Distribution of the probabilities of each pixel being PSGD (as given by CNN classification) when including the thermal band in the model.

The analysis of the features importance complement our understanding of the role of SST in determining PSGD pixels. Incorporating the thermal information in the model input yields PSGD pixels with higher probability in their classification (Figure 4.8). Consequently, the corresponding spatial distribution of PSGD pixels is more constrained and less likely to include false positives (see, for example, Figure 4.9 shows PSGD plumes with more than 90% probability). This results suggest that for future models the thermal band could be included when available. Nevertheless, we still recognize the value of this approach in enabling PSGD

detection when thermal information is not available. Additionally, feature importance might vary between areas, namely in cases where surface waters carry significant concentration of color-changing particles (e.g., suspended sediments, algal chlorophyll) or where bottom reflectance is confounding factor, thus highlighting the need for a more comprehensive training input collection.

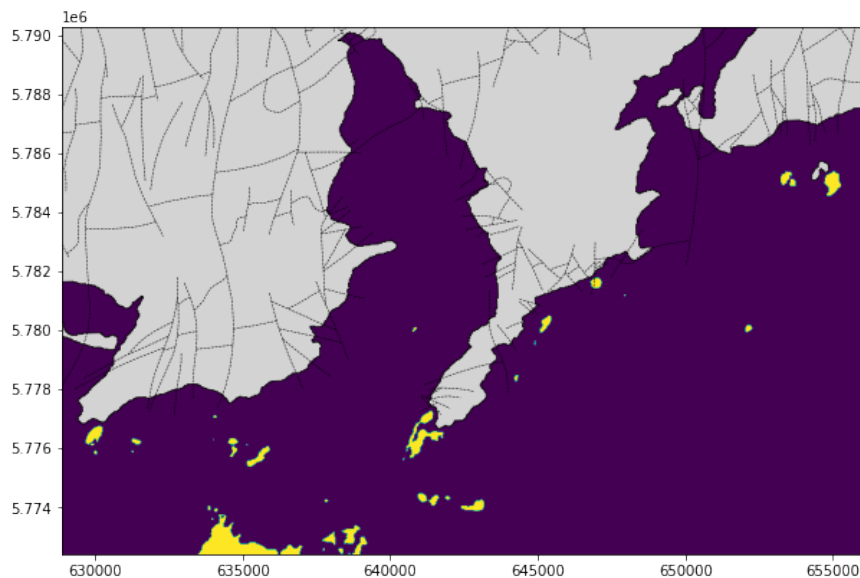


Figure 4.9: Pixels classified as PSGD by the CNN ($p > 97\%$) when including the thermal band in the model.

4.3.4 Limitations and Future Research Directions

As it is inherent to methods based on machine learning techniques, the method’s efficiency is highly dependent on the input data used at the training stage. The accuracy of the labeled data is directly related to the accuracy of the results. Here in particular, labeled plumes were inferred from SST patterns near locations of SGD occurrences mapped at a different date than the scene acquisition time. Therefore, uncertainty in the results is a function of the spatial and temporal distance between thermally anomalous plumes and previously identified SGD.

In addition to the accuracy of the training data, the size of the input data set is also paramount to achieve good results with a neural network approach. In this work, the size of the training data set is considered low for a neural network, although the number of

parameters to constrain is also smaller than a typical image-based neural network. Nevertheless, besides collecting more labeled data, additional data sources can be incorporated, including bathymetry (where a deeper profile corresponds to a diminished SGD potential), and distance to mapped faults (and corresponding projections; the closer to a known fault the higher the potential). Additionally, data augmentation can be extended to include spectral modifications, such as noise and simulation of other sensors characteristics (e.g., center wavelength and bandwidth, instantaneous field of view). There is also the question of the applicability of the trained network to an input data set from a different region or a different season. Therefore, an expanded training data set should also include data collected from varying locations (e.g., the southeastern and southwestern ends of the Irish coastline experience different currents than in the Waterford Harbour) and seasons (e.g., to capture different magnitudes of SGD flow). Accordingly, the test data set should also include data from yet another location and time of the year to assess how transferable the model is.

This approach is data-driven, hence improving data quality and quantity should be the priority for the next steps. Subsequently, by applying this method to multiple scenes, time series of detected PSGD can be generated and analyzed in a similar way as in Chapter 3. Then, plume area (and total area of plumes) can be tracked over time and compared to trends obtained through the methods in Chapter 3. If similar seasonal patterns are not verified (e.g., total plume area varying with groundwater levels), possible explanations could be found by analyzing variations in spectra—per band or in a composite or ratio way—over time; such analysis could assist in understanding the limitation of applying a model trained with data from a distinct season. Finally, as we have mentioned earlier, the implemented neural network is very simplistic and there are several avenues for improvement in terms of the hyper-parameters (e.g., learning rate, dropout), including using a pretrained model with remote sensing data sets [25].

4.4 Conclusions

SGD can drive changes in the optical and thermal spectrum of the sea surface—e.g., color and temperature. These changes can be very subtle and there is no guarantee of a linear relationship between color and temperature. This is a type of problem where machine learning techniques can improve upon parametric approaches. In this study, we demonstrate the ability of a neural network to capture SGD signatures embedded in remote sensing signals, and ultimately to detect PSGD through a simple, yet powerful framework, which can be extended to implement a machine learning-based product for global PSGD mapping.

5.0 Conclusions

In this dissertation, a common theme is presented throughout the three core chapters: submarine groundwater discharge (SGD) is observable from spaceborne sensors. Groundwater temperature has been used widely as a natural tracer, providing quantifiable evidence of groundwater discharge in coastal areas. Since approximately four decades ago, surface temperature can be measured from space, and it has a long history in SGD applications. However, the thermal signature that is produced by SGD influence is challenging to capture and non-exclusive to SGD. Therefore, SGD locations identified through this thermal response can only be considered estimates, or—as aforementioned throughout this dissertation—potential SGD (PSGD). Here, I improve upon the thermal-based approach by combining it with color. Ocean color refers to the reflectance in the optical spectrum from the oceans’ surface, and has been linked to various properties of the ocean (e.g., turbidity, sediments load, chlorophyll content). This is a novel approach that combines the thermal and optical responses to SGD to advance the remote sensing of groundwater flows.

A few concluding remarks should be outlined and emphasized. First, combining sea surface temperature and color results in a more reliable PSGD detection (Chapter 2). Through this combination, PSGD can be detected in locations where SGD has been proven to occur (Chapter 2), as well as in locations missed by approaches exclusively based on temperature (Chapters 2 and 3). Second, SGD signatures can be expressed as weaker or stronger responses (Chapter 3), which can vary across seasons (Chapter 3), possibly as a function of the magnitude of SGD flows and the local sea conditions (e.g., temperature, currents, other inputs). For this reason, it is advantageous to have a multi-method approach that leverages different sensitivities (Chapters 2 and 3) and that seeks SGD responses from different perspectives (Chapters 2 and 4). Third, whereas the thermal response alone can provide substantial evidence for SGD presence (Chapter 2), in particular where a time persisting signature is observed (Chapter 3), it should be evaluated in a manner that captures spatially (Chapter 2) and temporally (Chapter 3) distinct anomalies. Moreover, depending on location (Chapter 2) or season (Chapter 3), the thermal signature can be inadequate to reliably

detect PSGD. In such cases, a reduced dependency on the thermal signal can improve the ability to detect PSGD (Chapter 4).

With this work, a novel foundation is laid, and many opportunities for future research are identified. The relation between sea surface temperature, ocean color, and SGD, is complex and hindered by several confounding factors (e.g., upwelling, surface water inputs, bathymetry influence, clouds and cloud shade). Determining and separating the contribution of confounding factors to the observed remote sensing signals would improve the reliability of PSGD detection methods. Possibly, the influence of confounding factors can be reduced by quantitatively incorporating additional data sets, including bathymetry, currents (sea and wind), geological features (fractures, aquifer composition), as well as additional remote sensing data sources (e.g., Sentinel-2). In particular, additional reference spectra would contribute to overcoming the single-spectrum limitation recognized in Chapter 2. Extended field-based observations would fulfill a proper uncertainty assessment of PSGD and clarify the link between groundwater level (or, more generally, storage) and SGD flows. More practically, the temporal analysis of PSGD in Chapter 3 would surely benefit from a less restricted collection and processing of readily available remote sensing imagery, which could enable a complete understanding of PSGD seasonality (i.e., covering all seasons), and a climate change contextualization of PSGD variability. Potentially, the methods presented here form a prototype for the operationalization of PSGD detection, whereby Chapter 4 is essentially the proof of concept for a promising framework for global PSGD mapping. Finally, SGD research still has many opportunities to quantify SGD flows globally, to detail and demonstrate SGD exploitation as a water resource, and—perhaps more interestingly—how the global mapping of PSGD could be leveraged to improve the global water cycle assessment, and how it could be materialized in a broader, societal impact.

Appendix A Supporting Information for Chapter 2




Figure A1: Distribution of the distance between PSGD plumes and the coastline, for plumes derived from imagery acquired over Ireland’s south coast on (1) 2010-06-02 from Landsat 7 and on (2) 2016-06-01 from Landsat 8, as resulting from applying (a) derivative analysis and (b) angular distance.

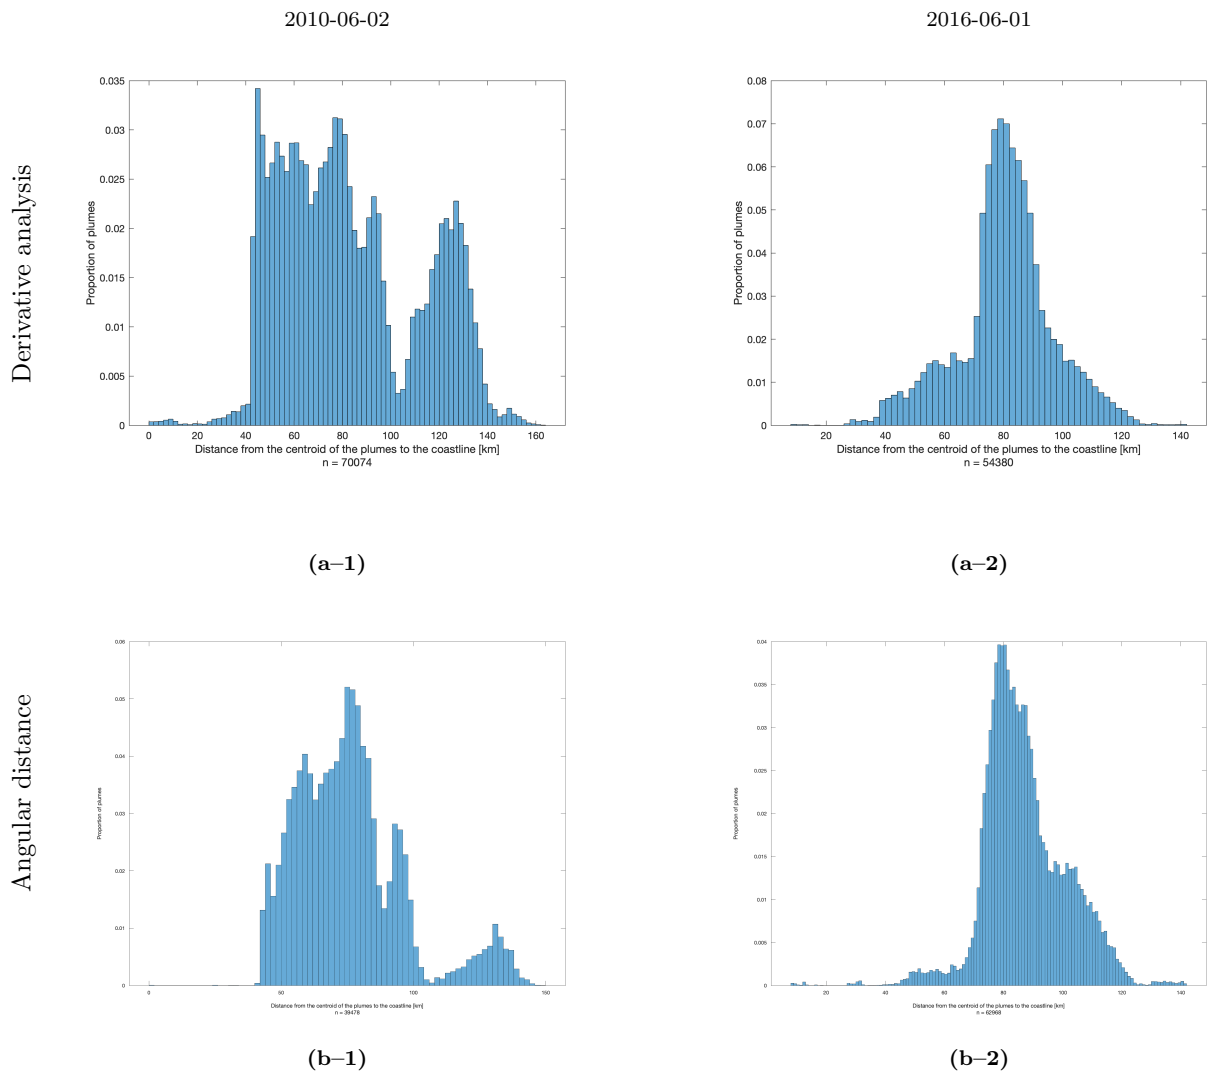


Figure A2: Distribution of the distance between PSGD plumes and the coastline, for plumes derived from imagery acquired over the Island of Hawai'i on (1) 2007-04-29 from Landsat 7 and on (2) 2019-04-06 from Landsat 8, as resulting from applying (a) derivative analysis and (b) angular distance.




Figure A3: Distribution of the area (in number of pixels) of PSGD plumes derived from imagery acquired over Ireland’s south coast on (1) 2010-06-02 from Landsat 7 and on (2) 2016-06-01 from Landsat 8, as resulting from applying (a) derivative analysis and (b) angular distance. Plumes with area above 30 pixels are included in the last bin.


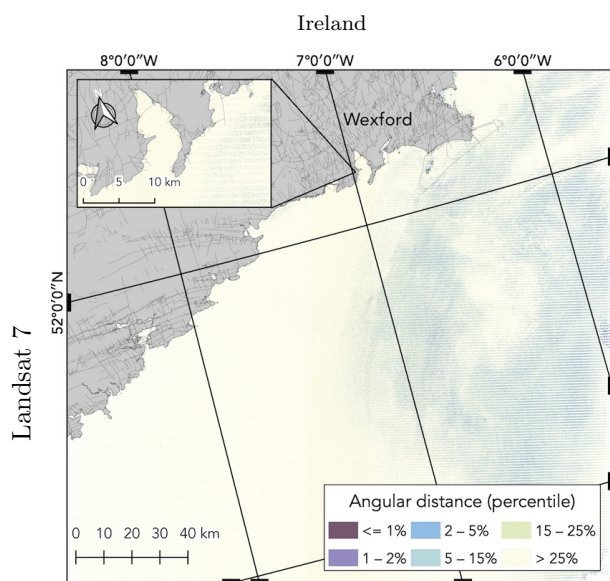
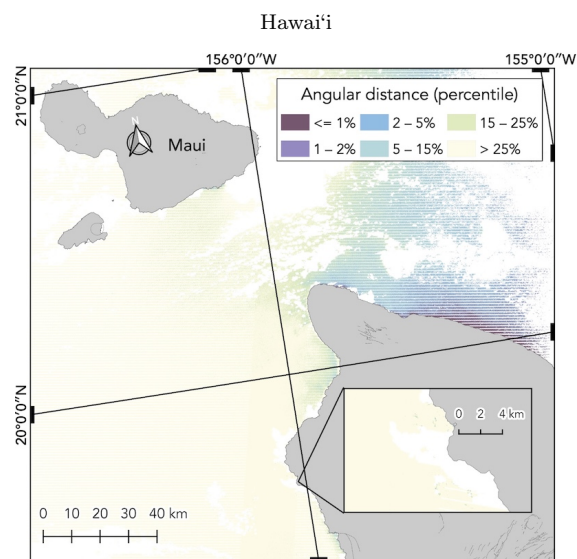


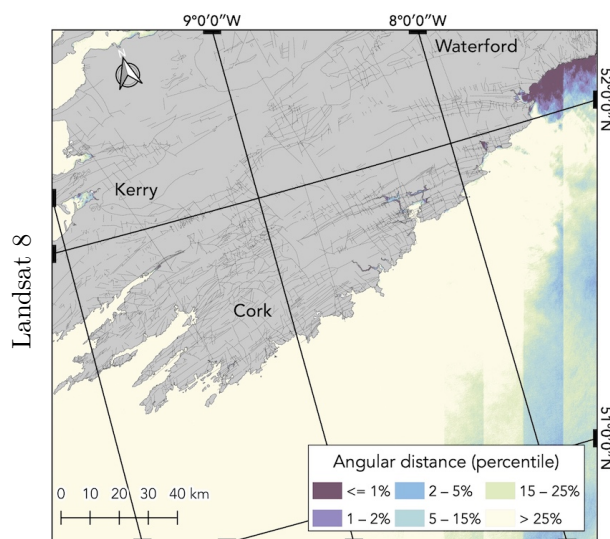
Figure A4: Distribution of the area (in number of pixels) of PSGD plumes derived from imagery acquired over the Island of Hawai'i on (1) 2007-04-29 from Landsat 7 and on (2) 2019-04-06 from Landsat 8, as resulting from applying (a) derivative analysis and (b) angular distance. Plumes with area above 30 pixels are included in the last bin.



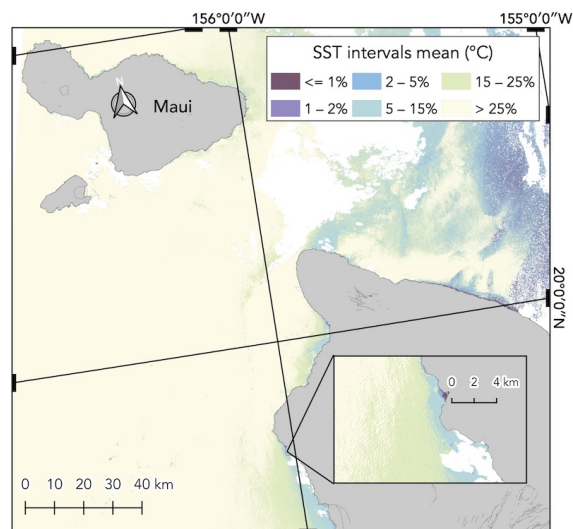
(a-1)



(a-2)



(b-1)



(b-2)

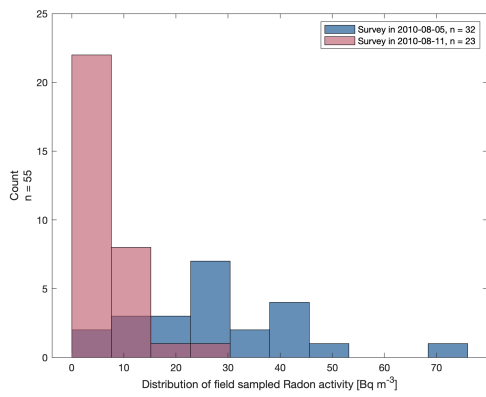
Figure A5: Angular distance maps, showing distribution between percentiles 1st, 2nd, 5th, 15th, and 25th percentiles, for imagery captured over Ireland's south coast on 2010-06-02 from Landsat 7 (L7) and on 2016-06-01 from Landsat 8 (L8), and over the Island of Hawai'i on 2007-04-29 from Landsat 7 and on 2019-04-06 from Landsat 8.

Table A1: Angular distance percentiles: 1st, 2nd, 5th, 15th, and 25th for imagery captured over Ireland’s south coast on 2010-06-02 from Landsat 7 (L7) and on 2016-06-01 from Landsat 8 (L8), and over the Island of Hawai’i on 2007-04-29 from Landsat 7 and on 2019-04-06 from Landsat 8.


Percentiles		1st	2nd	5th	15th	25th
Ireland	L7:	0.38405	0.39865	0.41973	0.45397	0.48035
	L8:	0.4608	0.48189	0.49759	0.52166	0.54667
Hawai’i	L7:	0.38577	0.39545	0.41117	0.44592	0.47783
	L8:	0.47181	0.49064	0.52952	0.63167	0.70292

Table A2: Number of pixels that are (a) within the two lowest SST intervals (SST Low), shortlisted by derivative analysis (DA), flagged as PSGD by derivative analysis (which are within the two lowest SST intervals; $DA \cap SST$), and percent of PSGD pixels relative to those within the two lowest SST intervals ($DA \mid SST \text{ Low}$), and relative to those shortlisted by derivative analysis ($SST \text{ Low} \mid DA$). Similarly, (b) shows the corresponding count and percents for angular distance (AD). (c) shows the number and percent of PSGD pixels (i.e., also considering the two lowest SST intervals) resulting from cross-checking derivative analysis and angular distance ($DA \cap AD$).

(a) Derivative analysis (DA).					
Case study	SST Low	DA	$DA \cap SST$	$DA \mid SST \text{ Low} (\%)$	$SST \text{ Low} \mid DA (\%)$
Ireland L7	5567612	7323719	2912886	52.32	39.77
Ireland L8	16816380	3343242	1969828	11.71	58.92
Hawai'i L7	6930116	168601	102275	1.48	60.66
Hawai'i L8	7604217	203147	89690	1.18	44.15
(b) Angular distance (AD).					
Case study	SST Low	AD	$AD \cap SST$	$AD \mid SST \text{ Low} (\%)$	$SST \text{ Low} \mid AD (\%)$
Ireland L7	5567612	226561	100710	1.81	44.45
Ireland L8	16816380	220345	166120	0.99	75.39
Hawai'i L7	6930116	148883	116904	1.69	78.52
Hawai'i L8	7604217	248209	214815	2.82	86.55
(c) PSGD resulting from derivative analysis (DA) cross-checked with angular distance (AD).					
Case study	$DA \cap AD$	$AD \mid DA (\%)$	$DA \mid AD (\%)$		
Ireland L7	88466	3.04	87.84		
Ireland L8	165649	8.41	99.72		
Hawai'i L7	27603	26.99	23.61		
Hawai'i L8	44852	50.01	20.88		



(a)



(b)

Figure A6: Distribution of radon activity (Bq m^{-3}) and salinity (psu) from field observations in (a) Waterford Harbour, Ireland [92], and in (b) four bays (Kailua, Kealahakua, Kiholo, and Honokohau Harbor) of the Island of Hawai'i [27], respectively. Field surveys in Ireland were conducted in the 5th and 11th of 2010-08, gathering a total of 55 data points that covered approximately 100 km^2 of one bay (Waterford Harbour). In Hawai'i, field surveys were conducted over 10 different dates between 2005-08 and 2007-05, gathering a total of 67 data points that covered approximately a 100 km coastline stretch over the four bays.

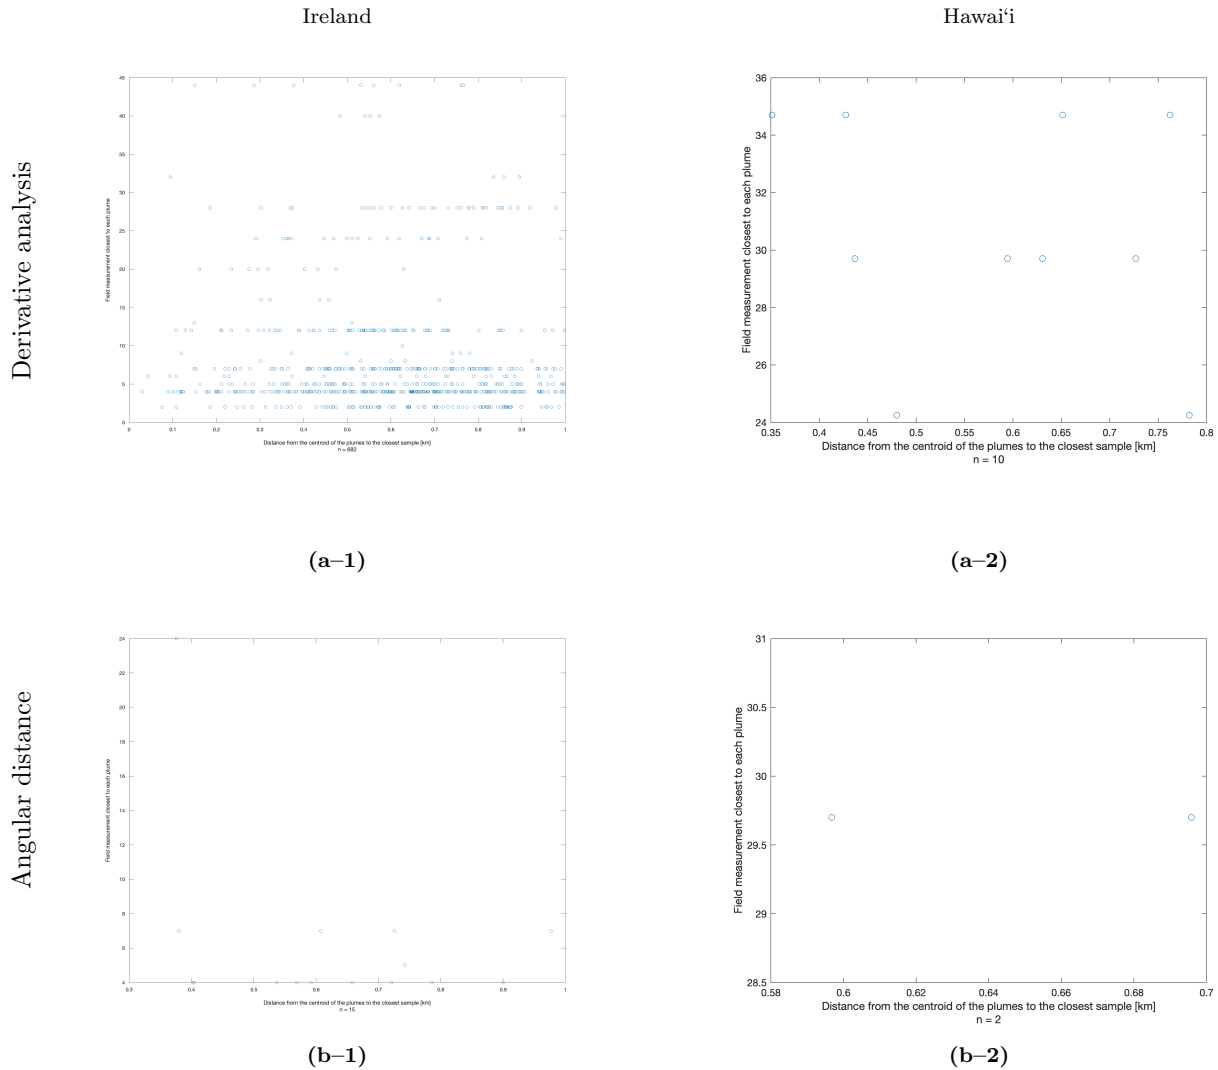


Figure A7: Distribution of the distance between PSGD plumes (from derivative analysis and angular distance) and the field samples taken in (1) Waterford Harbour, Ireland, and in (2) four bays of the Island of Hawai'i, respectively (horizontal axis), and the distribution of the corresponding sampled values (vertical axis). In Ireland, field observations measured radon activity (Bq m^{-3}), whereas in Hawai'i they measured salinity (psu). Only plumes located up to 5 km of samples are included here.

Chlorophyll Concentration, OC3 Algorithm

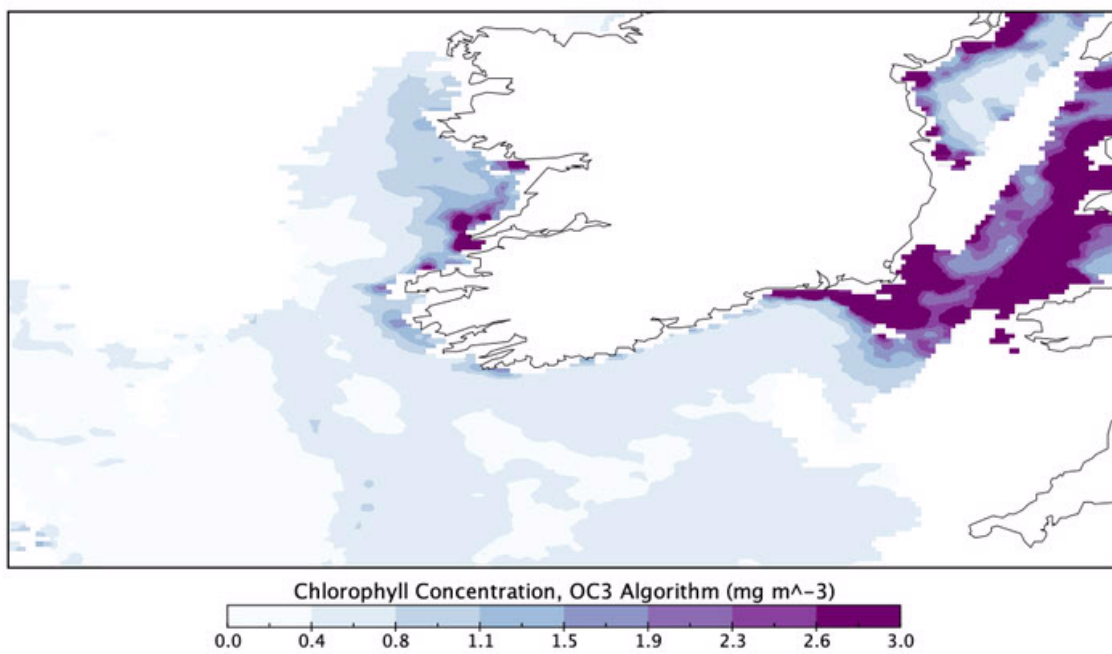


Figure A8: Chlorophyll concentration (mg m^{-3}) from the daily MODIS Level-3 product[56] of 2016-06-01 over south Ireland. Obtained from NASA Goddard Space Flight Center, Ocean Ecology Laboratory, Ocean Biology Processing Group.

Appendix B Supporting Information for Chapter 3

B.1 Introduction

This supporting information document includes more information regarding the data sets utilized: Landsat 8 and MODIS products, in addition to in situ observations of groundwater and sea surface temperature. We also include here additional results of our methods (non-essential for the conclusions presented in the main document).

B.2 Data

B.2.1 Remote Sensing

B.2.1.1 Table

Table B1: Product IDs of the 62 Landsat 8 and 29 MODIS scenes used in this study. Dates typeset in boldface indicate Landsat 8 overpasses (path 206, row 24) with corresponding groundwater temperature records within 30 days of the acquisition date. MODIS products were selected to only include the those dates (boldface) within the corresponding 8-day period. When there was no MODIS product available for a given date, the week prior was selected instead (indicated with a star in front of the MODIS Product ID).

Date	Season	Landsat 8 Product ID	MODIS Product ID
2013-07-04	Summer	LC08_L2SP_206024_20130704_20200912_02_T1	AQUA_MODIS.20130704_20130711.L3m.8D.SSTMMasked.SST.sst.4km
2013-07-20	Summer	LC08_L2SP_206024_20130720_20200912_02_T1	
2013-08-05	Summer	LC08_L2SP_206024_20130805_20200912_02_T1	
2014-04-18	Spring	LC08_L2SP_206024_20140418_20200911_02_T1	
2014-06-05	Spring	LC08_L2SP_206024_20140605_20200911_02_T1	AQUA_MODIS.20140602_20140609.L3m.8D.SSTMMasked.SST.sst.4km
2014-07-23	Summer	LC08_L2SP_206024_20140723_20200911_02_T1	
2014-09-09	Summer	LC08_L2SP_206024_20140909_20200911_02_T1	
2014-10-11	Autumn	LC08_L2SP_206024_20141011_20200910_02_T1	
2015-01-15	Winter	LC08_L2SP_206024_20150115_20200910_02_T1	
2015-01-31	Winter	LC08_L2SP_206024_20150131_20200909_02_T1	
2015-02-16	Winter	LC08_L2SP_206024_20150216_20200909_02_T1	
2015-03-20	Winter	LC08_L2SP_206024_20150320_20200909_02_T1	AQUA_MODIS.20150314_20150321.L3m.8D.SSTMMasked.SST.sst.4km
2015-04-05	Spring	LC08_L2SP_206024_20150405_20200909_02_T1	AQUA_MODIS.20150330_20150406.L3m.8D.SSTMMasked.SST.sst.4km
2015-04-21	Spring	LC08_L2SP_206024_20150421_20200909_02_T1	AQUA_MODIS.20150415_20150422.L3m.8D.SSTMMasked.SST.sst.4km
2015-08-27	Summer	LC08_L2SP_206024_20150827_20200908_02_T1	
2015-09-12	Summer	LC08_L2SP_206024_20150912_20200908_02_T1	AQUA_MODIS.20150906_20150913.L3m.8D.SSTMMasked.SST.sst.4km
2015-09-28	Autumn	LC08_L2SP_206024_20150928_20200908_02_T1	AQUA_MODIS.20150922_20150929.L3m.8D.SSTMMasked.SST.sst.4km
2015-10-14	Autumn	LC08_L2SP_206024_20151014_20200908_02_T1	AQUA_MODIS.20151008_20151015.L3m.8D.SSTMMasked.SST.sst.4km
2016-07-28	Summer	LC08_L2SP_206024_20160728_20200906_02_T1	
2016-08-29	Summer	LC08_L2SP_206024_20160829_20200906_02_T1	
2016-09-14	Summer	LC08_L2SP_206024_20160914_20200906_02_T1	AQUA_MODIS.20160913_20160920.L3m.8D.SSTMMasked.SST.sst.4km
2016-10-16	Autumn	LC08_L2SP_206024_20161016_20200905_02_T1	AQUA_MODIS.20161015_20161022.L3m.8D.SSTMMasked.SST.sst.4km
2016-11-17	Autumn	LC08_L2SP_206024_20161117_20200905_02_T1	
2017-01-04	Winter	LC08_L2SP_206024_20170104_20200905_02_T1	
2017-01-20	Winter	LC08_L2SP_206024_20170120_20200905_02_T1	
2017-02-05	Winter	LC08_L2SP_206024_20170205_20200905_02_T1	
2017-03-09	Winter	LC08_L2SP_206024_20170309_20200904_02_T1	AQUA_MODIS.20170306_20170313.L3m.8D.SSTMMasked.SST.sst.4km
2017-03-25	Spring	LC08_L2SP_206024_20170325_20200904_02_T1	AQUA_MODIS.20170322_20170329.L3m.8D.SSTMMasked.SST.sst.4km
2017-04-26	Spring	LC08_L2SP_206024_20170426_20200904_02_T1	AQUA_MODIS.20170423_20170430.L3m.8D.SSTMMasked.SST.sst.4km

(Continued on next page.)

Date	Season	Landsat 8 Product ID	MODIS Product ID
2017-07-31	Summer	LC08_L2SP_206024_20170731_20200903_02_T1	AQUA_MODIS.20170728_20170804.L3m.8D.SSTMMasked.SST.sst.4km
2017-09-01	Summer	LC08_L2SP_206024_20170901_20200903_02_T1	
2017-09-17	Summer	LC08_L2SP_206024_20170917_20200903_02_T1	AQUA_MODIS.20170914_20170921.L3m.8D.SSTMMasked.SST.sst.4km
2017-10-03	Autumn	LC08_L2SP_206024_20171003_20200903_02_T1	AQUA_MODIS.20170930_20171007.L3m.8D.SSTMMasked.SST.sst.4km
2017-11-04	Autumn	LC08_L2SP_206024_20171104_20200902_02_T1	
2018-01-07	Winter	LC08_L2SP_206024_20180107_20200902_02_T1	
2018-03-28	Spring	LC08_L2SP_206024_20180328_20200901_02_T1	
2018-04-29	Spring	LC08_L2SP_206024_20180429_20200901_02_T1	
2018-07-02	Summer	LC08_L2SP_206024_20180702_20200831_02_T1	
2018-08-19	Summer	LC08_L2SP_206024_20180819_20200831_02_T1	AQUA_MODIS.20180813_20180820.L3m.8D.SSTMMasked.SST.sst.4km
2018-09-04	Summer	LC08_L2SP_206024_20180904_20200831_02_T1	AQUA_MODIS.20180829_20180905.L3m.8D.SSTMMasked.SST.sst.4km
2018-10-06	Autumn	LC08_L2SP_206024_20181006_20200830_02_T1	AQUA_MODIS.20180930_20181007.L3m.8D.SSTMMasked.SST.sst.4km
2018-10-22	Autumn	LC08_L2SP_206024_20181022_20200830_02_T1	
2018-11-07	Autumn	LC08_L2SP_206024_20181107_20200830_02_T1	AQUA_MODIS.20181101_20181108.L3m.8D.SSTMMasked.SST.sst.4km
2018-12-09	Autumn	LC08_L2SP_206024_20181209_20200830_02_T1	AQUA_MODIS.20181203_20181210.L3m.8D.SSTMMasked.SST.sst.4km
2019-05-18	Spring	LC08_L2SP_206024_20190518_20200828_02_T1	
2019-06-19	Spring	LC08_L2SP_206024_20190619_20200827_02_T1	AQUA_MODIS.20190618_20190625.L3m.8D.SSTMMasked.SST.sst.4km
2019-07-05	Summer	LC08_L2SP_206024_20190705_20200827_02_T1	AQUA_MODIS.20190704_20190711.L3m.8D.SSTMMasked.SST.sst.4km
2019-08-06	Summer	LC08_L2SP_206024_20190806_20200827_02_T1	
2019-10-09	Autumn	LC08_L2SP_206024_20191009_20200825_02_T1	AQUA_MODIS.20191008_20191015.L3m.8D.SSTMMasked.SST.sst.4km
2019-11-10	Autumn	LC08_L2SP_206024_20191110_20200825_02_T1	
2020-03-01	Winter	LC08_L2SP_206024_20200301_20200822_02_T1	
2020-04-02	Spring	LC08_L2SP_206024_20200402_20200822_02_T1	
2020-06-05	Spring	LC08_L2SP_206024_20200605_20200825_02_T1	
2020-07-23	Summer	LC08_L2SP_206024_20200723_20200911_02_T1	*AQUA_MODIS.20200711_20200718.L3m.8D.SSTMMasked.SST.sst.4km
2020-09-09	Summer	LC08_L2SP_206024_20200909_20200919_02_T1	
2020-10-11	Autumn	LC08_L2SP_206024_20201011_20201016_02_T1	
2020-10-27	Autumn	LC08_L2SP_206024_20201027_20201106_02_T1	
2020-11-12	Autumn	LC08_L2SP_206024_20201112_20210317_02_T1	AQUA_MODIS.20201108_20201115.L3m.8D.SSTMMasked.SST.sst.4km
2021-04-05	Spring	LC08_L2SP_206024_20210405_20210409_02_T1	
2021-05-07	Spring	LC08_L2SP_206024_20210507_20210517_02_T1	AQUA_MODIS.20210501_20210508.L3m.8D.SSTMMasked.SST.sst.4km
2021-08-27	Summer	LC08_L2SP_206024_20210827_20210901_02_T1	AQUA_MODIS.20210821_20210828.L3m.8D.SSTMMasked.SST.sst.4km
2021-09-28	Autumn	LC08_L2SP_206024_20210928_20211001_02_T1	AQUA_MODIS.20210922_20210929.L3m.8D.SSTMMasked.SST.sst.4km

B.2.1.2 Histograms




Figure B1: Relative frequency (%) of the number of valid observations per pixel in each season, in the Landsat 8 data set, spanning 2013-07-04 to 2021-09-28.




Figure B2: Cumulative relative frequency (%) of the number of valid observations per pixel in each season, in the Landsat 8 data set, spanning 2013-07-04 to 2021-09-28.

B.2.1.3 Maps




Figure B3: Pixels with at least 5 valid observations (yellow), in the Landsat 8 data set, between 2013-07-04 and 2021-09-28. Pixels with less than 5 valid observations are represented in blue. White pixels represent land and areas outside of the scenes footprint.

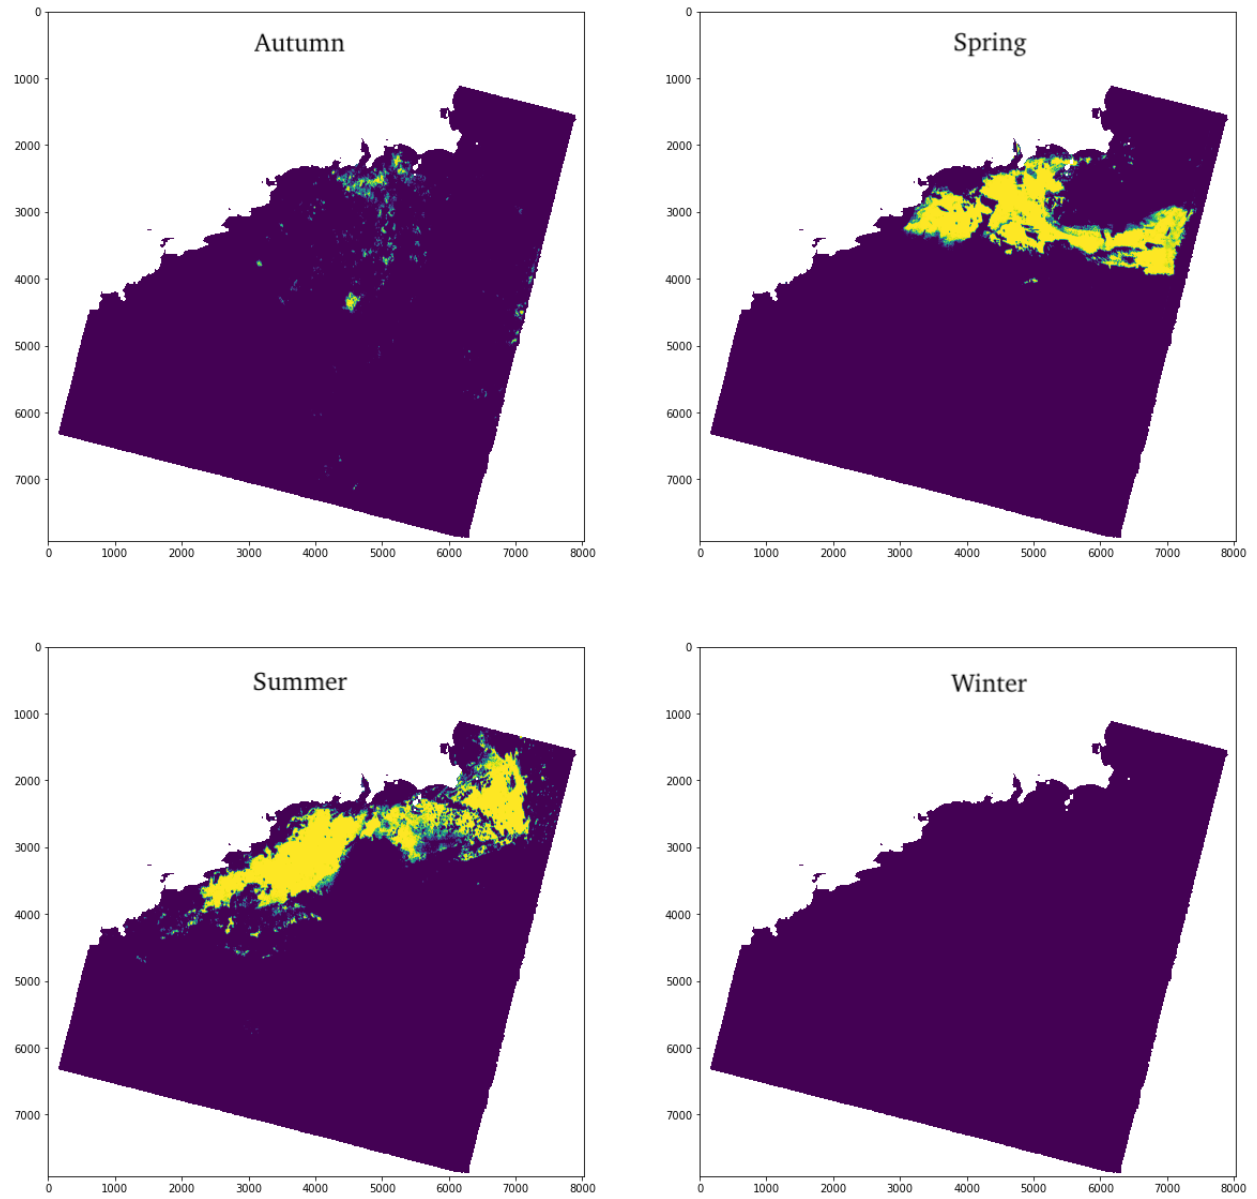


Figure B4: Pixels with at least 5 valid observations (yellow) in each season, in the Landsat 8 data set, between 2013-07-04 and 2021-09-28. Pixels with less than 5 valid observations are represented in blue. White pixels represent land and areas outside of the scenes footprint.

B.2.2 Groundwater Monitoring Stations

B.2.2.1 Table

Table B2: Percentiles of groundwater temperature measured from 20 monitoring wells in South Ireland. For each date, the 10th (GW p_{10}) and the 90th (GW p_{90}) percentiles are computed from different number of observations (GW Count), measured during an 8-day period which can start up to 5 days (Time Offset) before or after the scene acquisition date (Date).

Date	Season	Time Offset (days)	GW Count	GW p_{10} (°C)	GW p_{90} (°C)
2013-07-04	Summer	0	8	10.55	11.89
2014-06-05	Spring	0	11	10.20	13.30
2015-03-20	Winter	4	5	10.34	12.82
2015-04-05	Spring	2	10	10.39	11.41
2015-04-21	Spring	0	9	10.36	11.22
2015-09-12	Summer	2	11	10.60	11.30
2015-09-28	Autumn	0	15	11.00	12.10
2015-10-14	Autumn	1	2	11.27	11.83
2016-09-14	Summer	5	2	10.27	11.63
2016-10-16	Autumn	2	7	10.66	11.10
2017-03-09	Winter	4	12	10.72	11.86
2017-03-25	Spring	2	7	9.48	11.04
2017-04-26	Spring	0	11	10.20	11.10
2017-07-31	Summer	0	1	10.90	10.90
2017-09-17	Summer	1	14	10.70	12.03
2017-10-03	Autumn	1	9	10.50	12.06
2018-08-19	Summer	1	19	10.98	13.62
2018-09-04	Summer	1	2	10.84	13.56
2018-10-06	Autumn	2	11	10.50	11.80
2018-11-07	Autumn	0	5	10.64	12.96
2018-12-09	Autumn	1	10	10.09	11.18
2019-06-19	Spring	0	20	10.47	11.93
2019-07-05	Summer	1	4	10.27	15.13
2019-10-09	Autumn	1	10	11.20	12.25
2020-07-23	Summer	0	6	10.60	11.35
2020-11-12	Autumn	0	4	10.95	11.81
2021-05-07	Spring	1	8	10.00	11.03
2021-08-27	Summer	2	8	10.37	13.19
2021-09-28	Autumn	0	6	10.30	12.30

B.2.2.2 Map

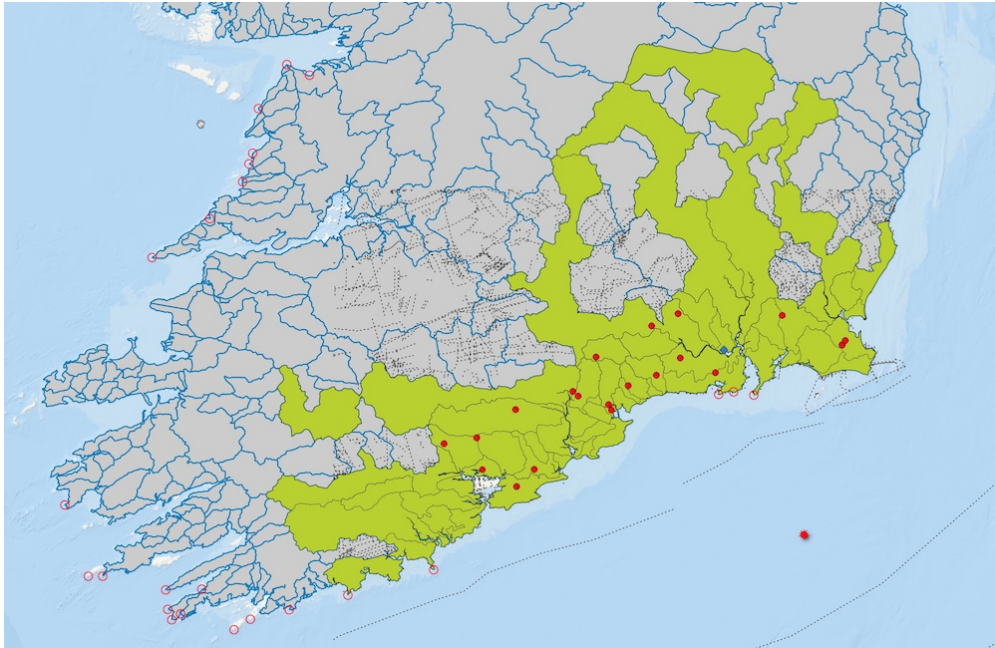


Figure B5: Subcatchments in South Ireland (blue outline), as defined by the Irish Environmental Protection Agency (EPA). Red circles show the location of the 20 boreholes in the Irish National Water Monitoring Stations that we used to retrieve groundwater temperature measurements. The corresponding subcatchments are represented in green. The red star shows the location of the Buoy M5 (51.69042, 6.704336; ~50 km offshore) in the Irish Buoy Network. Red circumferences show previously identified sites for potential SGD [92].

B.2.2.3 Graphs




Figure B6: Daily MODIS SST (blue; “SST, AQUA_MODIS, L3m.DAY.SST.sst.4km, Masked, SMI, NASA GSFC OBPG, R2019.0, Global, 0.041 66°, Science Quality, 2003-present (Daily)”) compared to the Buoy M5 (red). MODIS data displayed here were retrieved only from the pixel that corresponds to the buoy location.




Figure B7: Monthly aggregated measurements of groundwater temperature (°C) from 20 boreholes in the Irish National Water Monitoring Stations, sourced from the Irish Environmental Protection Agency (EPA; Resource identifier: EF.IE.EPA.MON.Waterstations).




Figure B8: Comparison between groundwater and sea surface temperature (SST; °C). Daily groundwater temperature measurements (purple dots and lines) are retrieved from 20 boreholes in the Irish National Water Monitoring Stations. The corresponding 10th (brown dots and lines) and 90th (pink dots and lines) percentiles are computed on a 8-day rolling window. Horizontal dashed gray lines show the overall 10th (bottom) and 90th (top) percentiles. Sub-daily SST measurements (blue dots) are retrieved from the Buoy M5 in the Irish Buoy Network, and are also averaged on a 8-day period (orange dots and lines). The 62 Landsat 8 overpasses are represented in vertical dashed lines. Vertical red lines correspond to Landsat 8 scenes without an matching groundwater measurement. Circles along the top and bottom represent whether PSGD should be in the coldest (green) or the warmest (red) SST interval, when comparing temperature measurements from the borehole stations and the buoy. This is only for illustration, since in practice we compared borehole measurements with each pixel of the MODIS SST scenes.




Figure B9: Same as Figure B8, but only for the period 2013-14.




Figure B10: Same as Figure B8, but only for the period 2016-18.

B.3 Results

B.3.1 Derivative Analysis

B.3.1.1 Histograms

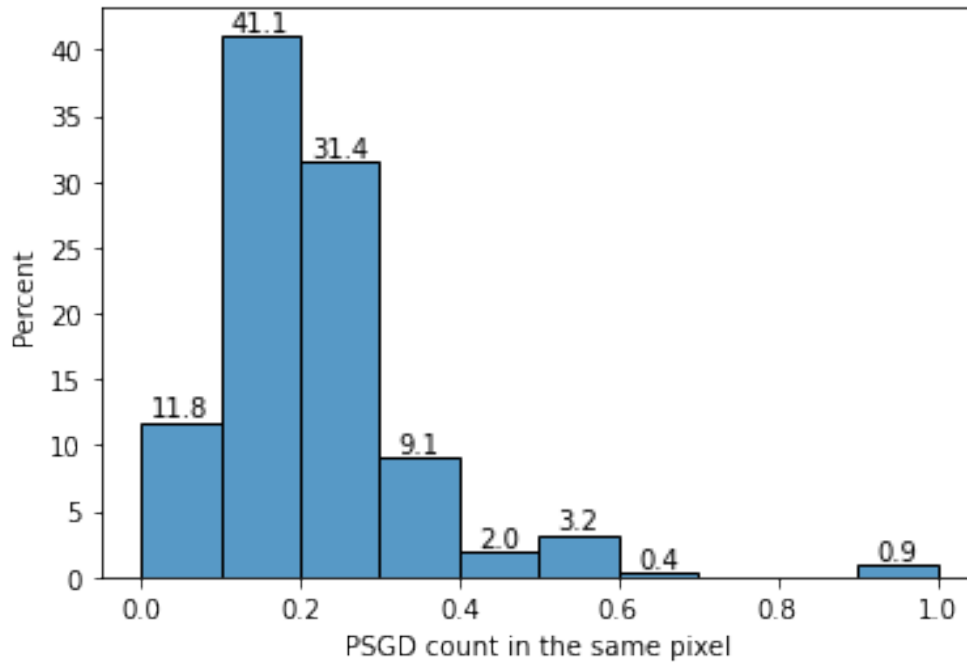


Figure B11: Relative frequency (%) of the number of times each pixel is flagged as PSGD by derivative analysis (with plume buffering), computed as a percent of the number of valid observations in the corresponding pixel.




Figure B12: Cumulative relative frequency (%) of the number of times each pixel is flagged as PSGD by derivative analysis (with plume buffering), computed as a percent of the number of valid observations in the corresponding pixel.




Figure B13: Relative frequency (%) of the number of times each pixel is flagged as PSGD, in each season, by derivative analysis (with plume buffering).

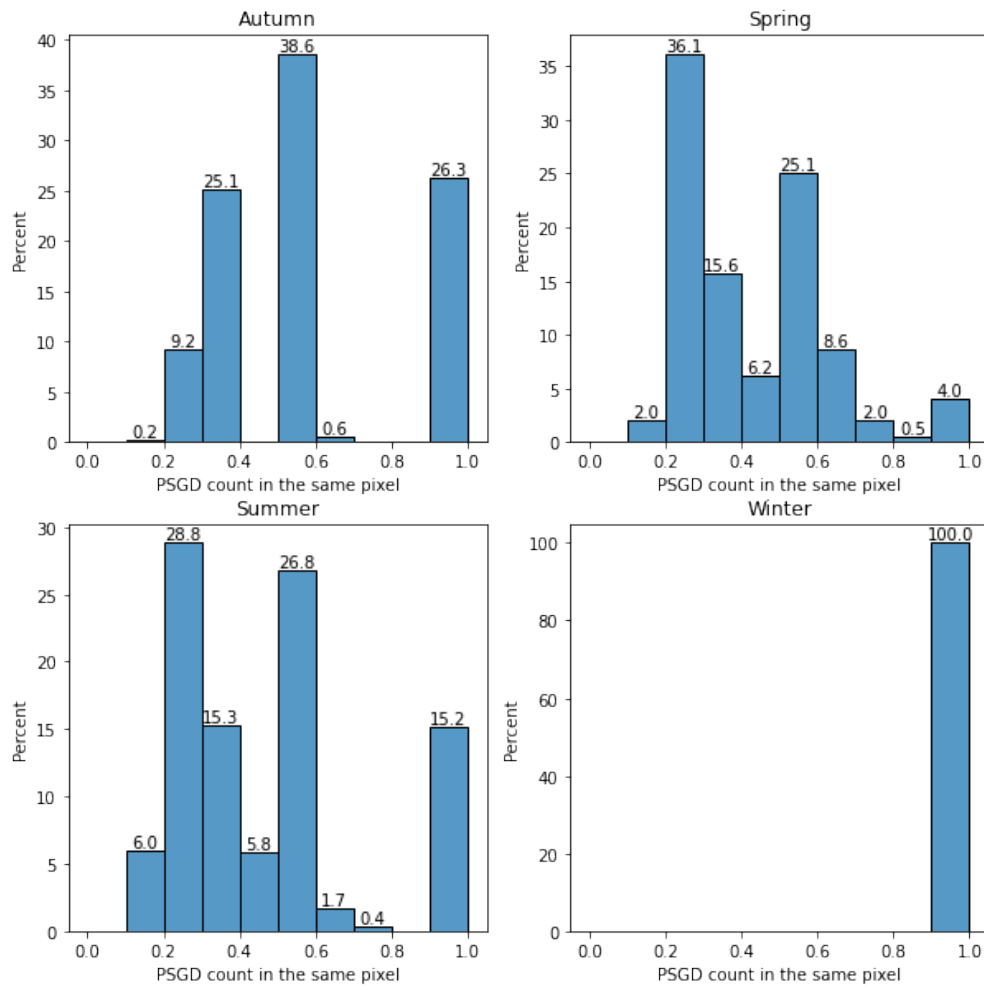


Figure B14: Relative frequency (%) of the number of times each pixel is flagged as PSGD, in each season, by derivative analysis (with plume buffering), computed as a percent of the number of valid observations in the corresponding pixel.




Figure B15: Relative frequency (%) of the number of times each pixel is flagged as PSGD, in each season, by derivative analysis (with plume buffering and crosschecked with SST), computed as a percent of the number of valid observations in the corresponding pixel.

B.3.1.2 Maps




Figure B16: PSGD detected by derivative analysis (without plume buffering and without crosschecking), and with temporal consistency of more than 50% of the valid observations. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents five categories of temporal consistency: > 50% (lighter orange) to > 90% (darker orange). Pixels with less than 5 valid observations are represented in gray. White pixels have at least 5 valid observation but no detectable PSGD. Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.

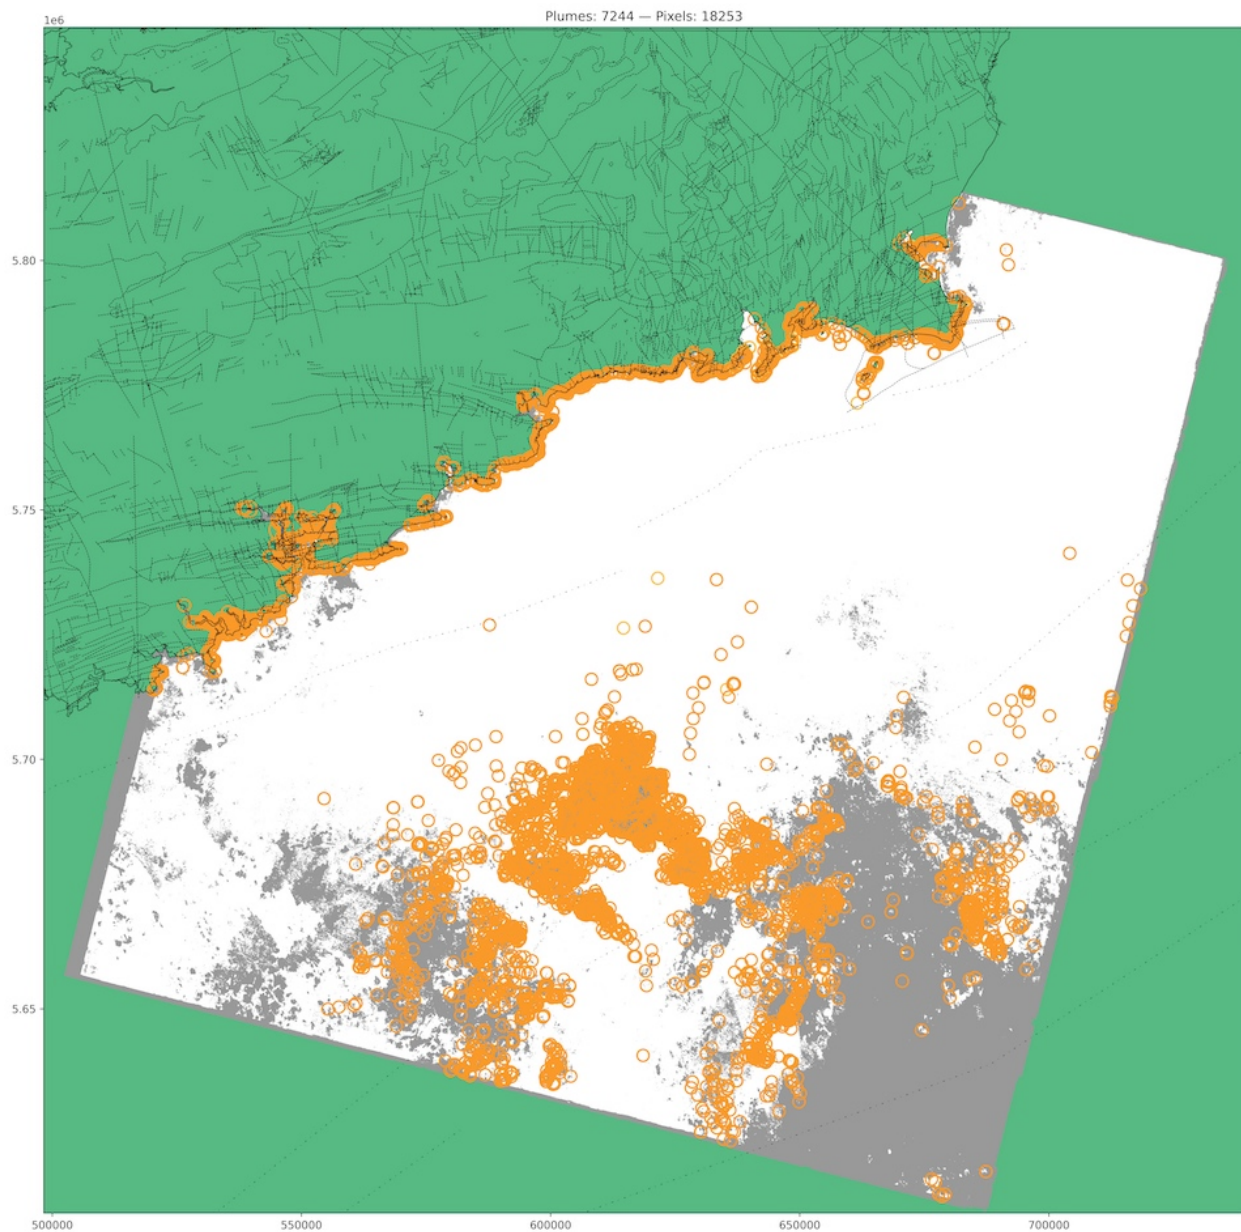


Figure B17: PSGD detected by derivative analysis (without plume buffering and without crosschecking), and with temporal consistency of more than 90% of the valid observations. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents five categories of temporal consistency: > 90% (lighter orange) to > 99% (darker orange). Pixels with less than 5 valid observations are represented in gray. White pixels have at least 5 valid observation but no detectable PSGD. Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.




Figure B18: PSGD detected by crosschecking derivative analysis and SST intervals (with plume buffering), and with temporal consistency of more than 70% of the valid observations. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents four categories of temporal consistency: > 70% (lighter orange) to > 95% (darker orange). Pixels with less than 5 valid observations are represented in gray. White pixels have at least 5 valid observation but no detectable PSGD. Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.




Figure B19: PSGD detected by crosschecking derivative analysis and SST intervals (with plume buffering), and with temporal consistency of more than 50% of the valid observations. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents five categories of temporal consistency: > 50% (lighter orange) to > 90% (darker orange). Pixels with less than 6 valid observations are represented in gray. White pixels have at least 6 valid observation but no detectable PSGD. Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.




Figure B20: PSGD detected by crosschecking derivative analysis and SST intervals (with plume buffering), and with temporal consistency of more than 70% of the valid observations. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents four categories of temporal consistency: > 70% (lighter orange) to > 95% (darker orange). Pixels with less than 6 valid observations are represented in gray. White pixels have at least 6 valid observation but no detectable PSGD. Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.

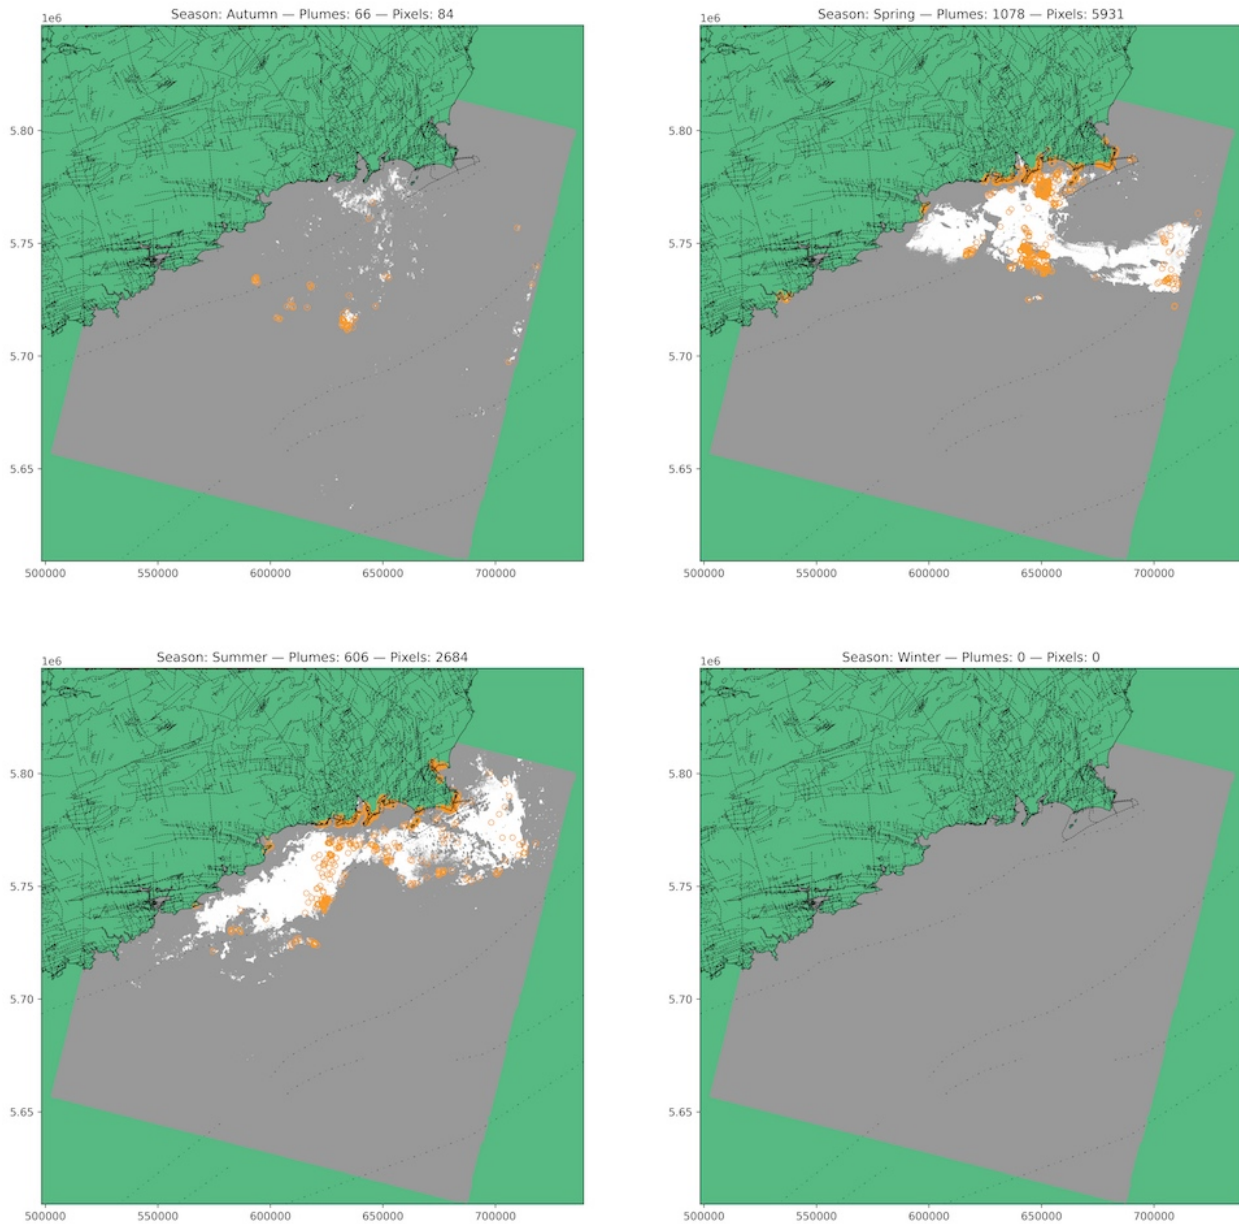


Figure B21: Seasonal PSGD detected by derivative analysis (without plume buffering and without cross-checking), and with temporal consistency of more than 90% of the valid observations. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents three categories of temporal consistency: > 90% (lighter orange) to > 99% (darker orange). Pixels with less than 5 valid observations are represented in gray. White pixels have at least 5 valid observation but no detectable PSGD. Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.




Figure B22: Seasonal PSGD detected by crosschecking derivative analysis and SST intervals (without plume buffering), and with temporal consistency of more than 50% of the valid observations. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents five categories of temporal consistency: > 50% (lighter orange) to > 90% (darker orange). Pixels with less than 5 valid observations are represented in gray. White pixels have at least 5 valid observation but no detectable PSGD (no PSGD detected in autumn and winter). Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.

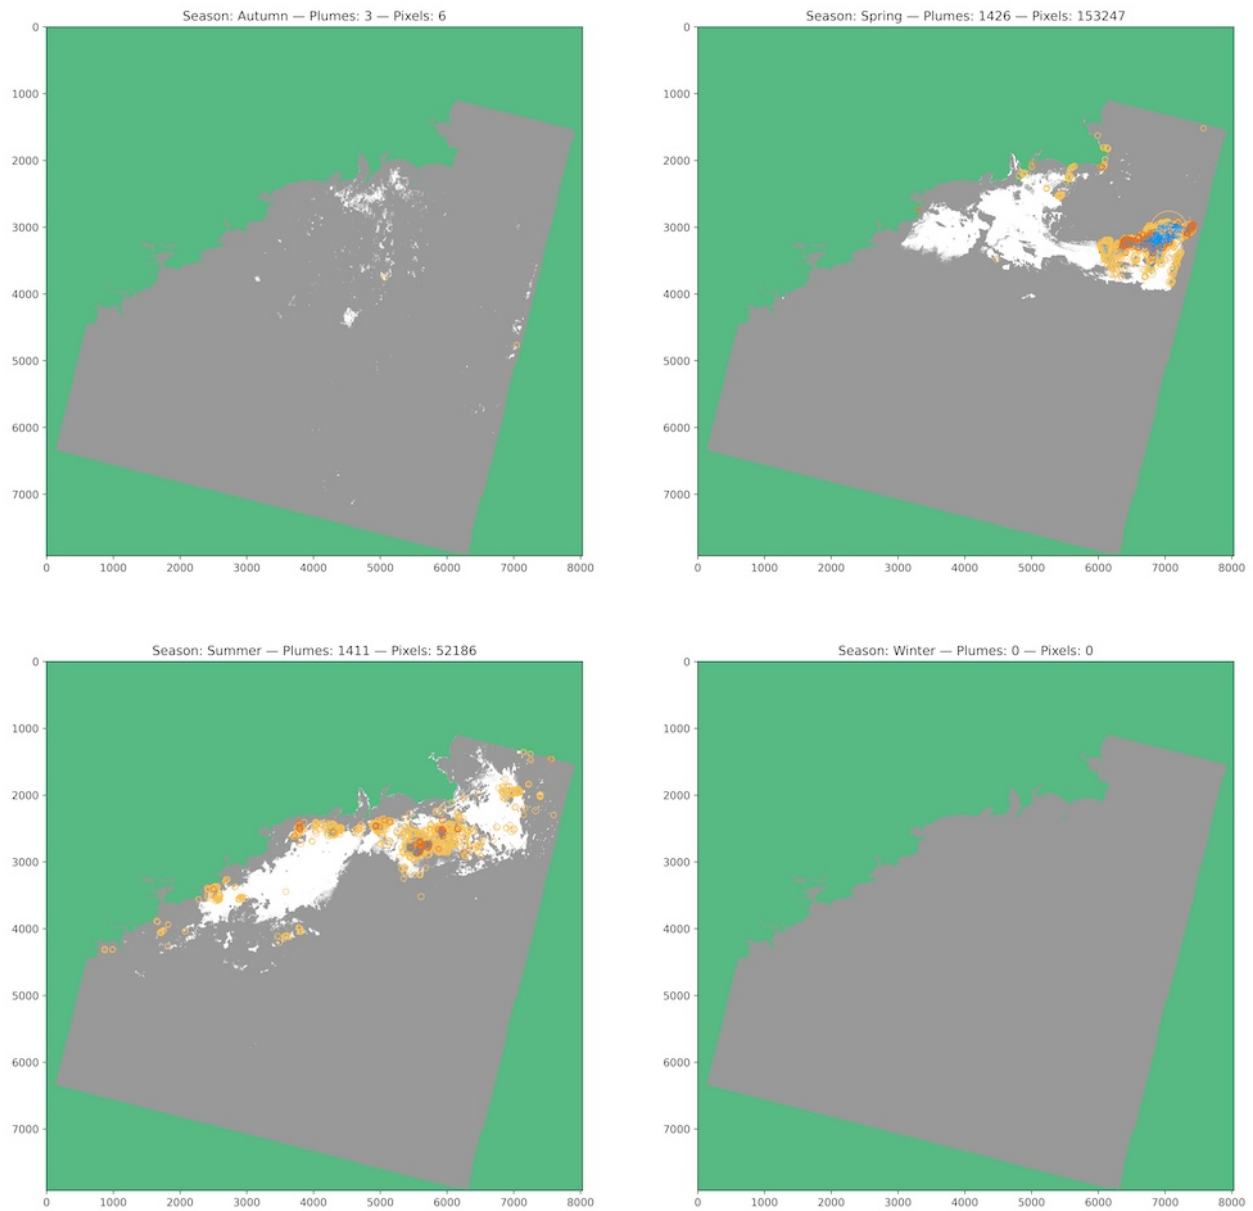


Figure B23: Seasonal PSGD detected by crosschecking derivative analysis and SST intervals (with plume buffering), and with temporal consistency of more than 50% of the valid observations. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents five categories of temporal consistency: > 50% (lighter orange) to > 90% (darker orange). Pixels with less than 5 valid observations are represented in gray. White pixels have at least 5 valid observation but no detectable PSGD (no PSGD detected in autumn and winter). Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.




Figure B24: Seasonal PSGD detected by crosschecking derivative analysis and SST intervals (with plume buffering), and with temporal consistency of more than 50% of the valid observations. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents five categories of temporal consistency: > 50% (lighter orange) to > 90% (darker orange). Pixels with less than 6 valid observations are represented in gray. White pixels have at least 6 valid observation but no detectable PSGD (no PSGD detected in autumn and winter). Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.

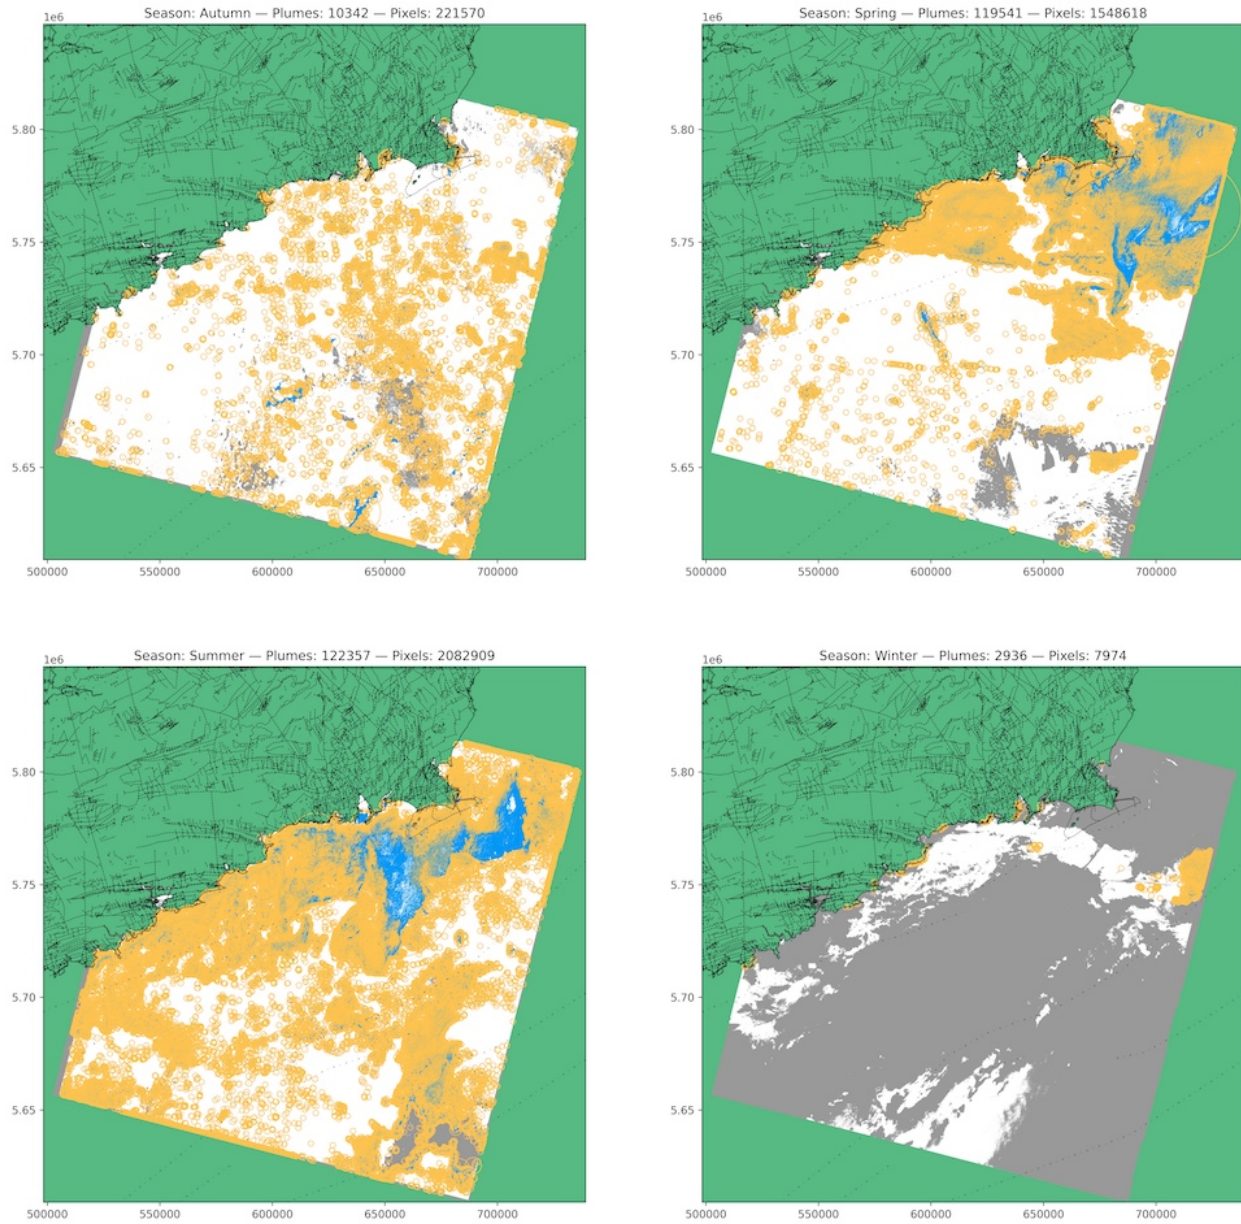


Figure B25: Seasonal PSGD detected by crosschecking derivative analysis and SST intervals (without plume buffering), but regardless of temporal consistency. PSGD plumes (blue pixels) are encircled (orange) to assist visualization. No data points are represented in gray. White pixels have no detectable PSGD. Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.

B.3.2 Angular Distance

B.3.2.1 Histograms




Figure B26: Relative frequency (%) of the number of times each pixel is flagged as PSGD by angular distance (with plume buffering), computed as a percent of the number of valid observations in the corresponding pixel.




Figure B27: Cumulative relative frequency (%) of the number of times each pixel is flagged as PSGD by angular distance (with plume buffering), computed as a percent of the number of valid observations in the corresponding pixel.




Figure B28: Relative frequency (%) of the number of times each pixel is flagged as PSGD, in each season, by angular distance (with plume buffering).




Figure B29: Relative frequency (%) of the number of times each pixel is flagged as PSGD, in each season, by angular distance (with plume buffering and crosschecked with SST), computed as a percent of the number of valid observations in the corresponding pixel.

B.3.2.2 Maps

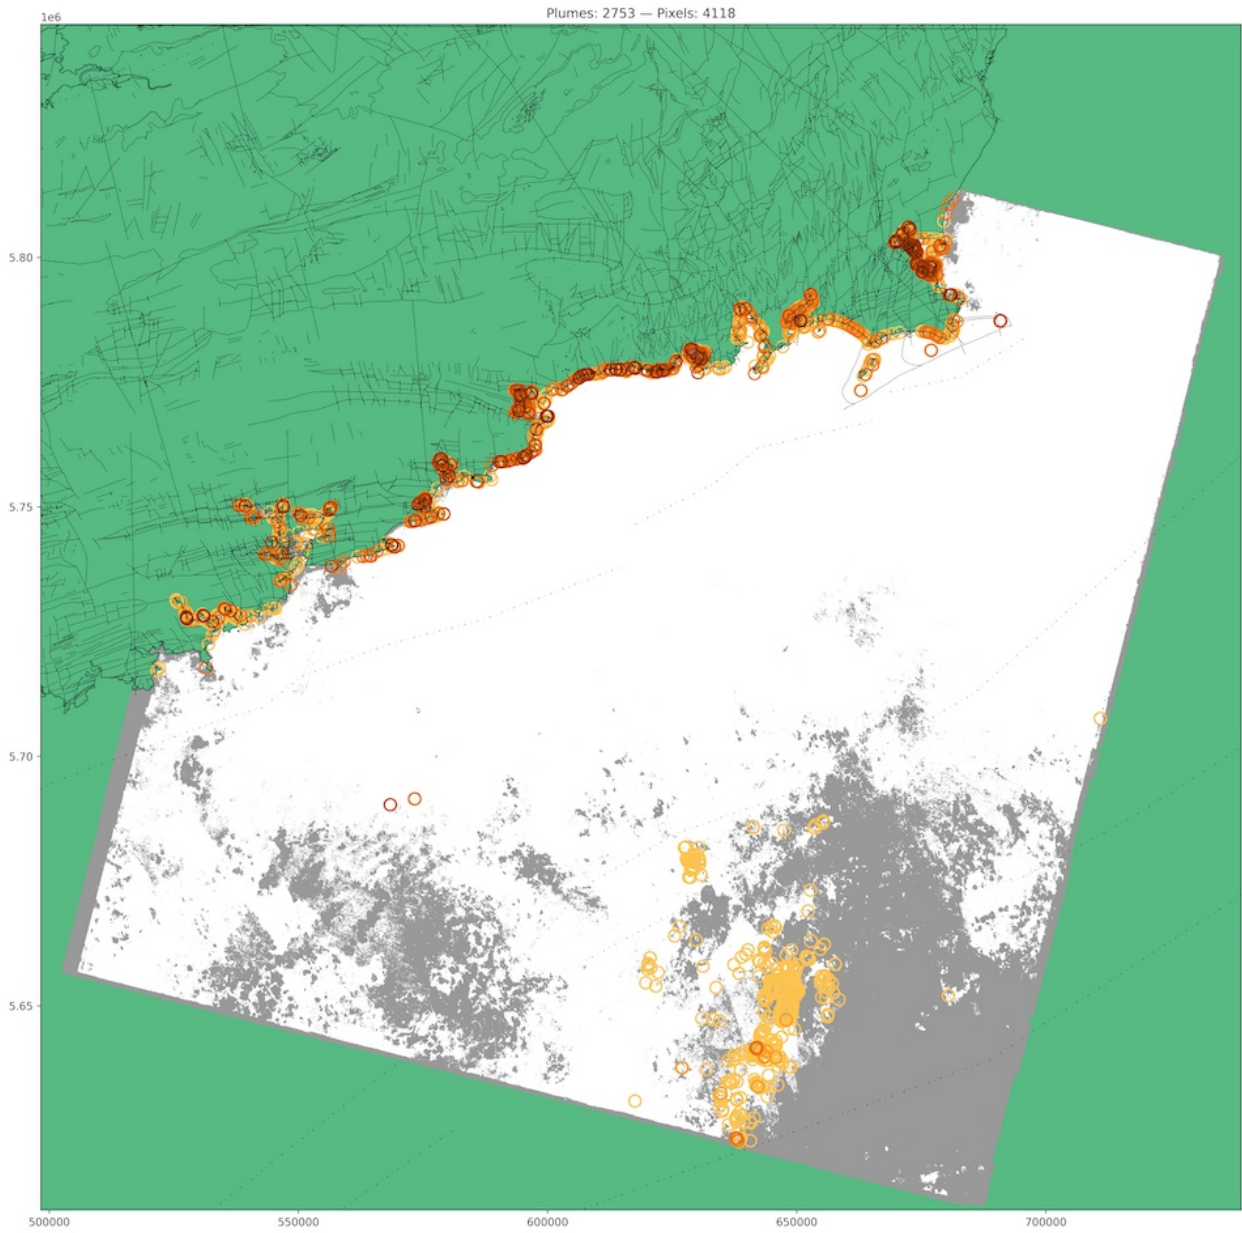


Figure B30: PSGD detected by angular distance (without plume buffering and without crosschecking), and with temporal consistency of more than 50% of the valid observations. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents five categories of temporal consistency: > 50% (lighter orange) to > 90% (darker orange). Pixels with less than 5 valid observations are represented in gray. White pixels have at least 5 valid observation but no detectable PSGD. Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.

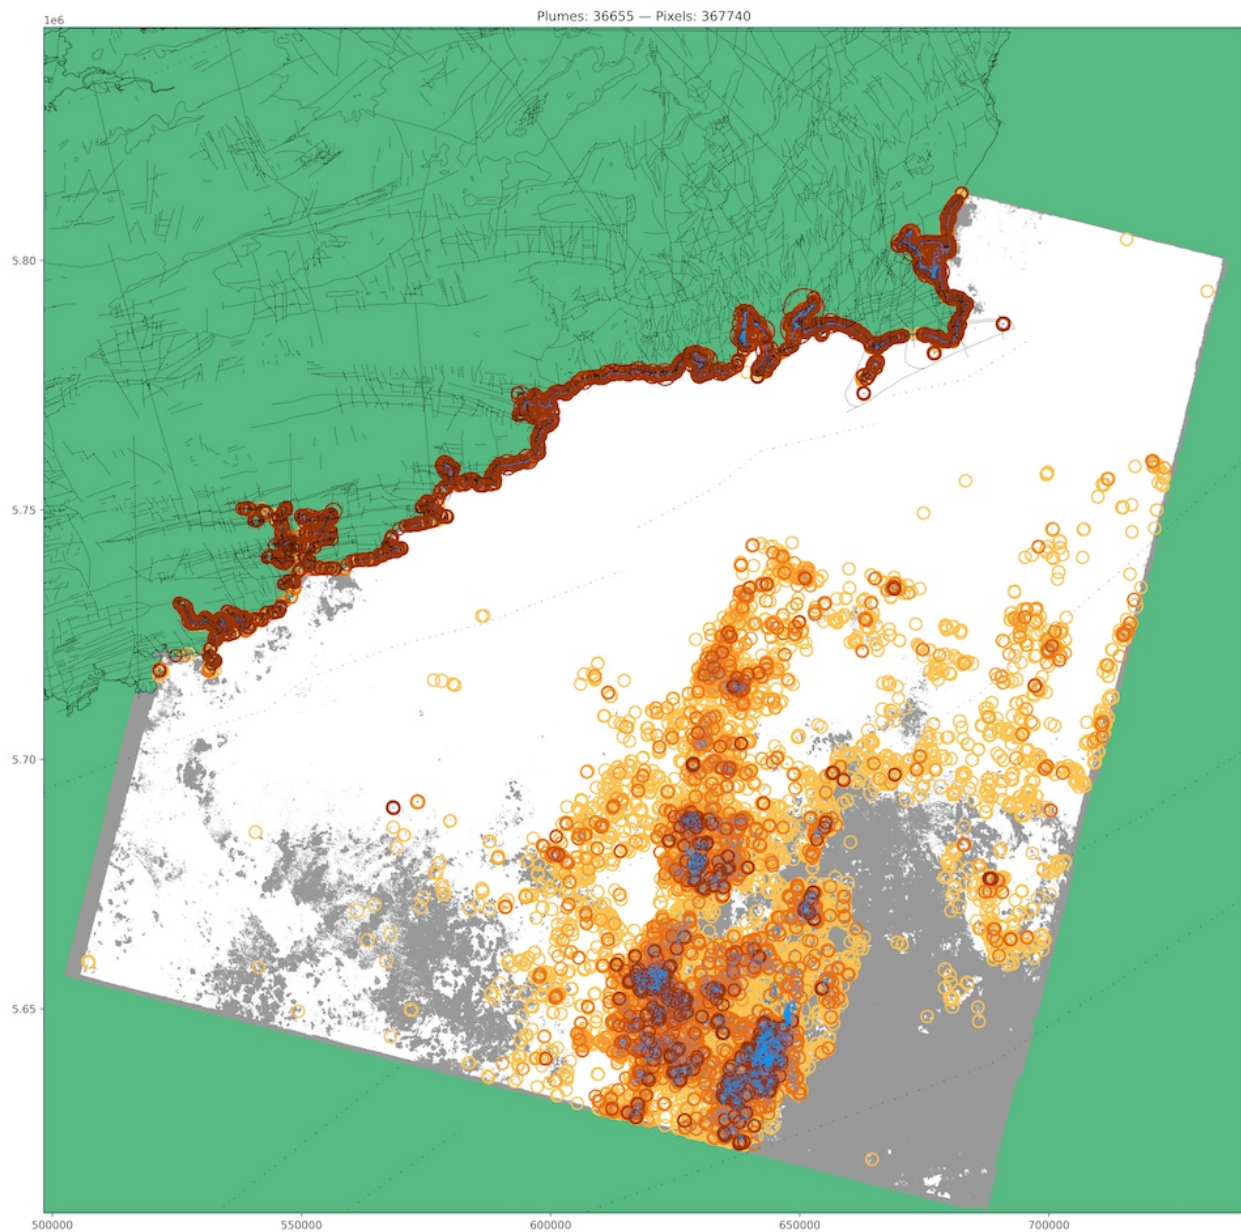


Figure B31: PSGD detected by angular distance (with plume buffering, but without crosschecking), and with temporal consistency of more than 50% of the valid observations. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents five categories of temporal consistency: > 50% (lighter orange) to > 90% (darker orange). Pixels with less than 5 valid observations are represented in gray. White pixels have at least 5 valid observation but no detectable PSGD. Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.




Figure B32: PSGD detected by crosschecking angular distance and SST intervals (with plume buffering), and with temporal consistency of more than 50% of the valid observations. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents five categories of temporal consistency: > 50% (lighter orange) to > 90% (darker orange). Pixels with less than 6 valid observations are represented in gray. White pixels have at least 6 valid observation but no detectable PSGD. Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.




Figure B33: Seasonal PSGD detected by angular distance (without plume buffering and without cross-checking), and with temporal consistency of more than 50% of the valid observations. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents five categories of temporal consistency: > 50% (lighter orange) to > 90% (darker orange). Pixels with less than 5 valid observations are represented in gray. White pixels have at least 5 valid observation but no detectable PSGD. Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.




Figure B34: Seasonal PSGD detected by crosschecking angular distance and SST intervals (with plume buffering), and with temporal consistency of more than 50% of the valid observations. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents five categories of temporal consistency: > 50% (lighter orange) to > 90% (darker orange). Pixels with less than 5 valid observations are represented in gray. White pixels have at least 5 valid observation but no detectable PSGD (no PSGD detected in autumn and winter). Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.



Figure B35: Seasonal PSGD detected by crosschecking angular distance and SST intervals (with plume buffering), and with temporal consistency of more than 50% of the valid observations. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents five categories of temporal consistency: > 50% (lighter orange) to > 90% (darker orange). Pixels with less than 6 valid observations are represented in gray. White pixels have at least 6 valid observation but no detectable PSGD (no PSGD detected in autumn and winter). Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.

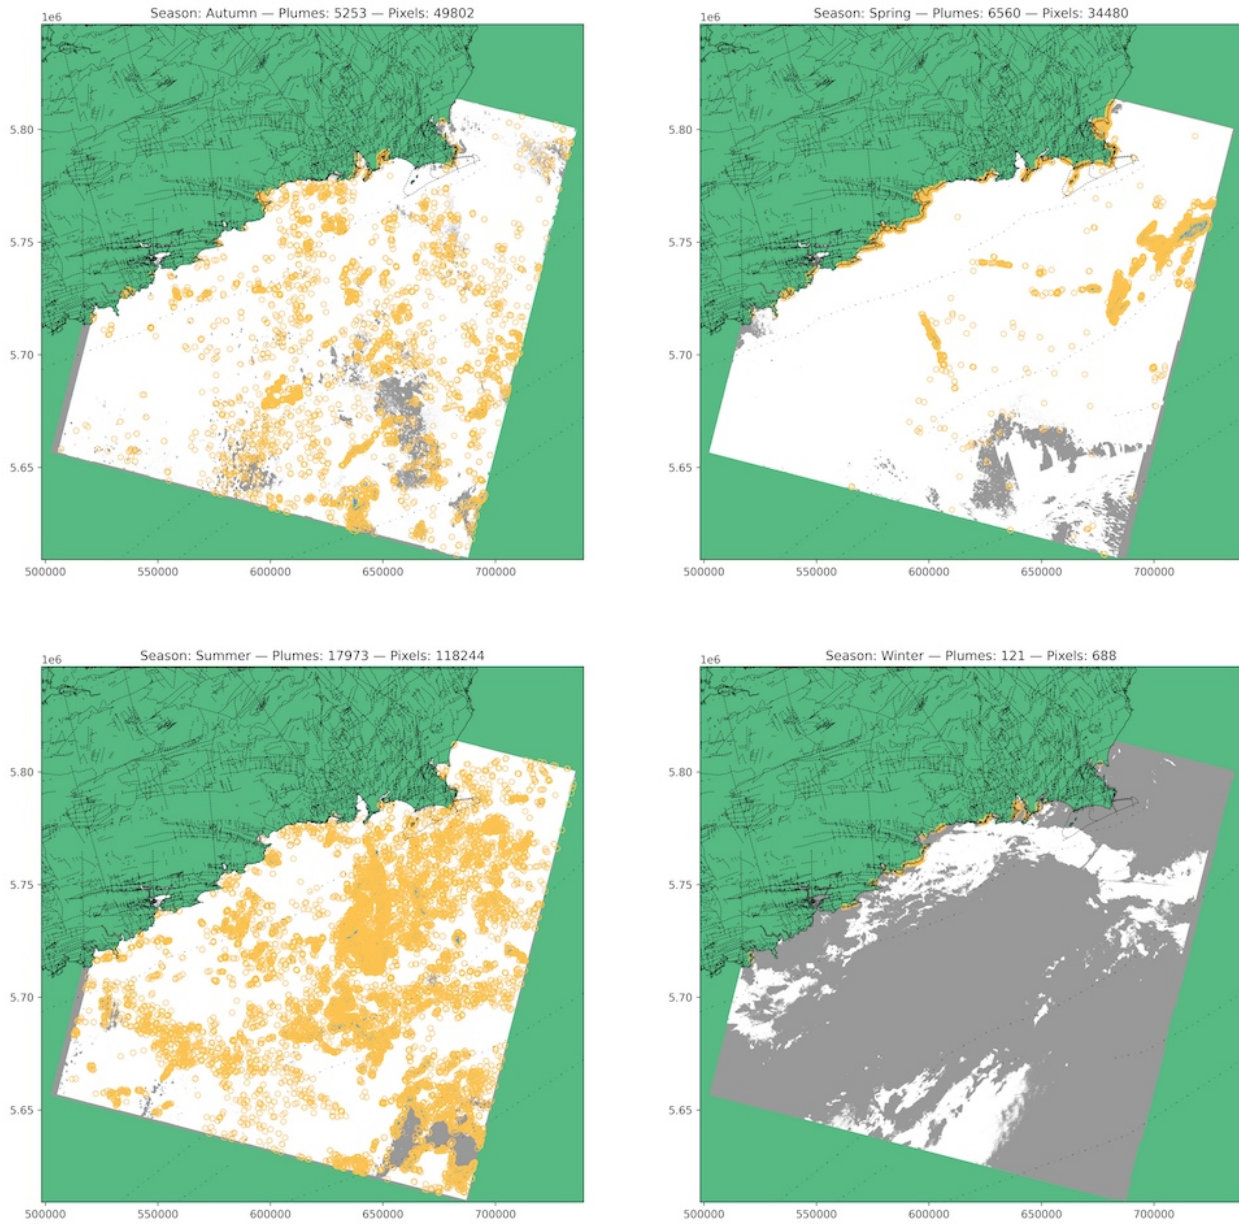


Figure B36: Seasonal PSGD detected by crosschecking angular distance and SST intervals (without plume buffering), but regardless of temporal consistency. PSGD plumes (blue pixels) are encircled (orange) to assist visualization. No data points are represented in gray. White pixels have no detectable PSGD. Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.

B.3.3 Derivative Analysis and Angular Distance

B.3.3.1 Histograms




Figure B37: Relative frequency (%) of the number of times each pixel is flagged as PSGD by crosschecking derivative analysis and angular distance (with plume buffering), computed as a percent of the number of valid observations in the corresponding pixel.




Figure B38: Cumulative relative frequency (%) of the number of times each pixel is flagged as PSGD by crosschecking derivative analysis and angular distance (with plume buffering), computed as a percent of the number of valid observations in the corresponding pixel.

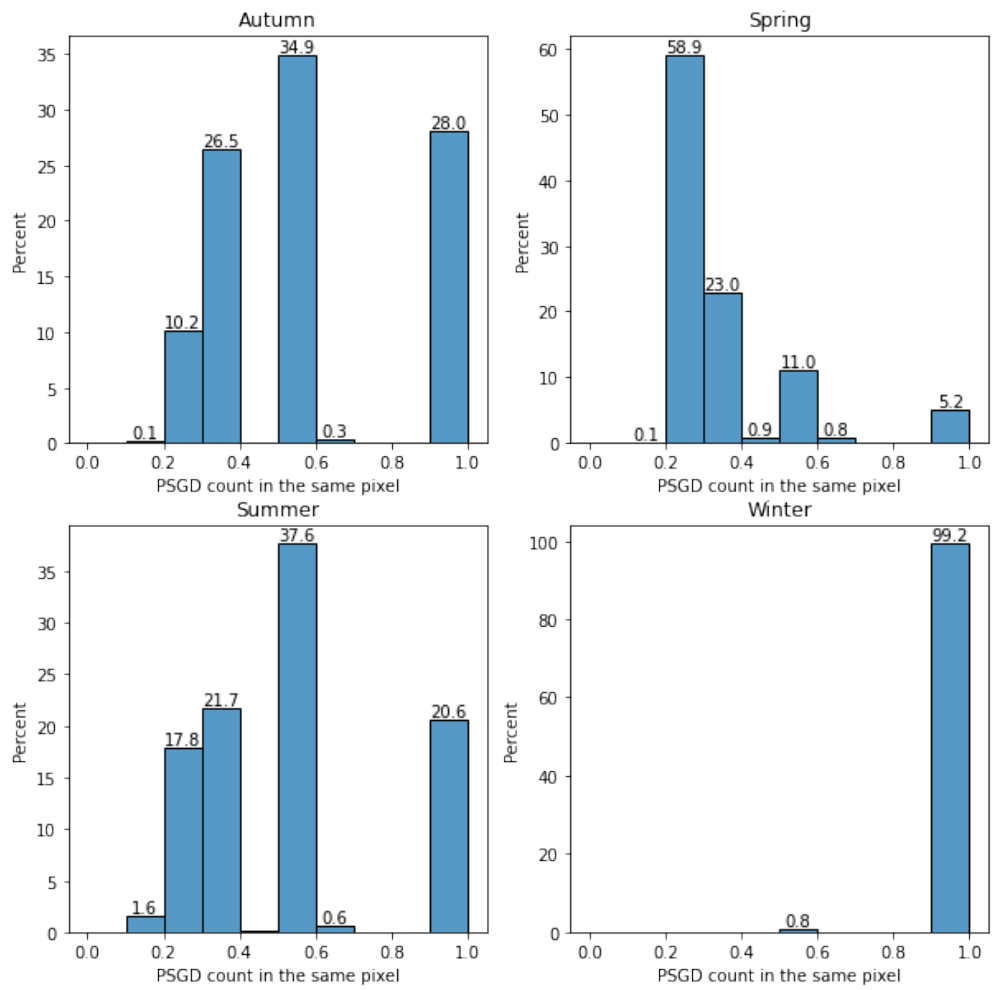


Figure B39: Relative frequency (%) of the number of times each pixel is flagged as PSGD, in each season, by crosschecking derivative analysis and angular distance (with plume buffering).

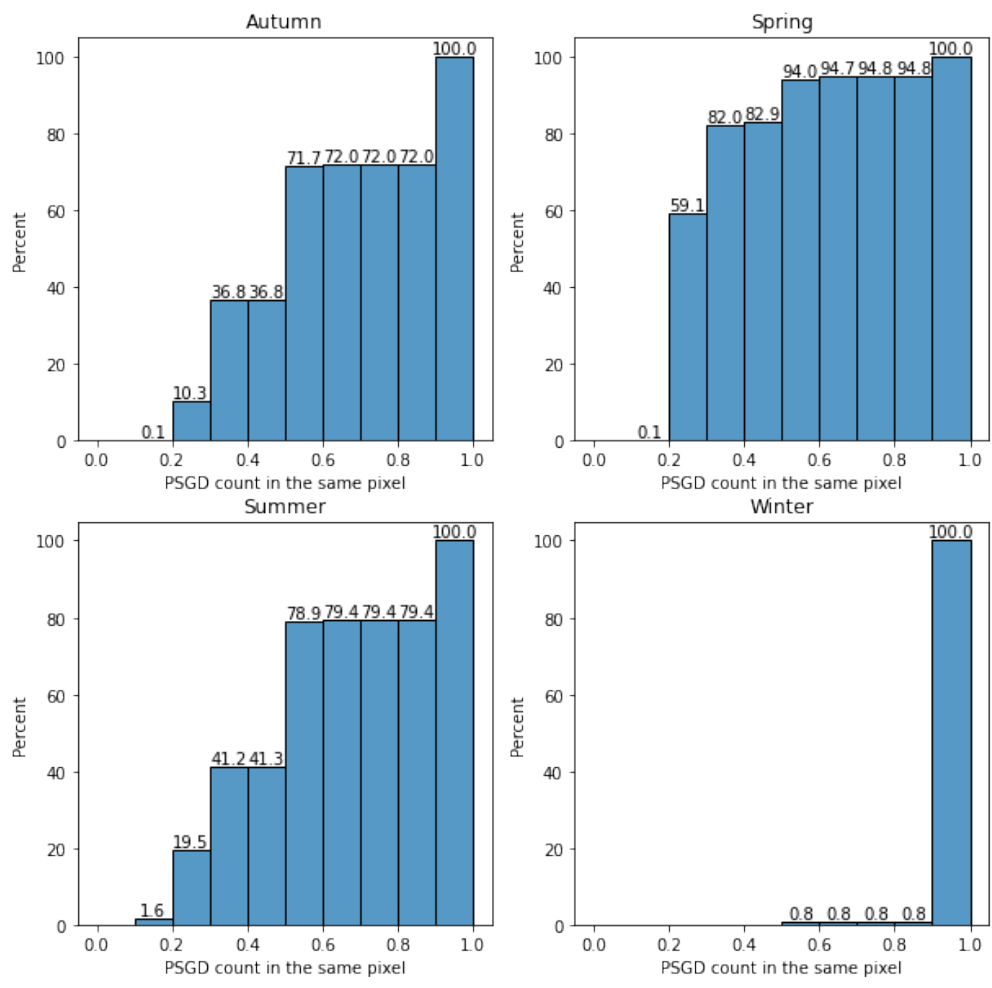


Figure B40: Relative frequency (%) of the number of times each pixel is flagged as PSGD, in each season, by crosschecking derivative analysis and angular distance (with plume buffering and crosschecked with SST), computed as a percent of the number of valid observations in the corresponding pixel.

B.3.3.2 Maps

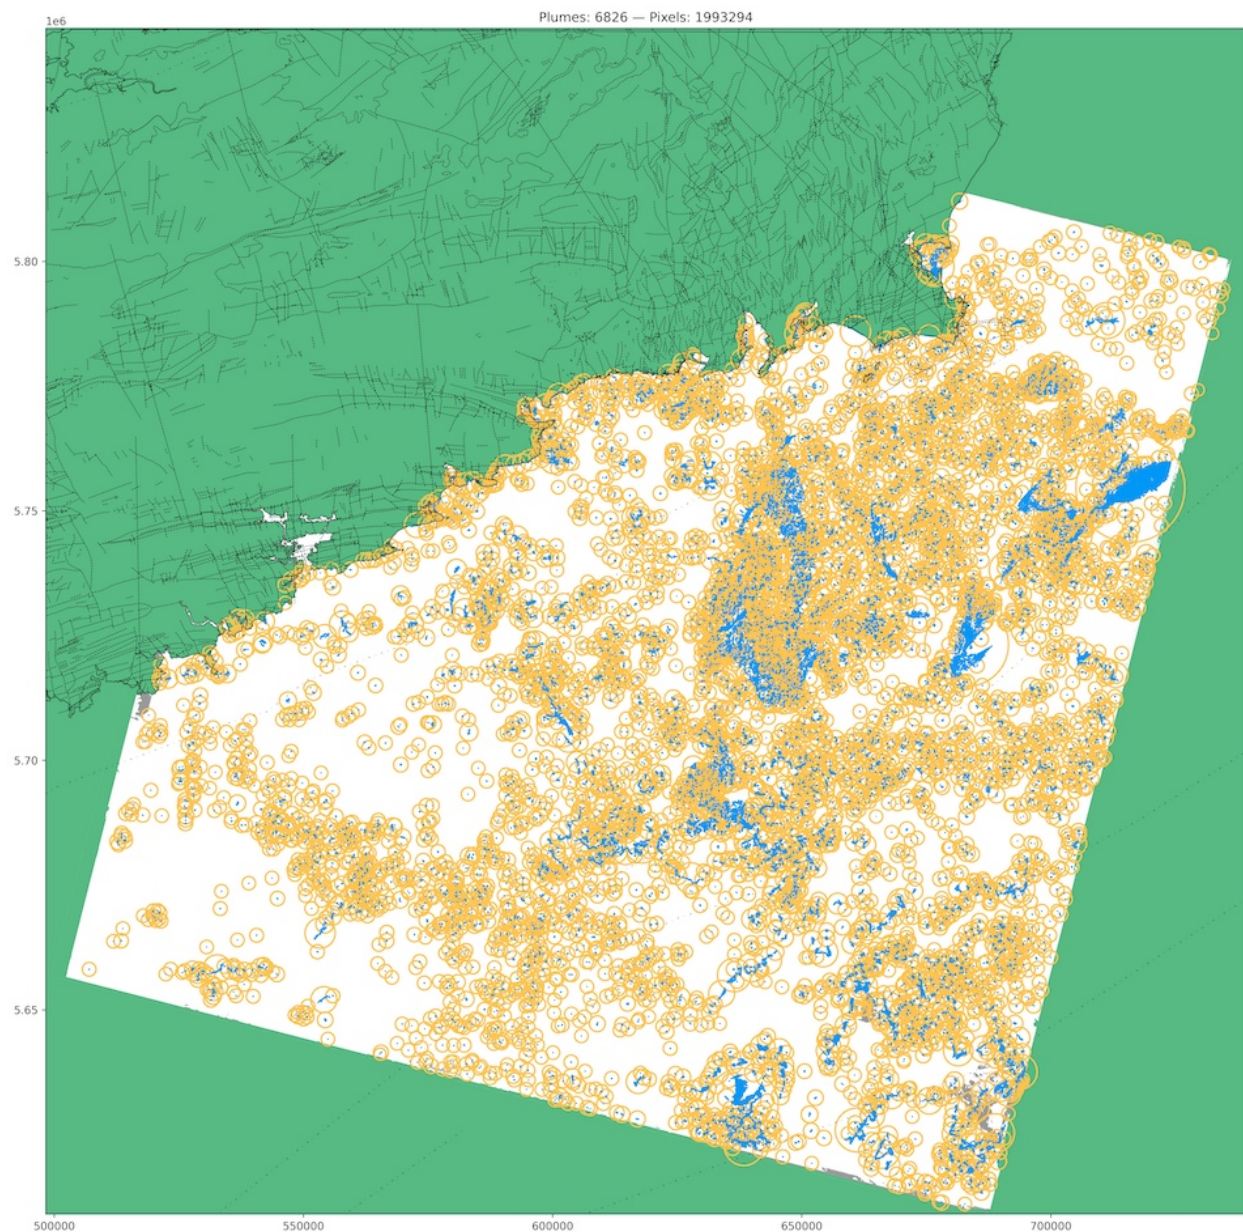


Figure B41: PSGD detected by crosschecking derivative analysis and angular distance (with plume buffering, without crosschecking with SST intervals), but regardless of temporal consistency. PSGD plumes (blue pixels) are encircled (orange) to assist visualization. No data points are represented in gray. White pixels have no detectable PSGD. Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.




Figure B42: PSGD detected by crosschecking derivative analysis and angular distance (with plume buffering, without crosschecking with SST intervals), and with temporal consistency of more than 50% of the valid observations. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents five categories of temporal consistency: > 50% (lighter orange) to > 90% (darker orange). Pixels with less than 6 valid observations are represented in gray. White pixels have at least 6 valid observation but no detectable PSGD. Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.




Figure B43: PSGD detected by crosschecking derivative analysis and angular distance (with plume buffering, without crosschecking with SST intervals), and with temporal consistency of more than 50% of the valid observations. The angular distance threshold is set at the 10th percentile, instead of the usual 5th percentile. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents five categories of temporal consistency: > 50% (lighter orange) to > 90% (darker orange). Pixels with less than 5 valid observations are represented in gray. White pixels have at least 5 valid observation but no detectable PSGD (no PSGD detected in autumn and winter). Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.




Figure B44: Seasonal PSGD detected by crosschecking derivative analysis and angular distance (with plume buffering, without crosschecking with SST intervals), and with temporal consistency of more than 50% of the valid observations. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents five categories of temporal consistency: > 50% (lighter orange) to > 90% (darker orange). Pixels with less than 6 valid observations are represented in gray. White pixels have at least 6 valid observation but no detectable PSGD (no PSGD detected in autumn and winter). Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.




Figure B45: PSGD detected by crosschecking derivative analysis, angular distance, and SST intervals, and with temporal consistency of more than 50% of the valid observations. Derivative analysis and angular distance are crosschecked with SST intervals separately, and only then the two are crosschecked together. This order of operations contrasts with the order applied elsewhere (derivative analysis and angular distance first, and then the result is crosschecked with SST intervals), showing minimal impact in the final result. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents five categories of temporal consistency: > 50% (lighter orange) to > 90% (darker orange). Pixels with less than 5 valid observations are represented in gray. White pixels have at least 5 valid observation but no detectable PSGD. Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.

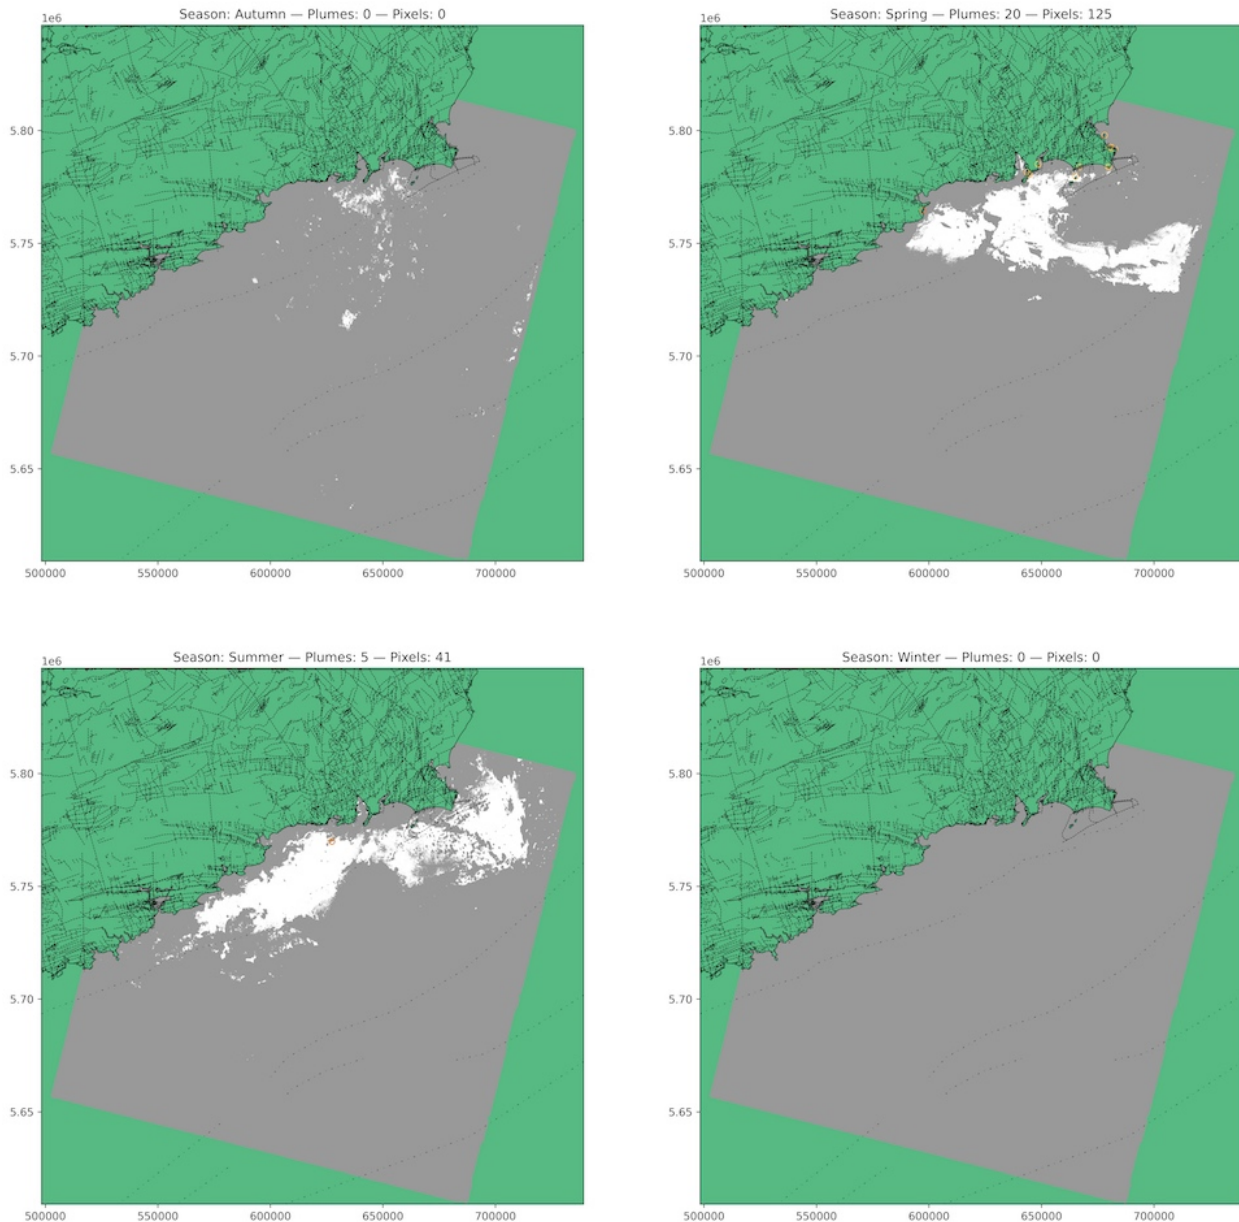


Figure B46: Seasonal PSGD detected by crosschecking derivative analysis, angular distance, and SST intervals, and with temporal consistency of more than 50% of the valid observations. Derivative analysis and angular distance are crosschecked with SST intervals separately, and only then the two are crosschecked together. This order of operations contrasts with the order applied elsewhere (derivative analysis and angular distance first, and then the result is crosschecked with SST intervals), showing minimal impact in the final result. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents five categories of temporal consistency: > 50% (lighter orange) to > 90% (darker orange). Pixels with less than 5 valid observations are represented in gray. White pixels have at least 5 valid observation but no detectable PSGD. Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.

B.3.4 Sea Surface Temperature

B.3.4.1 Maps




Figure B47: Mode of SST intervals from k -means clustering, where interval 1 is the coldest (dark blue) and interval 5 is the warmest (yellow). The pixel-wise mode shows the most commonly assigned interval for each pixel. White pixels represent no data points or masked out clouds. Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.

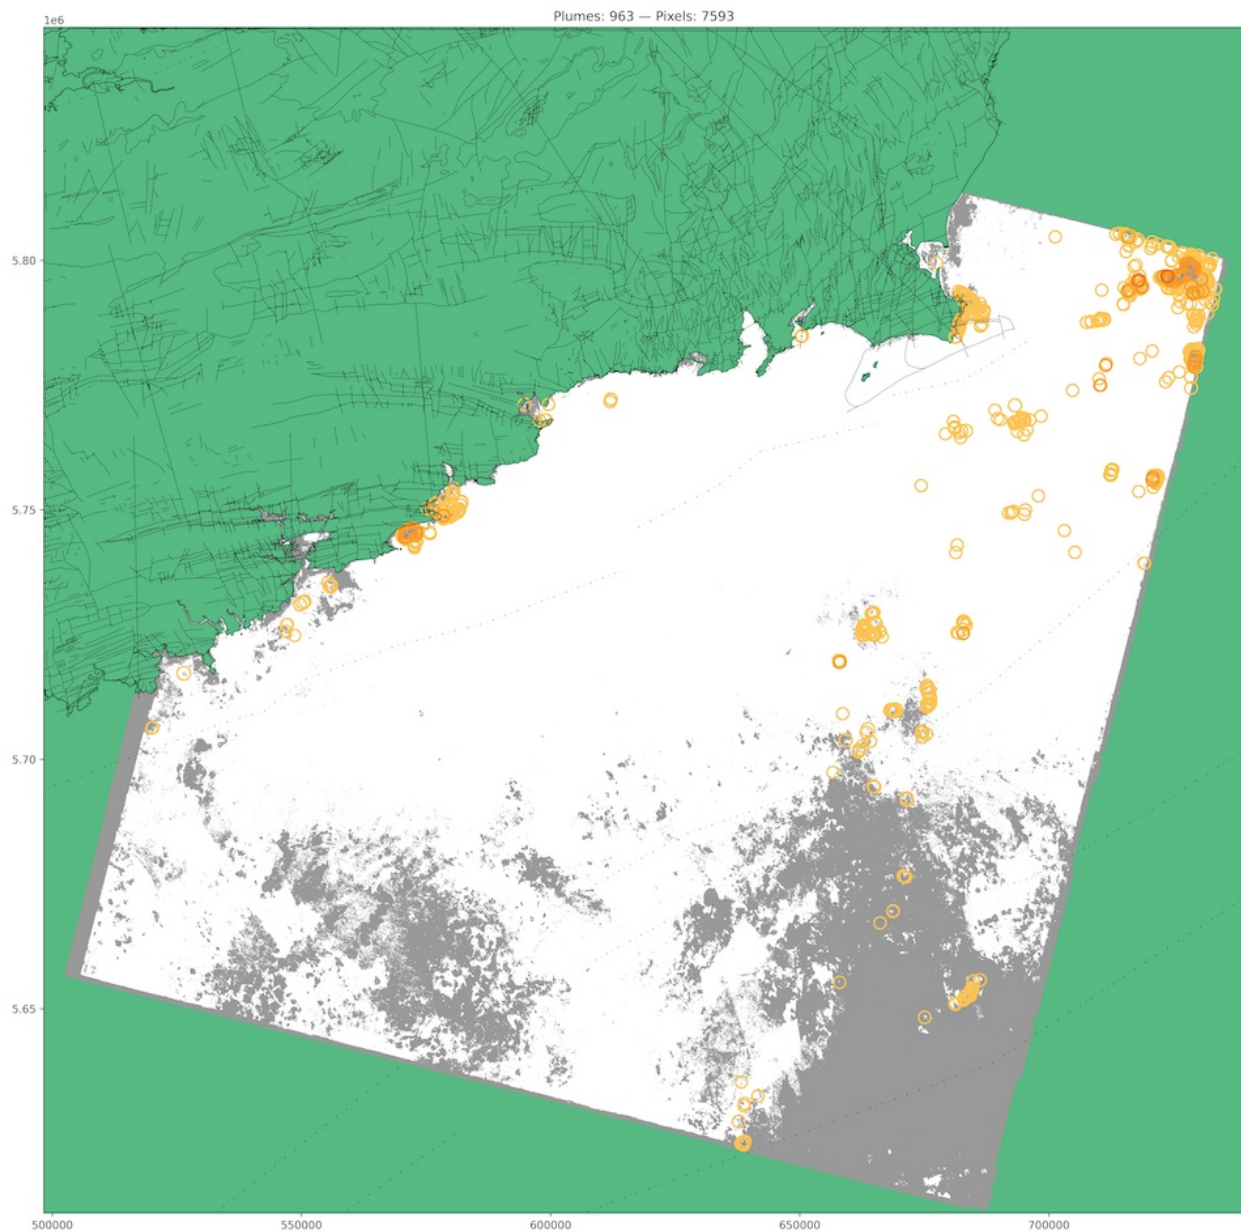


Figure B48: PSGD detected by SST intervals (without plume buffering and without crosschecking), and with temporal consistency of more than 50% of the valid observations. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents five categories of temporal consistency: > 50% (lighter orange) to > 90% (darker orange). Pixels with less than 5 valid observations are represented in gray. White pixels have at least 5 valid observation but no detectable PSGD. Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.




Figure B49: PSGD detected by SST intervals (without plume buffering and without crosschecking), and with temporal consistency of more than 50% of the valid observations. PSGD plumes (blue pixels) are encircled to assist visualization, where the color scale represents five categories of temporal consistency: > 50% (lighter orange) to > 90% (darker orange). Pixels with less than 5 valid observations are represented in gray. White pixels have at least 5 valid observation but no detectable PSGD. Dashed lines show bedrock faults [23]. Dotted dashed lines show offshore faults.




Figure B50: Maximum (warmer is higher) of the expected thermal signature of PSGD plumes in each season: cold plumes (green), warm plumes (yellow). The expected thermal signature is determined by comparing groundwater temperature to SST observations from MODIS. The thermal signature is undetermined when there are missing data in MODIS scenes, or when the gradient between groundwater temperature and SST is diminished (blue pixels).




Figure B51: Median of the expected thermal signature of PSGD plumes in each season: cold plumes (green), warm plumes (yellow). The expected thermal signature is determined by comparing groundwater temperature to SST observations from MODIS. The thermal signature is undetermined when there are missing data in MODIS scenes, or when the gradient between groundwater temperature and SST is diminished (blue pixels).

Bibliography

- [1] C. R. Aldwell and D. J. Burdon. Temperature of infiltration and groundwater. Conjunctive Water Use. In *Proc. Budapest Symp. July 1986. IAHS Publ. Number 156*, Budapest, 1986.
- [2] Alistair Allen and John Burgess. Developments in Geothermal Utilization in the Irish Republic. In *Proc. World Geotherm. Congr.*, number April, pages 25–29, 2010.
- [3] Theodoros Astaras and D. Oikonomidis. Remote sensing techniques to monitoring coastal plain areas suffering from salt water intrusion and detection of fresh water discharge in coastal, karstic areas: case studies from greece. In *Groundw. Ecosyst.*, pages 1–13. Kluwer Academic Publishers, Dordrecht, 2006.
- [4] Tor Haakon Bakken, Fridtjof Ruden, and Lars Erik Mangset. Submarine groundwater: A new concept for the supply of drinking water. *Water Resour. Manag.*, 26(4):1015–1026, 2012.
- [5] J.A. Barsi, J.L. Barker, and J.R. Schott. An Atmospheric Correction Parameter Calculator for a single thermal band earth-sensing instrument. In *IGARSS 2003. 2003 IEEE Int. Geosci. Remote Sens. Symp. Proc. (IEEE Cat. No.03CH37477)*, volume 5, pages 3014–3016, 2003.
- [6] Matthew W. Becker. Potential for satellite remote sensing of ground water. *Ground Water*, 44(2):306–318, 2006.
- [7] Susanne A Benz, Peter Bayer, and Philipp Blum. Global patterns of shallow groundwater temperatures. *Environ. Res. Lett.*, 12(3):034005, mar 2017.
- [8] Wilfried Brutsaert and John L. Nieber. Regionalized drought flow hydrographs from a mature glaciated plateau. *Water Resour. Res.*, 13(3):637–643, 1977.
- [9] W. C. Burnett, P. K. Aggarwal, A. Aureli, H. Bokuniewicz, J. E. Cable, M. A. Charette, E. Kontar, S. Krupa, K. M. Kulkarni, A. Loveless, W. S. Moore, J. A. Oberdorfer, J. Oliveira, N. Ozyurt, P. Povinec, A. M G Privitera, R. Rajar, R. T. Ramessur, J. Scholten, T. Stieglitz, M. Taniguchi, and J. V. Turner. Quantifying submarine groundwater discharge in the coastal zone via multiple methods. *Sci. Total Environ.*, 367(2-3):498–543, aug 2006.

- [10] W C Burnett, H Bokuniewicz, Markus Huettel, W Moore, and M Taniguchi. Groundwater and pore water inputs to the coastal zone. *Biogeochemistry*, 66:3–33, 2003.
- [11] Jaye E. Cable, William C. Burnett, Jeffrey P. Chanton, and Georges L. Weatherly. Estimating groundwater discharge into the northeastern Gulf of Mexico using radon-222. *Earth Planet. Sci. Lett.*, 144(3-4):591–604, 1996.
- [12] Júlio Caineta, Brian F Thomas, and Daniel J Bain. Submarine groundwater discharge detection through remote sensing: An application of Landsat 7 and 8 in Hawai'i and Ireland. *Remote Sens. Environ.*, 279(May):113109, sep 2022.
- [13] François Chollet. *Deep Learning with Python*. Manning Publications Co., Shelter Island, NY, 2018.
- [14] Denis Cohen, Mark Person, Peng Wang, Carl W. Gable, Deborah Hutchinson, Andee Marksamer, Brandon Dugan, Henk Kooi, Koos Groen, Daniel Lizarralde, Robert L. Evans, Frederick D. Day-Lewis, and John W. Lane. Origin and Extent of Fresh Paleowaters on the Atlantic Continental Shelf, USA. *Ground Water*, 48(1):143–158, jan 2010.
- [15] Claudia Corradino, Giuseppe Bilotta, Annalisa Cappello, Luigi Fortuna, and Ciro Del Negro. Combining radar and optical satellite imagery with machine learning to map lava flows at mount etna and fogo island. *Energies*, 14(1), 2021.
- [16] Emilio Custodio. Coastal aquifers of Europe: an overview. *Hydrogeol. J.*, 18(1):269–280, feb 2010.
- [17] C. F. D' Elia, K L Webb, and J W Porter. Nitrate-rich groundwater inputs to Discovery Bay, Jamaica: a significant source of N to local coral reefs? *Bull. Mar. Sci.*, 31(4):903–910, 1981.
- [18] Richard S. Deitchman and Steven P. Loheide. Ground-based thermal imaging of groundwater flow processes at the seepage face. *Geophys. Res. Lett.*, 36(14):L14401, jul 2009.
- [19] James Dooge. *Linear theory of hydrologic systems*. Agricultural Research Service, US Department of Agriculture, 1468, 1973.
- [20] Grant Ferguson and Tom Gleeson. Vulnerability of coastal aquifers to groundwater use and climate change. *Nat. Clim. Chang.*, 2(5):342–345, may 2012.

- [21] Food and Agriculture Organization of the United Nations. Seawater intrusion in coastal aquifers, guidelines for study, monitoring and control. Technical report, United Nations, Rome, Italy, 1997.
- [22] Alain P. Francés, Elsa C. Ramalho, Judite Fernandes, Michel Groen, Rui Hugman, Mohamed A. Khalil, Joel De Plaen, and Fernando A. Monteiro Santos. Contributions of hydrogeophysics to the hydrogeological conceptual model of the Albufeira-Ribeira de Quarteira coastal aquifer in Algarve, Portugal. *Hydrogeol. J.*, 23(7):1553–1572, nov 2015.
- [23] Geological Survey Ireland. GSI Bedrock Geology 100k Series, 2016.
- [24] Douglas G. Goodin, Luoheng Han Han, Rolland N. Fraser, Donald C. Rundquist, and Wesley A. Stebbins. Analysis of Suspended-Solids in Water Using Remotely Sensed High-Resolution Derivative Spectra. *Photogramm. Eng. Remote Sensing*, 59(4):505–510, 1993.
- [25] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.*, 12(7):2217–2226, 2019.
- [26] Scott Jasechko, S. Jean Birks, Tom Gleeson, Yoshihide Wada, Peter J. Fawcett, Zachary D. Sharp, Jeffrey J. McDonnell, and Jeffrey M. Welker. The pronounced seasonality of global groundwater recharge. *Water Resour. Res.*, 50(11):8845–8867, 2014.
- [27] Adam G. Johnson, Craig R. Glenn, William C. Burnett, Richard N. Peterson, and Paul G. Lucey. Aerial infrared imaging reveals large nutrient-rich groundwater inputs to the ocean. *Geophys. Res. Lett.*, 35(15):L15606, aug 2008.
- [28] Sònia Jou-Claus, Albert Folch, and Jordi Garcia-Orellana. Applicability of Landsat 8 thermal infrared sensor for identifying submarine groundwater discharge springs in the Mediterranean Sea basin. *Hydrol. Earth Syst. Sci.*, 25(9):4789–4805, 2021.
- [29] James O. Juvik, D. C. Singleton, and G. G. Clarke. Climate and Water Balance on the Island of Hawaii. In *Water Encycl.* John Wiley & Sons, Inc., Hoboken, NJ, USA, jul 2005.
- [30] Gholam A. Kazemi. Editor’s Message: Submarine groundwater discharge studies and the absence of hydrogeologists. *Hydrogeol. J.*, 16(2):201–204, 2008.

- [31] Jacque L. Kelly, Craig R. Glenn, and Paul G. Lucey. High-resolution aerial infrared mapping of groundwater discharge to the coastal ocean. *Limnol. Oceanogr. Methods*, 11(5):262–277, may 2013.
- [32] G. Kiely. Climate change in Ireland from precipitation and streamflow observations. *Adv. Water Resour.*, 23(2):141–151, 1999.
- [33] J. M. Krest, W. S. Moore, L. R. Gardner, and J. T. Morris. Marsh nutrient export supplied by groundwater discharge: Evidence from radium measurements. *Global Biogeochem. Cycles*, 14(1):167–176, mar 2000.
- [34] Alex Krizhevsky, Ilya Sutskever, and Hinton Geoffrey E. ImageNet Classification with Deep Convolutional Neural Networks. *Adv. Neural Inf. Process. Syst.* 25, pages 1–9, 2012.
- [35] F. A. Kruse, A. B. Lefkoff, J. W. Boardman, K. B. Heidebrecht, A. T. Shapiro, P. J. Barloon, and A. F.H. Goetz. The spectral image processing system (SIPS)-interactive visualization and analysis of imaging spectrometer data. *Remote Sens. Environ.*, 44(2-3):145–163, 1993.
- [36] Eun Young Kwon, Guebuem Kim, Francois Primeau, Willard S. Moore, Hyung Mi Cho, Timothy Devries, Jorge L. Sarmiento, Matthew A. Charette, and Yang Ki Cho. Global estimate of submarine groundwater discharge based on an observationally constrained radium isotope model. *Geophys. Res. Lett.*, 41(23):8438–8444, dec 2014.
- [37] Landsat Project Science Office. Landsat 7 Data Users Handbook. Technical Report November, USGS, Sioux Falls, SD, 2019.
- [38] Julie Laroche, Robert Nuzzi, Robert Waters, Kevin Wyman, Paul Falkowski, and Douglas Wallace. Brown Tide blooms in Long Island’s coastal waters linked to inter-annual variability in groundwater flow. *Glob. Chang. Biol.*, 3(5):397–410, oct 1997.
- [39] David J. Lary, Amir H. Alavi, Amir H. Gandomi, and Annette L. Walker. Machine learning in geosciences and remote sensing. *Geosci. Front.*, 7(1):3–10, jan 2016.
- [40] L.S.a Lau and J.F.b Mink. *Hydrology of the Hawaiian Islands*. 2006.
- [41] P. Lazzari, C. Solidoro, V. Ibello, S. Salon, A. Teruzzi, K. Béranger, S. Colella, and A. Crise. Seasonal and inter-annual variability of plankton chlorophyll and pri-

- mary production in the Mediterranean Sea: A modelling approach. *Biogeosciences*, 9(1):217–233, 2012.
- [42] Lukas W. Lehnert, Hanna Meyer, Wolfgang A. Obermeier, Brenner Silva, Bianca Regeling, and Jörg Bendix. Hyperspectral Data Analysis in R : The hsdar Package. *J. Stat. Softw.*, 89(12), 2019.
- [43] L. Li, D. A. Barry, F. Stagnitti, and J.-Y. Parlange. Submarine groundwater discharge and associated chemical input to a coastal sea. *Water Resour. Res.*, 35(11):3253–3259, nov 1999.
- [44] Scott M Lundberg and Su-In Lee. A Unified Approach to Interpreting Model Predictions. In I Guyon, U Von Luxburg, S Bengio, H Wallach, R Fergus, S Vishwanathan, and R Garnett, editors, *Adv. Neural Inf. Process. Syst.*, volume 31, Long Beach, CA, USA, 2017. Curran Associates, Inc.
- [45] U. Mallast, R. Gloaguen, J. Friesen, T. Rödiger, S. Geyer, R. Merz, and C. Siebert. How to identify groundwater-caused thermal anomalies in lakes based on multi-temporal satellite data in semi-arid regions. *Hydrol. Earth Syst. Sci.*, 18(7):2773–2787, jul 2014.
- [46] U. Mallast, C. Siebert, B. Wagner, M. Sauter, R. Gloaguen, S. Geyer, and R. Merz. Localisation and temporal variability of groundwater discharge into the Dead Sea using thermal satellite data. *Environ. Earth Sci.*, 69(2):587–603, may 2013.
- [47] Ulf Mallast, Friedhelm Schwonke, Richard Gloaguen, Stefan Geyer, Martin Sauter, and Christian Siebert. Airborne Thermal Data Identifies Groundwater Discharge at the North-Western Coast of the Dead Sea. *Remote Sens.*, 5(12):6361–6381, 2013.
- [48] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. {TensorFlow}: Large-Scale Machine Learning on Heterogeneous Systems, 2015.

- [49] Margaret McCaul, Jack Barland, John Cleary, Conor Cahalane, Tim McCarthy, and Dermot Diamond. Combining Remote Temperature Sensing with in-Situ Sensing to Track Marine/Freshwater Mixing Dynamics. *Sensors*, 16(9):1402, aug 2016.
- [50] Aaron Micallef, Mark Person, Christian Berndt, Claudia Bertoni, Denis Cohen, Brandon Dugan, Rob Evans, Amir Haroon, Christian Hensen, Marion Jegen, Kerry Key, Henk Kooi, Volker Liebetrau, Johanna Lofi, Brian J. Mailloux, Renée Martin-Nagle, Holly A. Michael, Thomas Müller, Mark Schmidt, Katrin Schwalenberg, Elizabeth Trembath-Reichert, Bradley Weymer, Yipeng Zhang, and Ariel T. Thomas. Offshore Freshened Groundwater in Continental Margins. *Rev. Geophys.*, 59(1):0–3, mar 2021.
- [51] Christoph Molnar. *Interpretable Machine Learning*. Independently published, 2022.
- [52] Willard S. Moore. The Effect of Submarine Groundwater Discharge on the Ocean. *Ann. Rev. Mar. Sci.*, 2(1):59–88, jan 2010.
- [53] N. Moosdorf and T. Oehler. Societal use of fresh submarine groundwater discharge: An overlooked water resource. *Earth-Science Rev.*, 171:338–348, aug 2017.
- [54] Nils Moosdorf, Thomas Stieglitz, Hannelore Waska, Hans H. Dürr, and Jens Hartmann. Submarine groundwater discharge from tropical islands: a review. *Grundwasser*, 20(1):53–67, 2015.
- [55] June A. Oberdorfer. Hydrogeologic modeling of submarine groundwater discharge: Comparison to other quantitative methods. *Biogeochemistry*, 66(1-2):159–169, 2003.
- [56] John E. O’Reilly. SeaWiFS postlaunch calibration and validation analyses, part 3. Technical report, 2000.
- [57] R. P. Pech, R. D. Graetz, and A. W. Davis. Reflectance modelling and the derivation of vegetation indices for an Australian semi-arid shrubland. *Int. J. Remote Sens.*, 7(3):389–403, 1986.
- [58] Paul R. Pinet. *Invitation to Oceanography*. Jones & Bartlett Learning, Burlington, MA, 7 edition, 2016.
- [59] V. E. A. Post. Fresh and saline groundwater interaction in coastal aquifers: Is our technology ready for the problems ahead? *Hydrogeol. J.*, 13(1):120–123, mar 2005.

- [60] Vincent E.A. Post, Jacobus Groen, Henk Kooi, Mark Person, Shemin Ge, and W. Mike Edmunds. Offshore fresh groundwater reserves as a global phenomenon. *Nature*, 504(7478):71–78, dec 2013.
- [61] J. T. Reager, A. S. Gardner, J. S. Famiglietti, D. N. Wiese, A. Eicker, and M.-H. Lo. A decade of sea level rise slowed by climate-driven hydrology. *Science (80-.)*., 351(6274):699–703, 2016.
- [62] M. Rodell, H. K. Beaudoin, T. S. L’Ecuyer, W. S. Olson, J. S. Famiglietti, P. R. Houser, R. Adler, M. G. Bosilovich, C. A. Clayson, D. Chambers, E. Clark, E. J. Fetzer, X. Gao, G. Gu, K. Hilburn, G. J. Huffman, D. P. Lettenmaier, W. T. Liu, F. R. Robertson, C. A. Schlosser, J. Sheffield, and E. F. Wood. The Observed State of the Water Cycle in the Early Twenty-First Century. *J. Clim.*, 28(21):8289–8318, nov 2015.
- [63] D. P. Roy, M. A. Wulder, T. R. Loveland, Woodcock C.E., R. G. Allen, M. C. Anderson, D. Helder, J. R. Irons, D. M. Johnson, R. Kennedy, T. A. Scambos, C. B. Schaaf, J. R. Schott, Y. Sheng, E. F. Vermote, A. S. Belward, R. Bindschadler, W. B. Cohen, F. Gao, J. D. Hipple, P. Hostert, J. Huntington, C. O. Justice, A. Kilic, V. Kovalskyy, Z. P. Lee, L. Lymburner, J. G. Masek, J. McCorkel, Y. Shuai, R. Trezza, J. Vogelmann, R. H. Wynne, and Z. Zhu. Landsat-8: Science and product vision for terrestrial global change research. *Remote Sens. Environ.*, 145:154–172, 2014.
- [64] Isaac R. Santos, Xiaogang Chen, Alanna L. Lecher, Audrey H. Sawyer, Nils Moosdorf, Valentí Rodellas, Joseph Tamborski, Hyung-Mi Cho, Natasha Dimova, Ryo Sugimoto, Stefano Bonaglia, Hailong Li, Mithra-Christin Hajati, and Ling Li. Submarine groundwater discharge impacts on coastal nutrient biogeochemistry. *Nat. Rev. Earth Environ.*, 2(5):307–323, may 2021.
- [65] Abraham. Savitzky and Marcel J. E. Golay. Smoothing and Differentiation of Data by Simplified Least Squares Procedures. *Anal. Chem.*, 36(8):1627–1639, jul 1964.
- [66] Audrey H. Sawyer, Cédric H. David, and James S. Famiglietti. Continental patterns of submarine groundwater discharge reveal coastal vulnerabilities. *Science (80-.)*., 353(6300):705–707, 2016.
- [67] Amin Shaban, Mohamad Khawlie, Chadi Abdallah, and Ghaleb Faour. Geologic controls of submarine groundwater discharge: Application of remote sensing to north Lebanon. *Environ. Geol.*, 47(4):512–522, 2005.

- [68] By David R Sherrod, John M Sinton, Sarah E Watkins, Kelly M Brunt, and U S Geological Survey. Geologic Map of the State of Hawai'i. *USGS Open File Rep.*, (2007-1089), 2007.
- [69] A.C.F. Silva, P Tavares, M Shapouri, T.Y. Stigter, J.P. Monteiro, M Machado, L. Cancela da Fonseca, and L Ribeiro. Estuarine biodiversity as an indicator of groundwater discharge. *Estuar. Coast. Shelf Sci.*, 97:38–43, jan 2012.
- [70] J. H. Simpson and J. R. Hunter. Fronts in the Irish Sea. *Nature*, 250, 1974.
- [71] Anthony J. Smith. Mixed convection and density-dependent seawater circulation in coastal aquifers. *Water Resour. Res.*, 40(8):66–73, aug 2004.
- [72] Ryo Sugimoto, Katsuhiko Kitagawa, Saori Nishi, Hisami Honda, Makoto Yamada, Shiho Kobayashi, Jun Shoji, Shinji Ohsawa, Makoto Taniguchi, and Osamu Tomimaga. Phytoplankton primary productivity around submarine groundwater discharge in nearshore coasts. *Mar. Ecol. Prog. Ser.*, 563:25–33, 2017.
- [73] Philip Sura, Matthew Newman, and Michael A. Alexander. Daily to decadal sea surface temperature variability driven by state-dependent stochastic heat fluxes. *J. Phys. Oceanogr.*, 36(10):1940–1958, 2006.
- [74] Philip Sura and Prashant D. Sardeshmukh. A global view of non-Gaussian SST variability. *J. Phys. Oceanogr.*, 38(3):639–647, 2008.
- [75] P W Swarzenski, J F Bratton, and J Crusius. Submarine ground-water discharge and its role in coastal processes and ecosystems. Technical Report 2004-1226, US Geological Survey, Reston, Virginia, 2004.
- [76] Tajdarul H Syed, James S Famiglietti, Don P Chambers, Josh K Willis, and Kyle Hilburn. Satellite-based global-ocean mass balance estimates of interannual variability and emerging trends in continental freshwater discharge. *Proc. Natl. Acad. Sci.*, 107(42):17916–17921, oct 2010.
- [77] Joseph J. Tamborski, A. Deanne Rogers, Henry J. Bokuniewicz, J. Kirk Cochran, and Caitlin R. Young. Identification and quantification of diffuse fresh submarine groundwater discharge via airborne thermal infrared remote sensing. *Remote Sens. Environ.*, 171:202–217, dec 2015.

- [78] Pierre Tandeo, Emmanuelle Autret, Bertrand Chapron, Ronan Fablet, and René Garello. SST spatial anisotropic covariances from METOP-AVHRR data. *Remote Sens. Environ.*, 141:144–148, 2014.
- [79] Makoto Taniguchi, William C. Burnett, Jaye E. Cable, and Jeffrey V. Turner. Investigation of submarine groundwater discharge. *Hydrol. Process.*, 16(11):2115–2129, aug 2002.
- [80] Makoto Taniguchi, Henrietta Dulai, Kimberly M. Burnett, Isaac R. Santos, Ryo Sugimoto, Thomas Stieglitz, Guebuem Kim, Nils Moosdorf, and William C. Burnett. Submarine Groundwater Discharge: Updates on Its Measurement Techniques, Geophysical Drivers, Magnitudes, and Effects. *Front. Environ. Sci.*, 7, oct 2019.
- [81] Gordon T. Taylor, Christopher J. Gobler, and Sergio A. Sañudo-Wilhelmy. Speciation and concentrations of dissolved nitrogen as determinants of brown tide *Aureococcus anophagefferens* bloom initiation. *Mar. Ecol. Prog. Ser.*, 312:67–83, 2006.
- [82] Julia Uitz, Hervé Claustre, Bernard Gentili, and Dariusz Stramski. Phytoplankton class-specific primary production in the world’s oceans: Seasonal and interannual variability from satellite observations. *Global Biogeochem. Cycles*, 24(3):1–19, 2010.
- [83] Sunil Varma, Jeffrey Turner, and Jim Underschultz. Estimation of submarine groundwater discharge into Geographe Bay, Western Australia. *J. Geochemical Explor.*, 101(1):107, apr 2009.
- [84] Sunil Varma, Jeffrey V Turner, and James R Underschultz. Estimation of Submarine Groundwater Discharge into Geographe Bay, Western Australia. Technical report, Commonwealth Scientific and Industrial Research Organisation, 2008.
- [85] Yoshihide Wada, John T. Reager, Benjamin F. Chao, Jida Wang, Min Hui Lo, Chunqiao Song, Yuwen Li, and Alex S. Gardner. Recent Changes in Land Water Storage and its Contribution to Sea Level Variations. *Surv. Geophys.*, 38(1):131–152, jan 2017.
- [86] Haizhou Wang and Mingzhou Song. Ckmeans.1d.dp: Optimal k-means clustering in one dimension by dynamic programming. *R J.*, 3(2):29–33, 2011.
- [87] Xuejing Wang, Hailong Li, Jiu Jimmy Jiao, D. A. Barry, Ling Li, Xin Luo, Chaoyue Wang, Li Wan, Xusheng Wang, Xiaowei Jiang, Qian Ma, and Wenjing Qu. Submarine fresh groundwater discharge into Laizhou Bay comparable to the Yellow River flux. *Sci. Rep.*, 5(1):8814, aug 2015.

- [88] GuiFeng Wei, DanLing Tang, and Sufen Wang. Distribution of chlorophyll and harmful algal blooms (HABs): A review on space based studies in the coastal environments of Chinese marginal seas. *Adv. Sp. Res.*, 41(1):12–19, 2008.
- [89] Adrian D. Werner, Mark Bakker, Vincent E.A. Post, Alexander Vandenbohede, Chunhui Lu, Behzad Ataie-Ashtiani, Craig T. Simmons, and D.A. Barry. Seawater intrusion processes, investigation and management: Recent advances and future challenges. *Adv. Water Resour.*, 51:3–26, jan 2013.
- [90] Alicia M. Wilson. Fresh and saline groundwater discharge to the ocean: A regional perspective. *Water Resour. Res.*, 41(2):1–11, feb 2005.
- [91] Authors Jean Wilson, Carlos Rocha, and Catherine Coxon. Report No . 172 Combining earth observation and geochemical tracing techniques for groundwater detection and evaluation in Ireland. Technical Report 172, 2016.
- [92] Jean Wilson and Carlos Rocha. Regional scale assessment of Submarine Groundwater Discharge in Ireland combining medium resolution satellite imagery and geochemical tracing techniques. *Remote Sens. Environ.*, 119:21–34, apr 2012.
- [93] Lin Yan, David P. Roy, Hankui Zhang, Jian Li, and Haiyan Huang. An automated approach for sub-pixel registration of Landsat-8 Operational Land Imager (OLI) and Sentinel-2 Multi Spectral Instrument (MSI) imagery. *Remote Sens.*, 8(6):520, jun 2016.
- [94] I. S. Zektser and Hugo A. Loaiciga. Groundwater fluxes in the global hydrologic cycle: past, present and future. *J. Hydrol.*, 144(1-4):405–427, apr 1993.
- [95] Igor S. Zektser, V. A. Ivanov, and A. V. Meskheteli. The problem of direct groundwater discharge to the seas. *J. Hydrol.*, 20(1):1–36, sep 1973.
- [96] Liangpei Zhang, Lefei Zhang, and Bo Du. Deep Learning for Remote Sensing Data: A Technical Tutorial on the State of the Art. *IEEE Geosci. Remote Sens. Mag.*, 4(2):22–40, jun 2016.
- [97] Qi Zhang and Penglin Zhang. An uncertainty descriptor for quantitative measurement of the uncertainty of remote sensing images. *Remote Sens.*, 11(13), 2019.