JNCC Report
No. 682

Towards an operational wildfire and muirburn monitoring system for Scotland. Report for the Caroline Herschel Framework Partnership Agreement for Copernicus User Uptake (Work Package six)

Duncan Blake, Alastair Graham, Alun Jones, Kostas Sideris
Karen Frake and Gwawr Jones

August 2021

ISSN 0963 8091

Framework Partnership Agreement for Copernicus User Uptake
FPA no.: 275/G/GRO/COPE/17/10042

Scotland's Nature Agency
Buidheann Nàdair na h-Alba

For further information please contact:
Joint Nature Conservation Committee
Monkstone House
City Road
Peterborough PE1 1JY
https://jncc.gov.uk/

This report should be cited as:
Blake, D., Graham, A., Jones, A., Sideris, K. and Frake, K. & Jones, G. 2021. Towards an operational wildfire and muirburn monitoring system for Scotland. Report for the Caroline Herschel Framework Partnership Agreement for Copernicus User Uptake (Work Package six). *JNCC Report No. 682*. JNCC, Peterborough, ISSN 0963-8091.

Author affiliation:
Duncan Blake (NatureScot)
Alastair Graham (Geoger Ltd)
Alun Jones (JNCC)
Kostas Sideris (JNCC)
Karen Frake (NatureScot)
Gwawr Jones (JNCC)

Acknowledgments:
This project was funded by the Caroline Herschel Framework Partnership Agreement on Copernicus User Uptake. The CEDA (Centre for Environmental Data Analysis) provided support with the use of JASMIN cloud computing facilities and data access via the CEDA Archive. Our thanks to colleagues at JNCC and NatureScot for their support and input into this project. We are grateful to Paula Lightfoot at JNCC for help with quality assuring the report.

JNCC EQA Statement:
This report is compliant with JNCC's Evidence Quality Assurance Policy
https://jncc.gov.uk/about-jncc/corporate-information/evidence-quality-assurance/

Executive summary

Vegetation burning in upland habitats forms part of moorland management regimes and also occurs naturally as wildfires. Mapping and monitoring the extent of upland burning gives insight into carbon emissions, biodiversity and natural capital accounting. Previous studies have shown that burn scars can be mapped using indices derived from Sentinel-2 imagery.

This project aimed to use Sentinel-2 data and a cloud computing infrastructure to develop an operational burn mapping and monitoring system for the whole of Scotland.

Outline of method

- **Site selection**
  The Isle of Skye and the Eastern Cairngorms were selected as primary study sites based on known burn activity and availability of cloud-free Sentinel-2 data. Lammer Law in the Southern Highlands was used as an independent site to test reproducibility.

- **Index selection**
  Pre- and post-burn Sentinel-2 imagery was selected for both sites and used to generate five burn indices and two vegetation indices. Investigation of the range of index values for burn scars and other landcover change classes led to selection of the following indices for burn detection:
    - Difference in Soil Adjusted Vegetation Index (dSAVI)
    - Post fire Normalised Burn Ratio (NBR)
    - Difference in Normalised Burn Ratio 2 (NBR2)

- **Identification of core burn pixels**
  Pixels were classified as ‘core burn’ if they met the following criteria:
    - dSAVI >= 0.2853 AND NBR >= 0.2395 AND dNBR2 <= 0.8

- **Identification of extended burn area**
  Pixels were classified as ‘extended burn area’ if they met the following criteria
    - dSAVI >= 0.206748 AND NBR >= 0.173447 AND dNBR2 <= 0.8

  Extended burn areas that overlapped core burn pixels were extracted and converted to a vector shapefile of burn extents.

- **Automation and scaling-up of method for national burn detection**
  The burn detection code was adapted to run on JASMIN cloud computing facilities with direct access to Sentinel-2 data in the CEDA (Centre for Environmental Data Analysis) archive covering the whole of Scotland.

- **Evaluation of outputs and recommendations for further development**
  The method was tested on Sentinel-2 data for Scotland for April 2020, using imagery with less than 95% cloud cover. Outputs were compared with burn maps produced by NatureScot following wildfires in April 2020. Findings were used to outline the steps required to develop this method into an operational national burn mapping system.

The initial investigation successfully detected the majority of both wildfire and managed burns at the two primary study sites and at the independent site. The national-scale test took less than 4 hours to analyse 130 pairs of Sentinel-2 images. It identified several burns which occurred in April, including some previously unknown burns. However, it also generated false positives, notably along the coast and in ploughed fields. Furthermore,

burns could only be identified if they occurred between two consecutive dates which both had cloud-free imagery.

Recommendations were made for further development to address over- and under-prediction of burnt areas. These include masking agricultural fields and the intertidal area, investigating the impact on accuracy of applying/not applying a cloud mask, and either lowering the cloud-cover threshold or splitting the imagery into a grid of smaller tiles to maximise the availability of usable imagery.

These enhancements should lead to a system which could be implemented operationally. A global threshold approach will never be perfect and will always produce false positives and false negatives. However, by enabling rapid, cost-effective, automated analysis of Sentinel-2 imagery, this system will improve knowledge of the extent, location and time periods of burning across Scotland.

Code development for this project was carried out in Python and Jupyter notebooks by specialist staff at NatureScot and by JNCC business associate Alastair Graham. All scripts have been made publicly available via GitHub.

Contents

1 Background

Wildfires are areas of uncontrolled burning of vegetation in rural areas. In Scotland they occur mainly on peatland or moorland, where they can cause significant damage to biodiversity, agricultural land, forestry and areas used for recreation and contribute to climate change. Muirburn is the controlled burning of moorland for land management purposes and is an extensive practice across Scotland.

This project developed and tested an automated workflow to map burn extents from Sentinel-2 imagery in Scotland. National monitoring of burn extents will help NatureScot understand the impact of wildfires and muirburn on carbon emissions, biodiversity and natural capital accounting. Using remotely sensed data will allow a national picture of burning to be developed with improved spatial and temporal detail.

This will help inform:

- The scale, distribution, timing and frequency of muirburn activity.
- Damage caused by wildfires to habitats and wildlife on protected sites (NatureScot has a duty to assess this).
- Good land management practices and behaviours that could be promoted to reduce the likelihood or impact of wildfires.

At a technical level the aim of the project is to make progress on an operational workflow to map wildfire and muirburn extent nationally with the possibility of frequent updates. To date NatureScot has only mapped specific, already known, wildfires, especially where they have damaged protected sites. Recent literature suggests that it should be possible to develop a more automated system using a time series of Sentinel-2 images at a national scale, for example the case study on the 2017 Italy wildfires (Filipponi 2019).

The objectives were:

- Refine burn extent mapping methods and investigate anomalies.
- Make the process more automated.
- Include muirburn (previously only wildfires have been targeted).
- Scale to a national level.
- Improve timeliness of burn extent products.

This work was delivered as part of JNCC's Copernicus Project, which was launched in September 2019 to increase uptake of Copernicus data and services across the UK via capacity building and cross-border collaboration. Through a set of work packages including training sessions, thematic workshops, and development of practical applications, the project aimed to facilitate the use of earth observation (EO) data to deliver public environmental functions more efficiently or effectively across multiple policy areas.

JNCC's Copernicus Project is funded by the European Commission under the Caroline Herschel Framework Partnership Agreement on Copernicus User Uptake (FPCUP), which was established in 2018 to increase the use of Copernicus data, products and services.
Methodology.

2 Methodology

2.1 Selection of test areas and imagery

Test areas were selected to evaluate suitable indices and thresholds that could be included within an operational workflow. The Isle of Skye and Eastern Cairngorms were chosen as they represent very different habitats, topography and land management regimes.

An area of the eastern Cairngorms around Balmoral was chosen as it displayed good examples of heather muirburn of varying shapes and sizes. A pair of clear Sentinel-2 granules from 12 March 2019 and 16 May 2019 (Figure 1) were used as pre- and post-fire images (although there is some snow cover in the March granule). The site also had reasonably concurrent aerial photography coverage. Burn polygons were manually digitised by comparing the two S2 images to identify burns that had occurred between them. Underlying aerial photography was used to verify the digitisation.

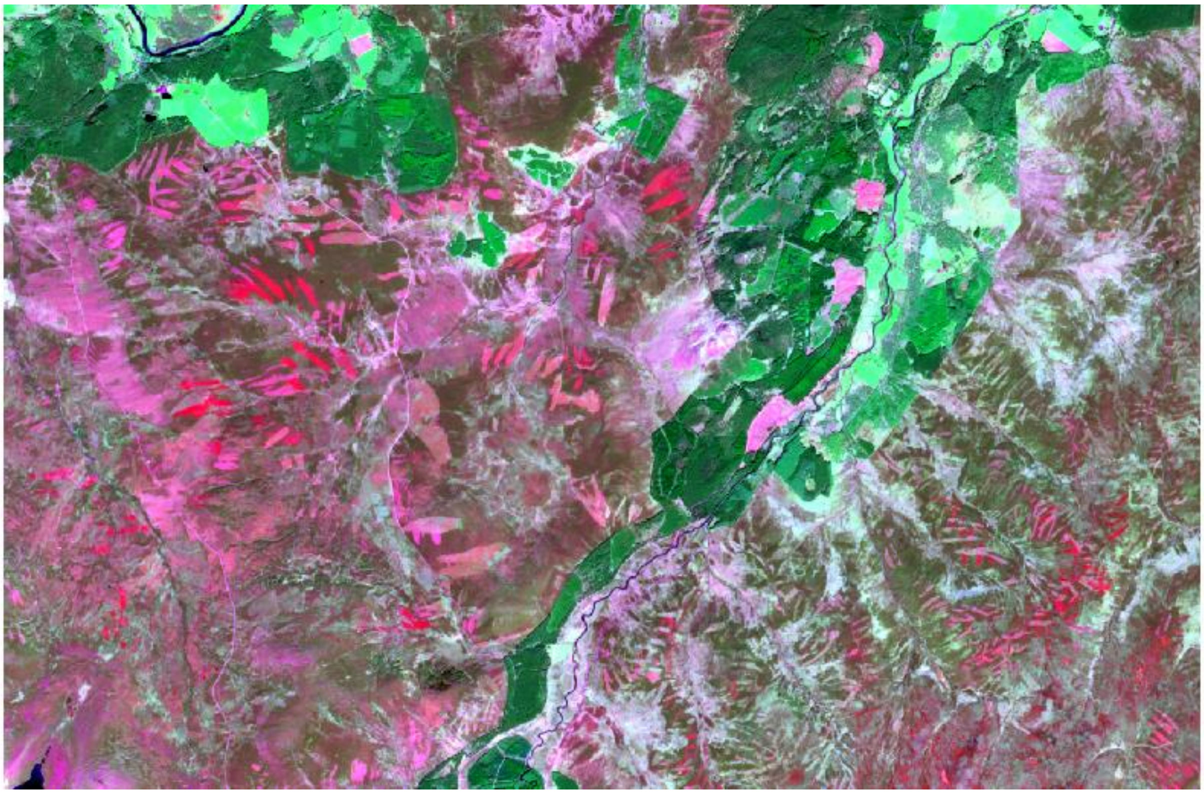

Figure 1: Sentinel-2 image of eastern Cairngorm study site 16 May 2019 (bands SWIR2, NIR, green). Image width = 13.5 km. Brighter pink areas represent recent burn scars.

Secondly the island of Skye was used (Figure 2) as this represents a more west coast sheep grazing land management system and steeper topography. Also, a previous mapping exercise by NatureScot using pre- and post-fire images from 25 February 2018 and 17 March 2018 had already mapped a series of wildfires and burns across the island during this time period.

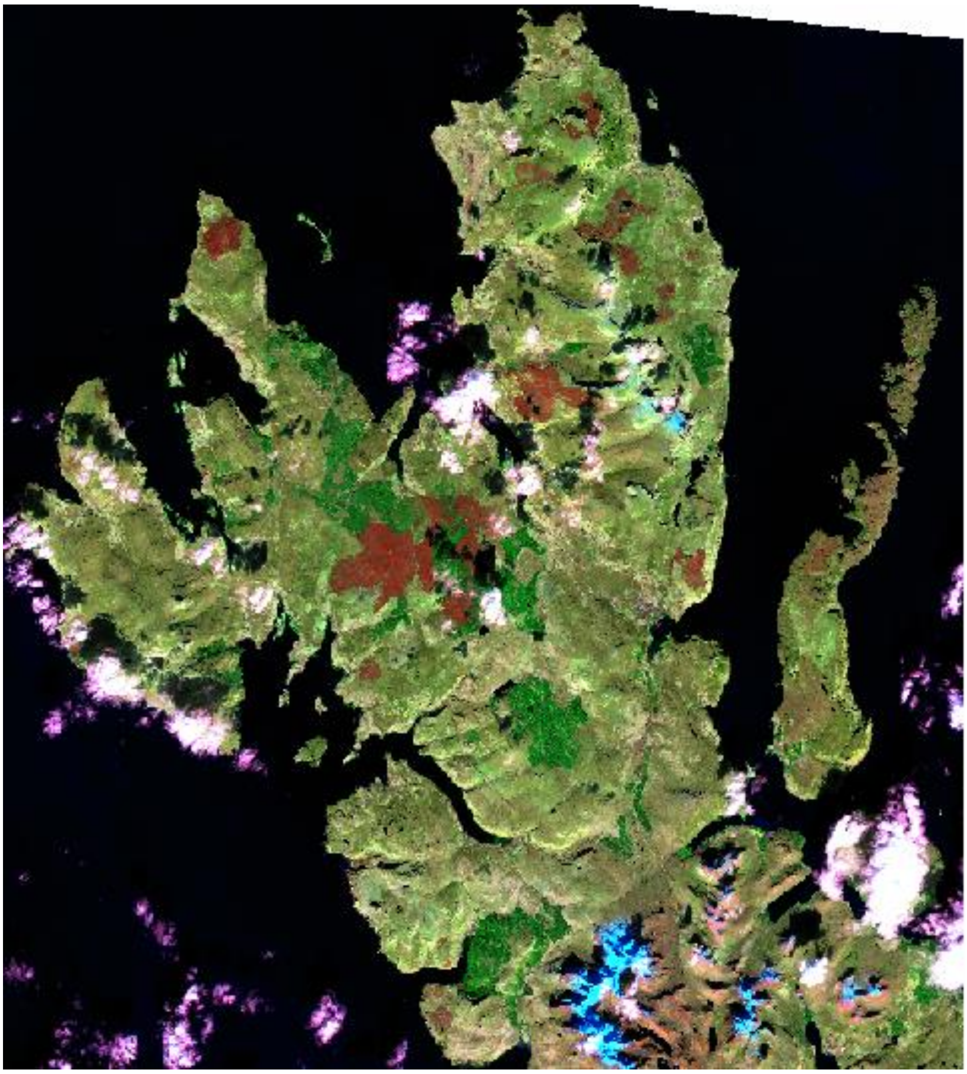

Figure 2: Sentinel-2 image of Skye 17 March 2018 (bands SWIR2, NIR, green). Image width = 52km. Burn scars are visible as dark brown patches.

2.2 Examination of indices on two test sites

The following seven indices were calculated for the two test sites on pre- and post- fire images and difference indices were also calculated by subtracting the pre-fire index from the post-fire index (Table 1).

Table 1: Indices calculated from Sentinel-2 imagery.

Index name	Acronym	Python formula
Normalised Burn Ratio	NBR	(SWIR1 - NIR)/(SWIR1 + NIR)
Normalised Burn Ratio 2	NBR2	(SWIR2 - SWIR1)/(SWIR2 + SWIR1)
Normalised Difference Vegetation Index	NDVI	-1*(NIR – Red)/(NIR + Red)
Normalised Mid Infrared Burn Index	nMIRBI	(((10*SWIR2)-(9.8*SWIR1+2)) / ((10* SWIR2)+(9.8* SWIR1+2)))
Char Soil Index	CSI	(-1 * (NIR/SWIR2))
Burned Area Index for Sentinel-2	BAIS2	(-1*(((1 - sqrt((re6 * re7 * NIR8A)/ red)) * (((SWIR2 - NIR8A) / (sqrt(SWIR2 + NIR8A))) + 1))))
Soil Adjusted Vegetation Index	SAVI	(-1 * (1.5 * ((NIR - red) / (NIR + red + 0.5))))

A Python script produced a 21 band geoTiff containing pre-fire, post-fire and difference indices (Table 2).

Table 2: List of pre-fire, post-fire and difference indices generated from Sentinel-2 imagery.

Band	Index name
1	pre_nbr
2	post_nbr
3	dnbr
4	pre_nbr2
5	post_nbr2
6	dnbr2
7	pre_ndvi
8	post_ndvi
9	dndvi
10	pre_nmirbi
11	post_nmirbi
12	dnmirbi
13	pre_csi
14	post_csi
15	dcsi
16	pre_bais2
17	post_bais2
18	dbais2
19	pre_savi
20	post_savi
21	dsavi

Some research studies used a spectral separability index to compare burnt and unburnt pixels to try and deduce which indices are the most effective at detecting burns (Smiraglia *et al.* 2020; Filipponi 2019). However, selecting scattered non-burn pixels around the burn areas or even randomly across the image is unlikely to pick out those areas where confusion is most likely.

Instead, training data were digitised across the two areas through visual interpretation of satellite and aerial imagery to include burn extents along with other land cover classes burns could be confused with and also ‘change classes’ between the two images (such as changes in crops, tree cover, clouds, shadow or snow).

Where confusion is most likely depends on whether it is the change index (d) or the post-fire index that is being considered. The classes most likely to show confusion were considered to be:

- Bare peat
- Shadow
- Rock
- Water
- Bare fields
- Felled areas

The changes most likely to be confused are:
- Crop harvesting
- Felling

- Changes in cloud cover, cloud shadow, topographic shadow and snow cover between dates.

Not all these classes or change combinations are evident in every image pair (for example as both of the time periods are spring there is plenty of snow to no snow change but not *vice versa)*. These are shaded in grey in Table 3, which shows the number and total area of training data polygons digitised at the two trial sites.

Table 3: Number and total area of training data polygons per land cover class. Classes and change combinations which are not present in a given image pair are shaded (and contain '0').

Class	Class ID	Skye		Cairngorms		Combined
		No. polygons	Area (ha)	No. polygons	Area (ha)	Total (ha)
Bare peat	11	6	0.2	14	1.3	1.5
Bare fields	12	0	0	2	12.8	13
Burn extents	13	36	5,571	86	98.3	5,669
Felled forest	14	1	14	2	13	27
Other (veg)	15	6	238	5	42	280
Rock	16	8	24	4	1.5	26
Shadow	17	6	57	8	4.7	62
Water	18	4	42	3	0.9	43
Cloud2nocloud	21	0	0	2	7	7
Nocloud2cloud	22	4	148	0	0	148
Shadow2noshadow	23	3	28	6	5	33
Noshadow2shadow	24	7	150	0	0	150
Snow2nosnow	25	7	2	8	36	38
Nosnow2snow	-	0	0	0	0	0
Harvesting	26	0	0	2	1	1
Felling	27	0	0	3	1	1

Box plots showing the range of index values for each class were created from the training data for each site and also from the combined training dataset from both sites to examine which indices might be best used to separate burns from other classes.

2.3 Thresholding to identify core burn pixels

The aim of thresholding is to classify pixels as burnt if their index values fall above or below a certain index threshold or combination of thresholds. These thresholds were based empirically on the training data from the two test sites combined to try and obtain values that would work best across the country.

Applying the same threshold(s) to imagery from all dates and locations has clear limitations in terms of geographic variability, habitat variability and image variability from different orbits. A literature review of dynamic thresholding was undertaken, however the techniques described were complex and often require the creation of training data which would be difficult to implement in an automated national system (Appendix 3: Literature Review).

Given the time constraints of this project it was decided to use fixed thresholds to start with, with a view to refining this if the results were not sufficiently accurate.

Interpretation of the box plots (see Appendix 1 for the complete set of box plots) especially those containing the combined data from both sites led to the conclusion that:

- dSAVI (first plot, Figure 3) showed the best separability with the other change classes (especially snow to no snow) though with considerable overlap with field to bare field and a little overlap with trees to no trees and no cloud to cloud.
- With the two test areas combined the post-fire indices were more confused than analysis of a single site but NBR (second plot, Figure 3) or NBR2 showed the best separability.

Owing to the observations above a combined rule of the **difference in SAVI** and the **post fire NBR** index was trialled (with a threshold of their median value used to identify core burn pixels). This combination minimises errors of commission particularly in the Skye test area. This combination of using a difference index and a post-fire index has also been used in other studies (Filipponi 2019).

However, some errors where present if a change in cloud cover persisted. The **difference in NBR2** index (third plot, Figure 3) was used to remove many of these errors.

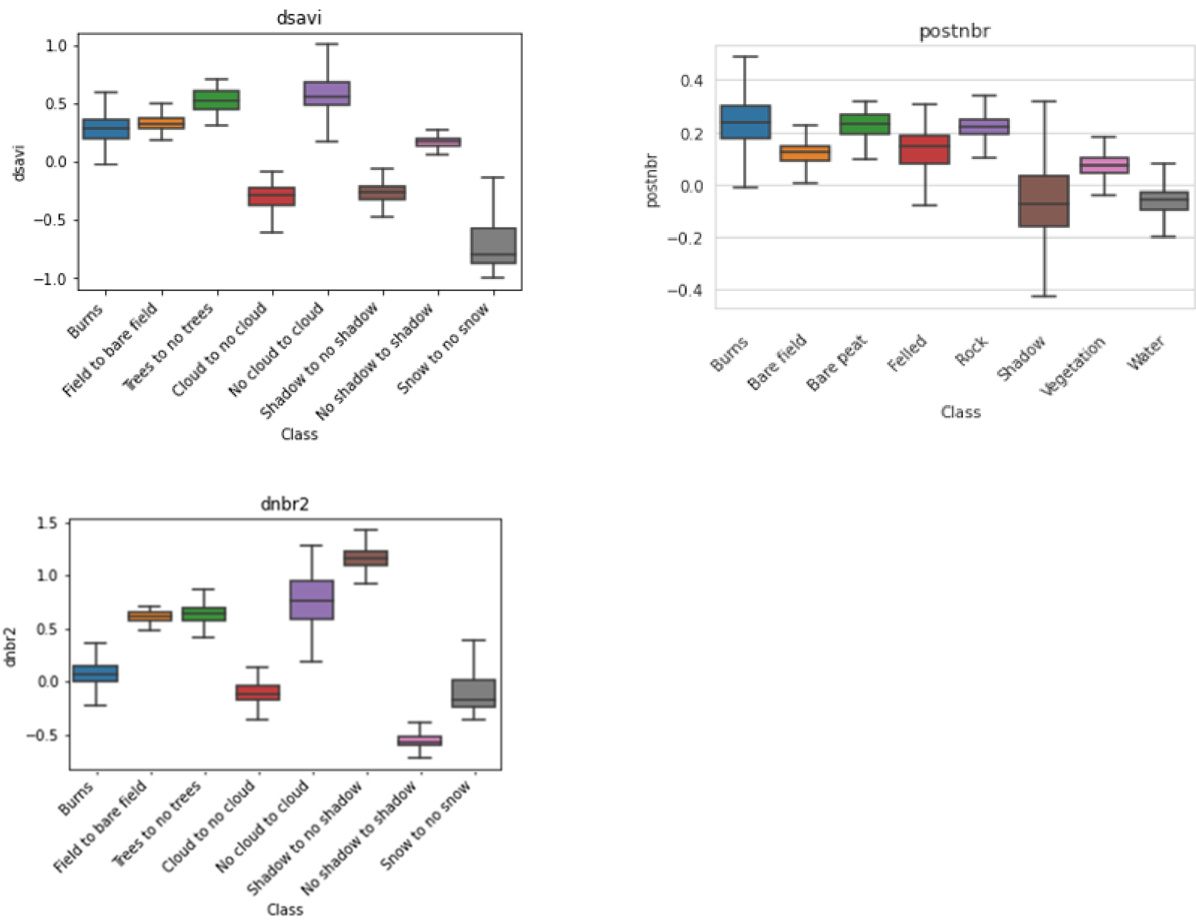

Figure 3: Box plots from the combined test areas for the 3 indices used in the prototype system.

Final thresholds applied were:

dSAVI >= 0.2853 AND NBR >= 0.2395 AND dNBR2 <= 0.8

2.4 Apply lower thresholds to map extent of burn areas

The thresholds used in Section 2.3 were necessarily quite strict to avoid too many errors of commission. As a result, a large portion of each burn area may be missed. One approach would be to implement a region growing algorithm to grow the core burn pixels. However, an easily implementable algorithm in core Python modules could not be found. Some code was investigated (https://github.com/charmichokshi/Region-Growing-Algorithm-on-RGB-Image) but it was not clear how it worked and would not work directly with GeoTIFF files.

However, a simpler approach is to create a second layer using lower thresholds:

dSAVI >= 0.206748 AND NBR >= 0.173447 AND dNBR2 <= 0.8

This layer includes more pixels in the burn areas, approximating more closely to the actual burn extents, but will also identify other pixels that are not burns. To exclude these, the pixels identified by applying the lower thresholds are clumped together into objects and only objects that have core burn pixels identified are retained to produce the final burn layer.

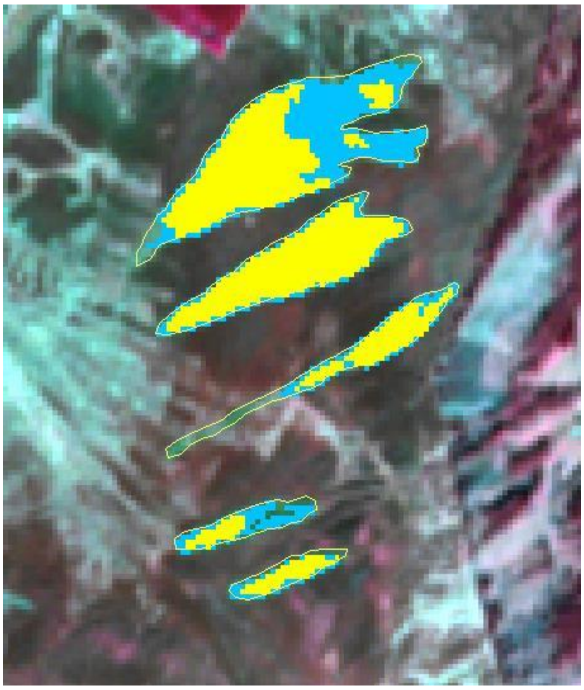

Figure 4: Area of Eastern Cairngorms showing the effect of the lower threshold 'region growing' method - core burn pixels in yellow, extended burn areas in blue, yellow outline manually digitised burn boundary.

The result is converted to a vector shapefile of estimated burn extents with dates of the images used as attributes.

2.5 Scale up the code to run nationally

Initial testing of the previous steps was carried out using Python Jupyter notebooks in a Google Colab environment. The code was then ported to run at scale within the JASMIN ‘super-data-cluster’. A JASMIN virtual machine was set up using MobaXterm software to access it. This gives direct access to the imagery residing in the CEDA archive which forms the Simple ARD Service developed for Scotland and Northern Ireland. The code was set up to run on the grid of 26 Sentinel-2 tiles shown in Figure 5:

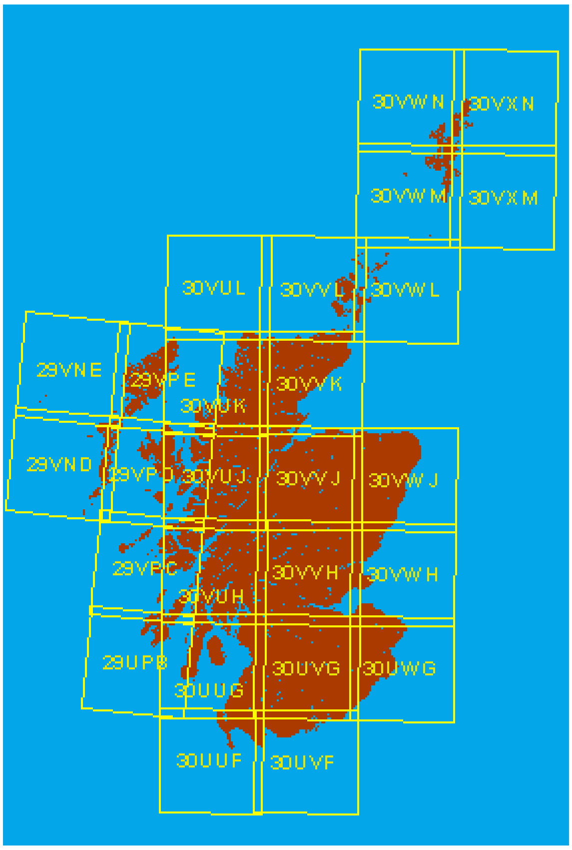

Figure 5: S2 tiles that are processed within the operational system. The background shows the extent of the raster mask, areas of sea and inland water with a value of 0 and land with a value of 1.

When running nationally a mask is used to exclude areas of the imagery covered by sea (a land polygon created from the OS MasterMap Mean Low Water Spring line and the Scottish border was used) or inland water (based on OS VectorMap District). This is combined with the cloud and topographic shadow masks for each Sentinel-2 granule.

For testing purposes, the code in JASMIN was run on Sentinel-2 data for Scotland for April 2020. This process examined the 154 full Sentinel-2 granules available for that month. Granules with the same tile number are arranged in date order and each granule is compared with the most recent granule captured before it. This resulted in 130 granule comparisons. Processing time was 3 hours 51 minutes so roughly **two minutes per comparison.**

One practical point to note is that a run of the code can time out if the remote computer goes to sleep or there is disruption to the connection. This can be avoided by the use of Linux ‘screens’ as described in Appendix 2.

2.6 Evaluate outputs to inform recommendations for methodological improvements

The burn maps produced for the whole of Scotland were evaluated with reference to known burn data from April 2020 (six fires mapped manually by NatureScot affecting designated sites) to check for errors of commission and omission. More widely a sample of areas flagged as burns were checked against the underlying S2 imagery to see if they were likely to be burns.

The results of this evaluation were used to inform recommendations for future development and improvements to the methodology.

3 Code location and summary

3.1 Python code for testing the method

A Github repo called burn-mapping has been set up to host the Jupyter notebooks: https://github.com/duncansnh/burn-mapping.

This contains the following code:

- **CUU_burn_extent_indices.ipynb** - generates 21 indices on the Sentinel-2 image pairs (7 from image 1, 7 from image 2 and 7 difference indices) for testing purposes and boxplot generation.
- **CUU_burn_extent_pixel_box_plots.ipynb** - generates boxplots for the difference indices and classes and boxplots for the postfire image indices and land cover classes. Can be run on either pilot site by changing input parameters.
- **CUU_burn_extent_pixel_box_plots_both_sites.ipynb** - as above but combined data from both pilot sites to produce the box plots.
- **CUU_burn_extent_image_thresholding.ipynb** - code to threshold the image based on multiple thresholds of indices and removes clumps of small numbers of pixels. Also applies the 'region growing' and outputs a shapefile. Note the method used in this script was improved later in the operational code to keep outputs as rasters until the final export step.

3.2 Python code for operational system

Operational code was developed by Alastair Graham, Geoger Ltd. in: https://github.com/ajggeoger/JNCCBurnCalcs

This was forked and developed further in: https://github.com/Scottish-Natural-Heritage/GIG-JNCC-Muirburn.

Full details on how to run the code in JASMIN are in Appendix 2, however the code essentially goes through the following process:

- Working directories set and checked.
- Log file started
- Count files in working directory
- **picklecheck** – returns list of previously processed images if it exists from output directory or creates an empty list.
- **getdatalist** – list of files to be processed from working directory sorted by granule and date (oldest at beginning of list, most recent at end)
    - Calls **cleanlistfunc** – removes files already processed by comparing lists.
- Removes granules which are not full granules by getting the file size and keeping if **>1Gb***
- If list is less than 2 it stops processing
- Loop through list of files to process:-
    - Last image becomes a 'post fire' image (using pop)
        - Calculates cloud raster name
        - Gets the transform and dimensions of the image.

            - Calls **maskimage** with image path, transform, no_rows, no_cols and cloud path as inputs.
                - Imports land/sea/inland water raster to the extent of the image
                - Creates a cloud mask of ones and zeros
                - Creates a topographic shadow mask of ones and zeros
                - Reads in the S2 bands required and multiplies them by the land/sea, cloud and shadow layers to create zeros where there is cloud, water or shadow.
                - Returns S2 array and profile
        - Next to last image becomes a ‘pre fire’ image (using pop)
            - Processes same as post image.
        - If granule name of pre and post image match:
            - Calculate indices
            - Calculate thresholds returning core burn pixels (sievedarray)
            - Calculate region growing threshold and return array (burnedarray).
            - (Optionally) save rasters of seeds and burn areas.
            - Save shapefile of outputs.
        - If list of files to process is equal to or greater than one:
            - Post fire image becomes pre fire image
- Writes out pickled list of files processed, and text file equivalent.
- Cleans up temporary files.
- Stops timer.

***** Please note: The **1GB** rule chooses granules with full image coverage to minimise issues surrounding division by null.

4 Results

4.1 Test sites

4.1.1 East Cairngorms

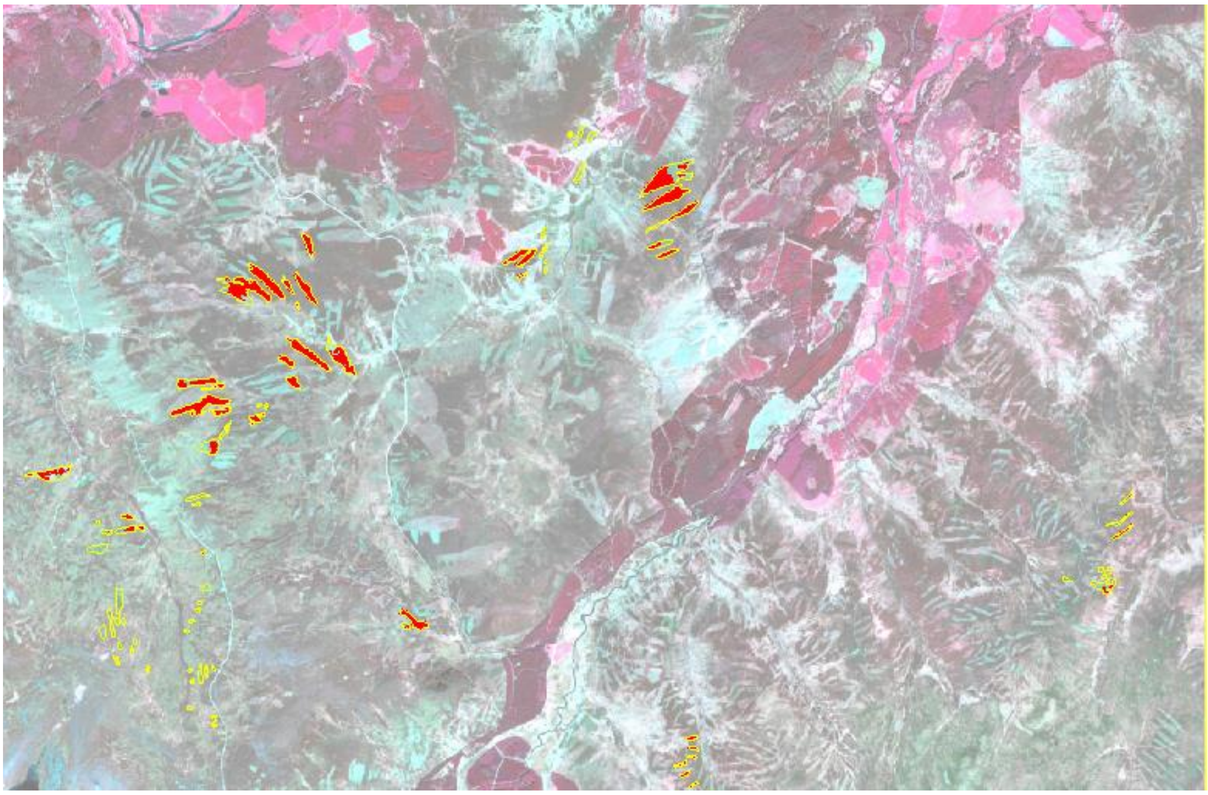

Figure 6: East Cairngorms test site showing manually digitized burn extents in yellow, pixels identified by the automated method in red.

Figure 6 shows burn locations at the East Cairngorms site that were successfully identified by the automated analysis of S2 imagery).

42 of 86 burn scar polygons (49%) were detected or 65 ha of 98.3 ha (66%) of the total burnt area was highlighted with the automated method (Table 4).

Table 4: Number and area of burn scars detected through pixel thresholding and through manual mapping at the East Cairngorms test site.

	Pixel thresholding	Manually mapped	Percentage agreement
Number of burn scars	42	86	49%
Area of burn scars (ha)	60	98.3	61%
Commission errors (ha)	5	0	N/A

Twenty-three burn scars in the south west corner were partially covered in snow in the pre-fire image so the threshold approach could not identify them.

4.1.2 Skye

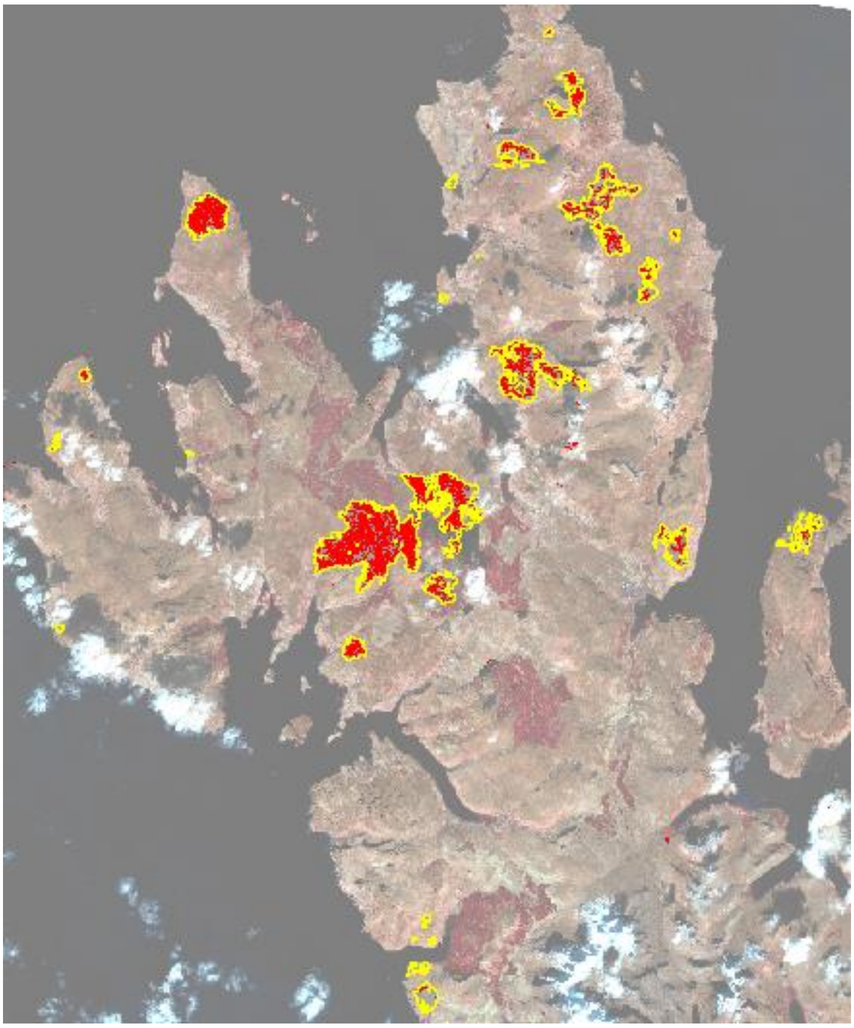

Figure 7: Skye test site showing digitized burn extents in yellow, pixels identified by the automated method in red.

Figure 7 shows the mapping of burn scars on Skye in red using the global thresholds. The results of comparing the outputs with the results of manual mapping are shown in Table 5.

Table 5: Number and area of burn scars detected through pixel thresholding and through manual mapping at the Skye test site.

	Pixel thresholding	Manually mapped	Percentage agreement
Number of burn scars	25	36	69%
Area of burn scars (ha)	3,365	5,571	60%
Commission errors (ha)	169	0	N/A

A couple of smaller burns in the south are not so well identified. Figure 8 shows an additional burn site that had been missed by manual assessment but also some commission errors on the edges of clouds (note cloud and shadow were not masked for the test sites but were masked for the national burn mapping which will eliminate some of these errors though not all).

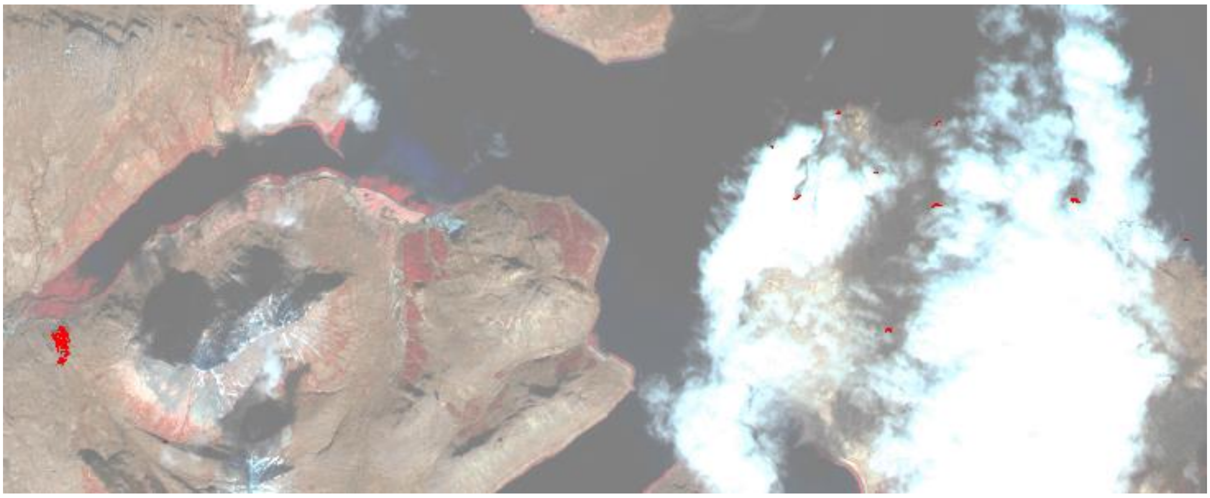

Figure 8: Additional burn area and commission errors on the edge of clouds.

4.1.3 Independent site test – Lammer Law

As a test the thresholds derived from the analysis of the combined Skye and Cairngorms data were used on a third area – Lammer Law in the Lammermuir Hills – as previous work by NatureScot had already mapped burn patches from a pair of granules captured on 20 September 2019 and 19 April 2020. Some of the patches are quite small, only a few pixels.

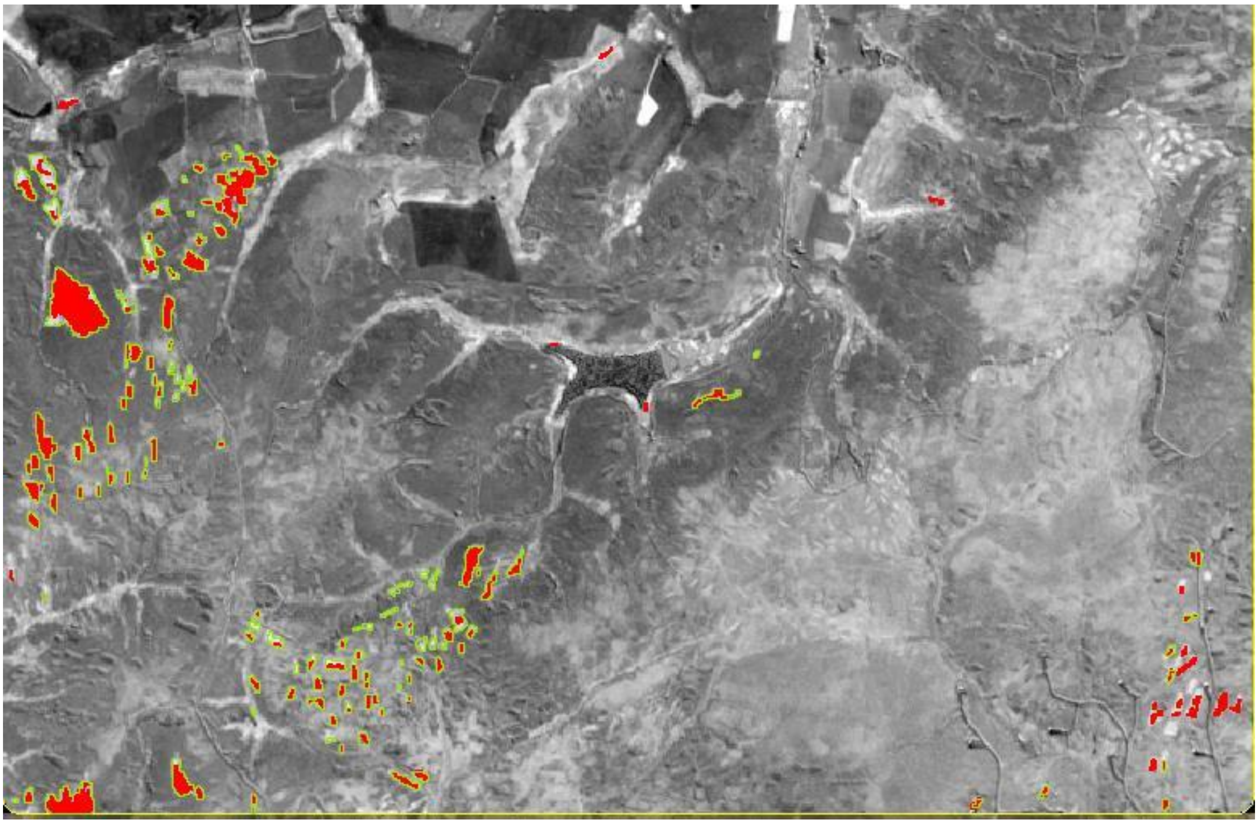

Figure 9: Lammer Law muirburn patches. Green outlines are from previous NatureScot mapping, red areas are the pixels identified by the automated approach with the thresholds defined in Section 3.3.

The results of comparing the outputs with the results of manual mapping are shown in Table 6.

Table 6: Number and area of burn scars detected through pixel thresholding and through manual mapping at the Lammer Law independent test site.

	Pixel thresholding	Manually mapped	Percentage agreement
Number of burn scars	89	115	77%
Area of burn scars (ha)	25	52	48%
Commission errors (ha)	2.5	0	N/A

There were 23 additional core burn pixel areas that did not intersect a previously mapped burn area. However, of these 13 are almost certainly burns (south east corner of Figure 9) so the commission errors in the table are an overestimate.

Of the remaining ten predicted burn areas that did not coincide with previously mapped burns:

- four could feasibly be burns (unfortunately there is no concurrent aerial photography to confirm)
- three represented one area along a watercourse (reason unknown)
- three were areas of water (this was run before an inland water mask was implemented in the operational code).

Note this test was carried out only on the core burn areas before the lower thresholds were implemented to extend the burn extents, therefore the area of burn scars detected is probably an underestimate.

4.2 Running nationally

Scaling the automated burn detection process to run nationally introduces additional issues, both in terms of the land cover types and changes encountered and owing to the method of granule comparisons.

4.2.1 Overall outputs for Scotland for April 2020

Combining the outputs of the 130 comparisons for April 2020 generated 4,017 polygons amounting to 807 hectares. A visual assessment of Figure 10 immediately shows many of these are along the coast and may be due to changes in cliff shadow, tidal water or slight image misalignments.

Figure 10: Overall results from April 2020 analysis showing coastal errors of commission in red.

4.2.2 Successfully detected burns

False Colour S2 imagery (NIR, Red and Green bands) was used as a reference to evaluate whether the areas mapped as burns had been correctly identified because burned areas show up clearly as dark patches in this band combination. A selection of sites where burns were known to have occurred were investigated and the results shown below.

Loch Morar

Figure 11 shows burnt areas automatically detected by a comparison of indices derived from Sentinel-2 data acquired on 20 and 25 April 2020.

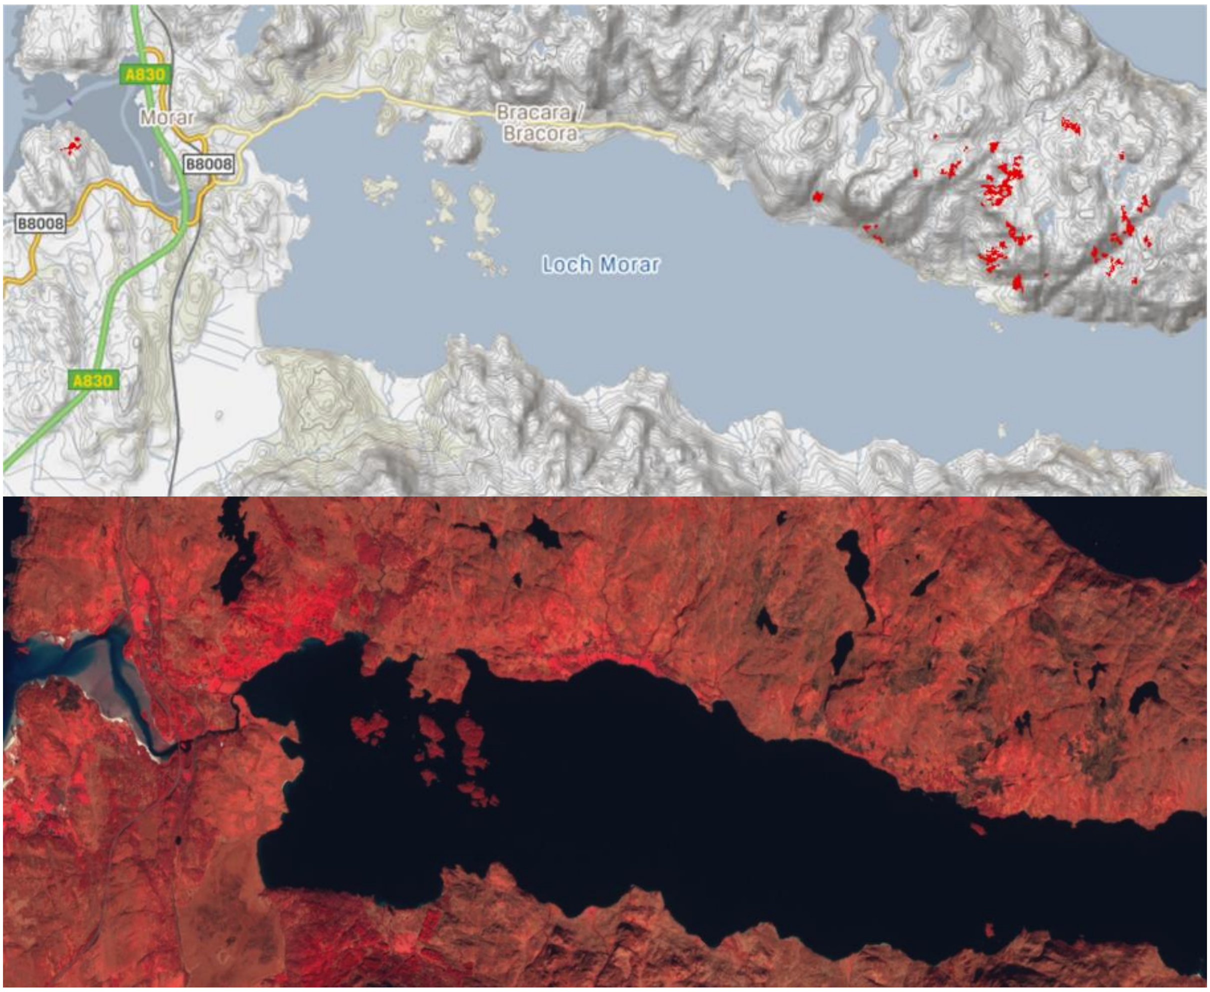

Figure 11: Top image - Burns automatically detected between 20 and 25 April 2020 at Loch Morar, Bottom image - false colour infrared image from 25 April for comparison. Imagery © Sentinel Hub.

Looking in more detail (Figure 12) there is a good correlation between the detected burns and the false colour Sentinel-2 imagery. They are an underestimate of the total extent that would be gained if manually digitising, missing some of the less severely burned areas, but this has to be balanced against the increase in false positives if the thresholds were relaxed further.

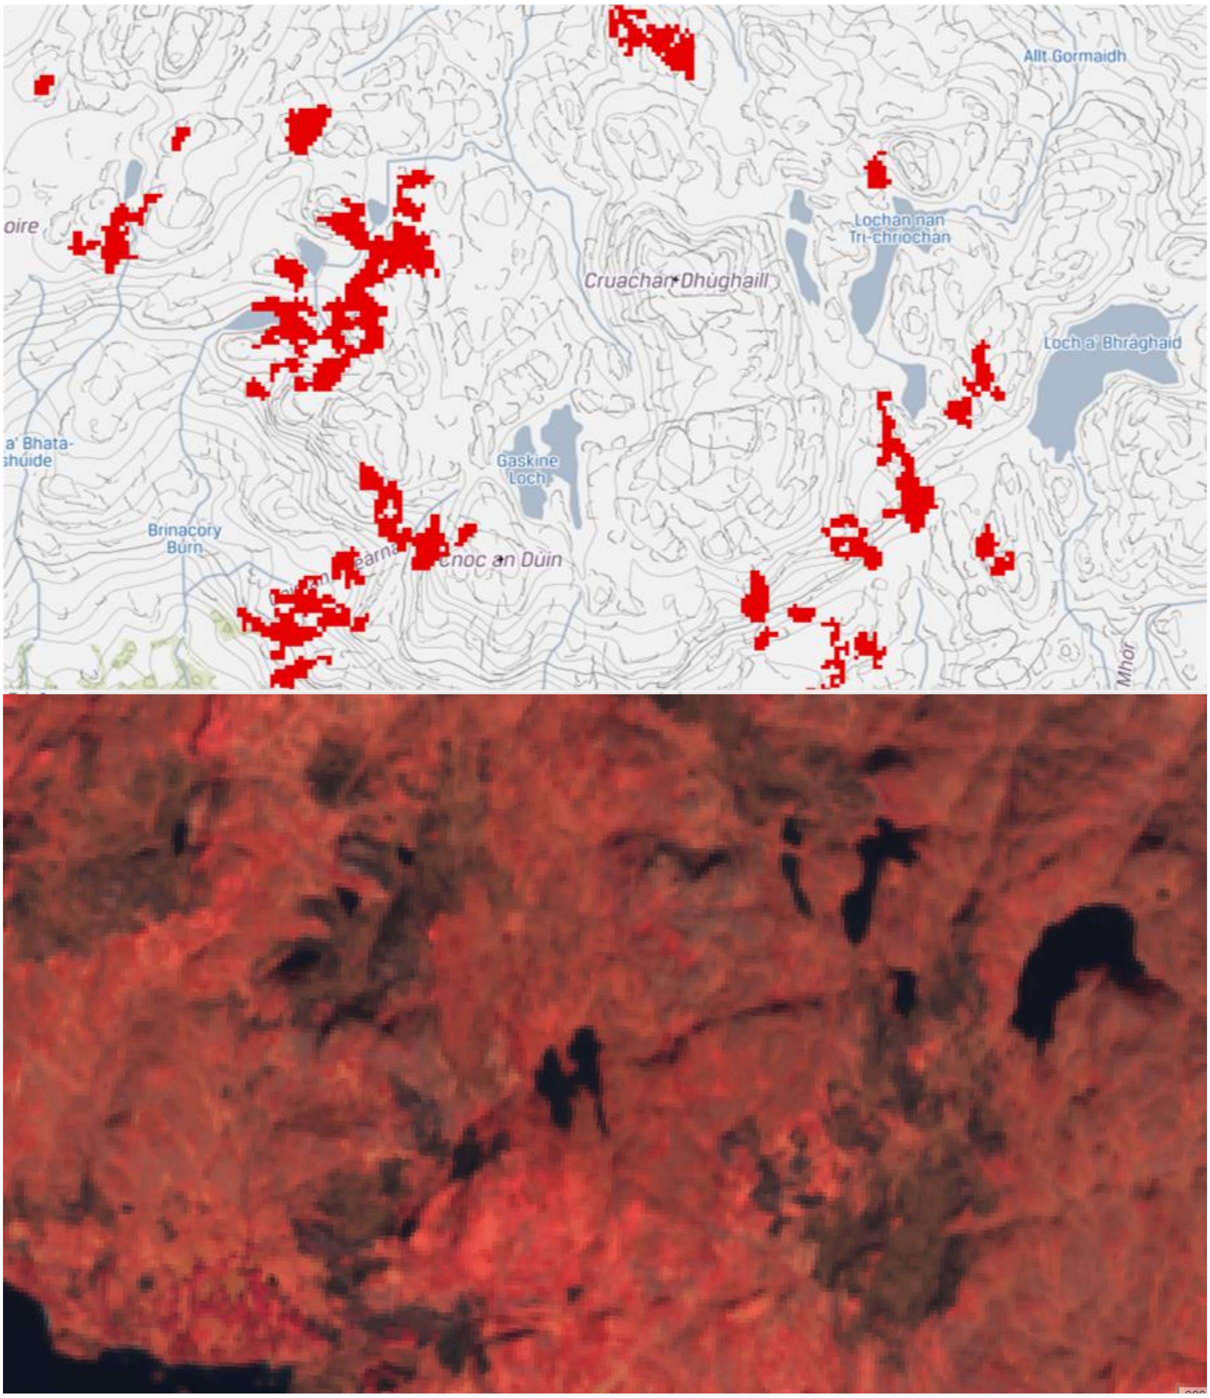

Figure 12: Top image - Burns automatically detected between 20 and 25 April 2020 at Loch Morar, Bottom image - false colour infrared image from 25 April for comparison. Imagery © Sentinel Hub.

Loch Garve

A large wildfire was detected that took place between 17 and 20 April 2020, though it is at a granule boundary (shown in yellow below) so only half of it has been detected. The western side of the fire was not detected as the pre-fire granule of 17 April was captured on a different orbit and was a partial granule so was filtered out of the comparison list as described in Section 4.2. Furthermore, the previous pre-fire full granule on 15 April had very dense cloud cover so the area was not visible (see Section 5.2.2 for further explanation of this issue).

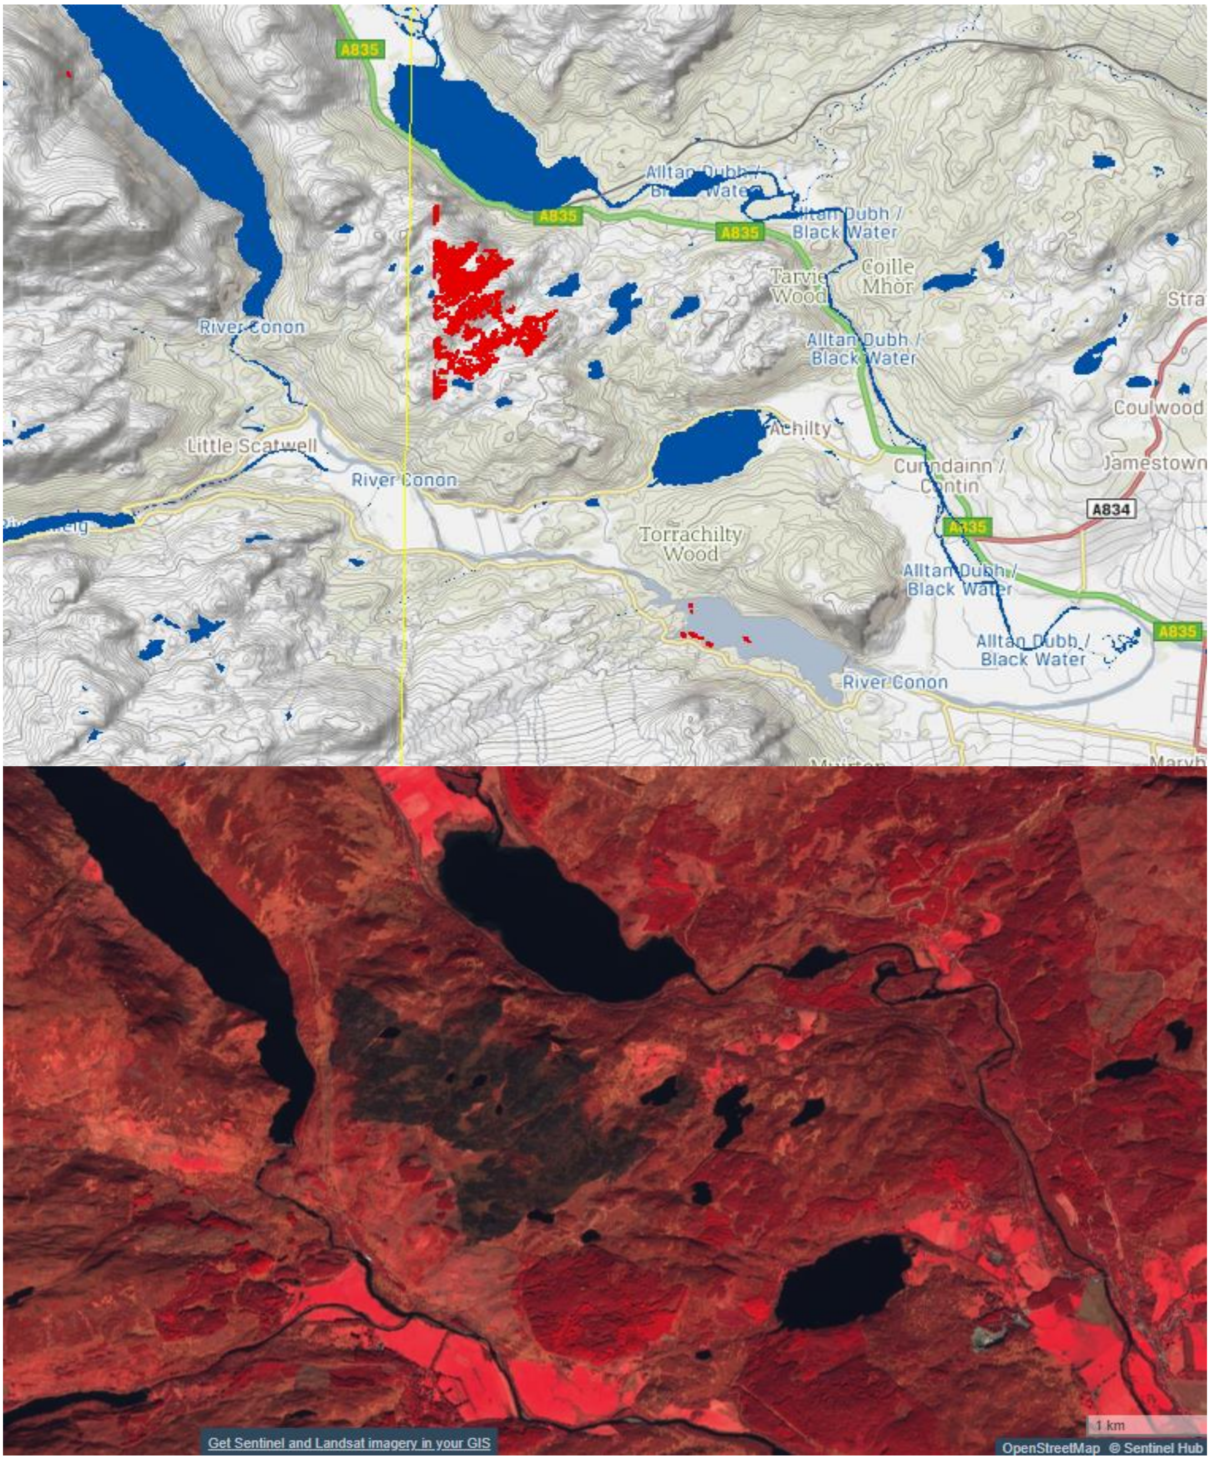

Figure 13: Top image - Wildfire at Garve between 17 and 20 April 2020. Note an error in the inland water mask has led to false positives at the loch on the River Conon. Imagery © Sentinel Hub.

4.2.3 Known wildfires in April 2020

As an example of an area with known wildfires in April 2020, Tinto Hills in Lanarkshire, was examined. Burning was known to have occurred on 11 April 2020 (Figure 14). Sentinel-2 imagery for granule T30UVG captured on the following dates was processed, the dates shown in green are pre-fire images, those in black are post-fire images (Table 7).

Table 7: Pre- and post-fire images processed for Sentinel-2 granule T30UVG, showing percentage cloud cover and visibility of burn area. Pre-fire dates (2 April and 7 April) are given in green font.

Date	Cloud cover from S2 metadata	Burn area visible	Burns detected
2 April	78%	Partially	N/A
7 April	82%	No	No
17 April	72%	Yes	No
19 April	0%	Yes	No
22 April	0.5%	Yes	No
27 April	1%	No	No

The process currently works by a pairwise analysis of S2 granules in date order. Although the burns are clearly visible in 3 of the images taken in April (for example Figure 15) burns can only be automatically detected by comparison with a pre-fire image and use of the difference in the SAVI (dSAVI). In this case the pre-fire image from 7 April was very cloudy (Figure 16) and the burn area had been masked out. Therefore, the comparison with 17 April could not take place and the burns were not detected.

If the image from 17 April had been compared with the pre-fire image of 2 April some of the burn scars would have been detected but because of the pairwise process this does not occur. Comparisons between later images, for example between 17 and 19 April would not detect the burn either because dSAVI would be very low as the burn had already occurred.

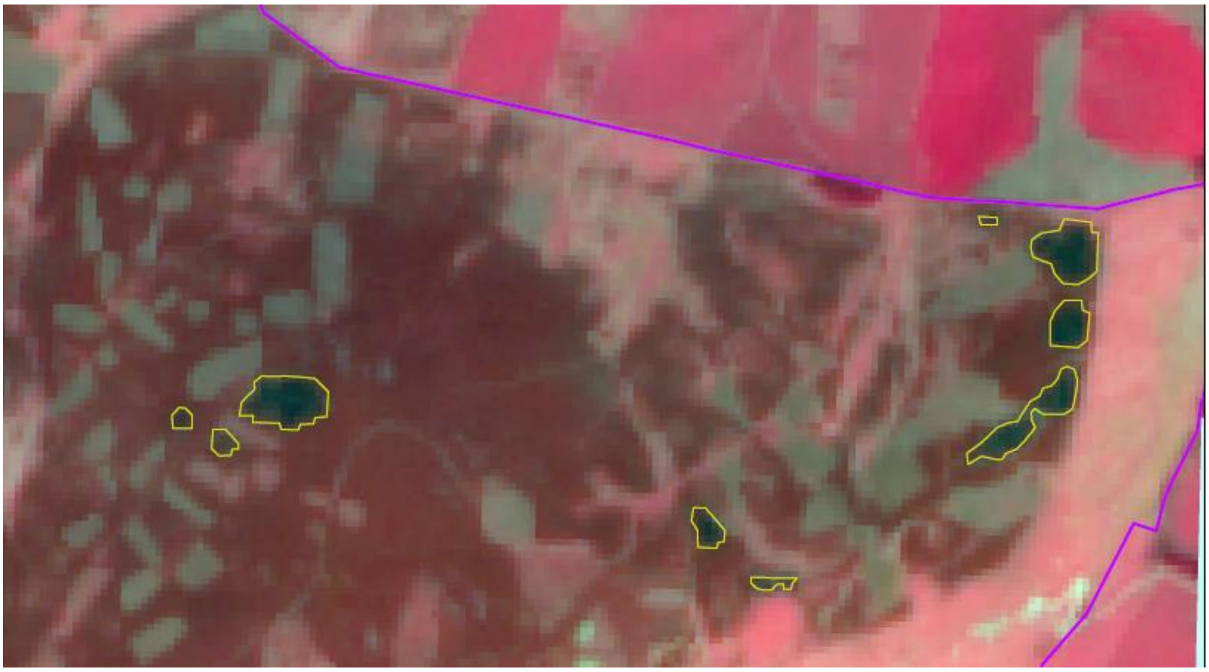

Figure 14: Burn scars from fires that occurred at Tinto Hills on 11 April 2020 manually delineated on false colour infrared Sentinel-2 imagery. Image width 2km. Imagery © Sentinel Hub.

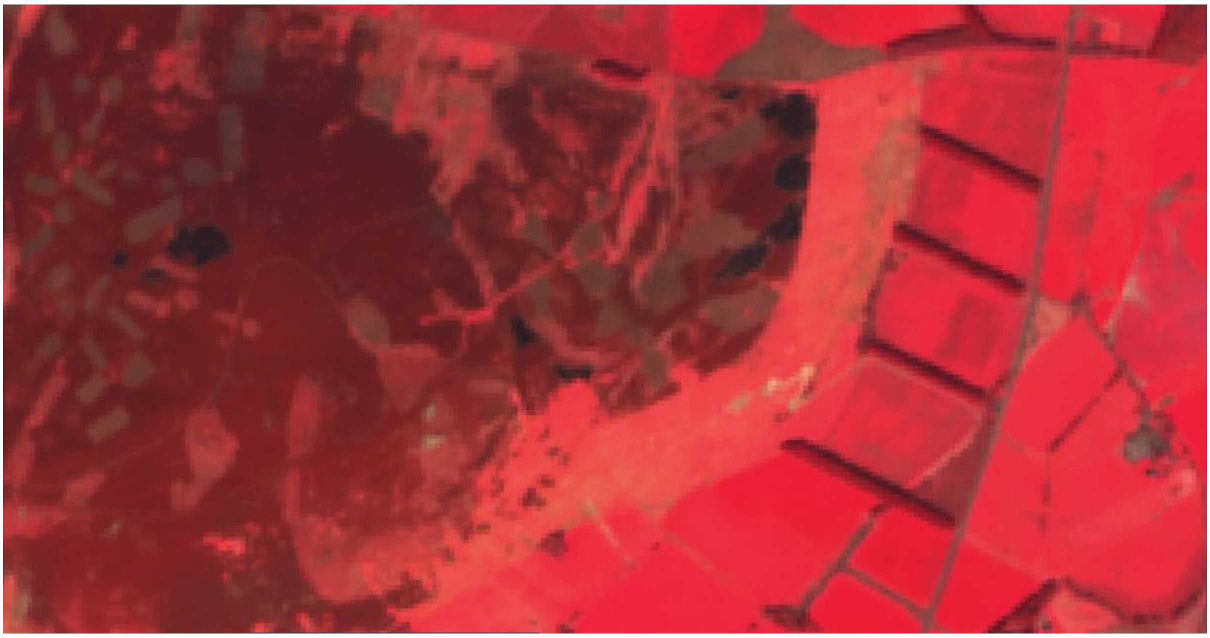

Figure 15: False colour infrared S2 image of Tinto Hills from 19 April 2020. Imagery © Sentinel Hub.

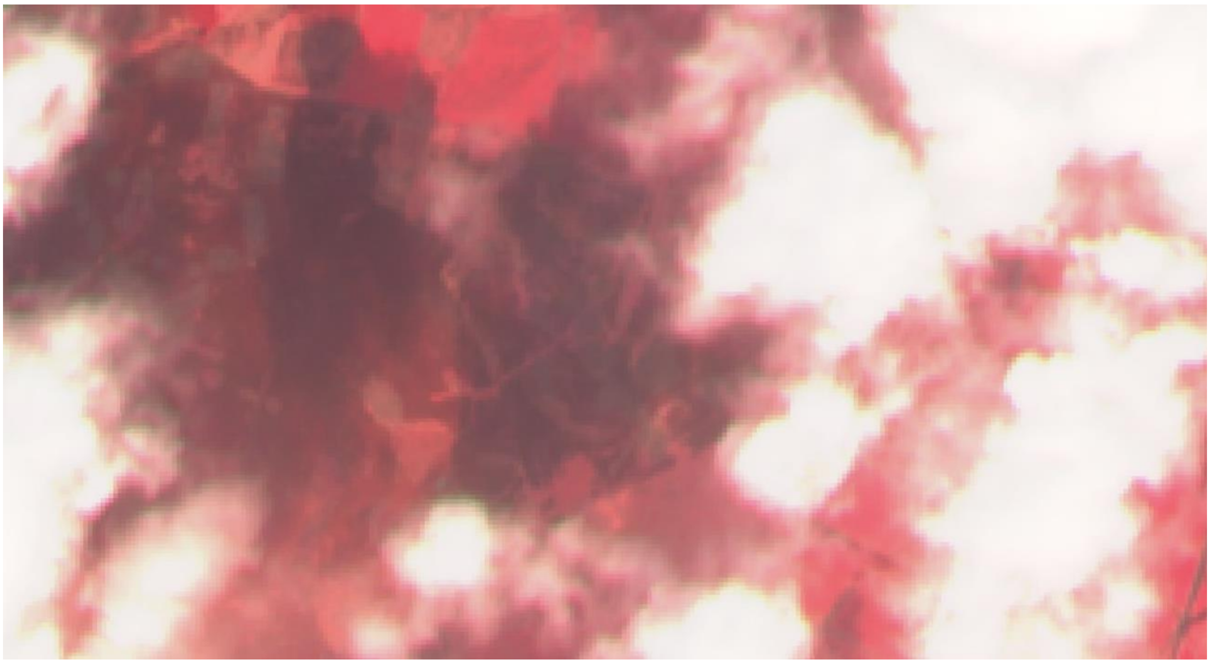

Figure 16: False colour infrared S2 image from 7 April 2020. Imagery © Sentinel Hub.

The same issue affected the other five wildfires NatureScot had mapped in April.

4.2.4 False positives in automated burn detection

The national burn data for April 2020 was further checked for errors of commission and some common types of error are noted:

Agricultural fields: In some place fields that have been ploughed (Figure 17) have been highlighted as potential burn scars:

Figure 17: Example of ploughed fields that were highlighted as potential burns. © Bluesky International Ltd. & Getmapping Plc. (2021).

Cloud edges: although the further threshold on the dNBR2 index reduced the number of false positives at the edge of clouds and cloud shadow some still remain (Figure 18).

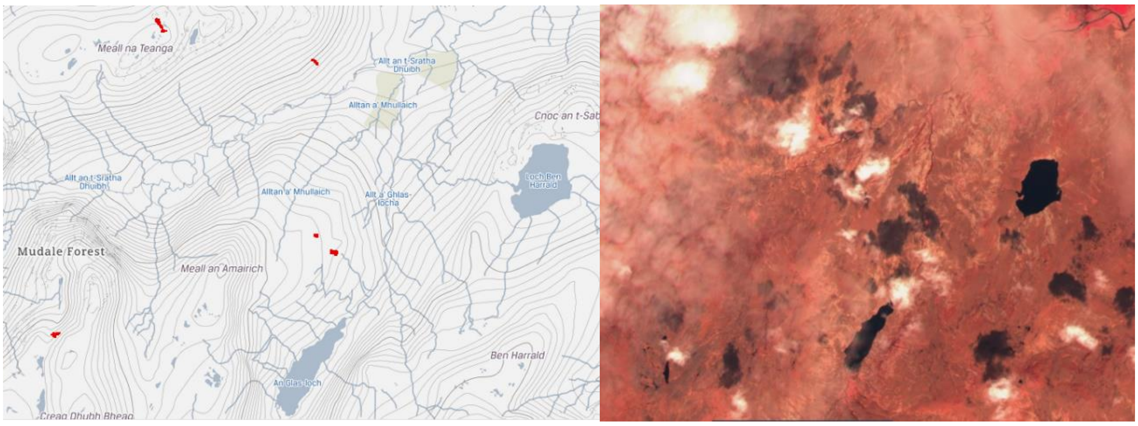

Figure 18: Cloud/shadow areas incorrectly classified as 'burn' through automated analysis of a clear pre fire image on 20 April and a partially cloudy post fire image on 25 April (right hand image © Sentinel Hub).

Coastal/tidal false positives

In some areas there are lots of false positives around the coast (see Figure 19). This is an issue with using the Mean Low Water Spring line as the land sea boundary. Mostly this is a result in changes to the tide level meaning the shore is emersed in one image and underwater in another. The change in index values can be similar to a burn on land. Using Mean High-Water Spring would eliminate the vast majority of these, as also shown in Figure 19.

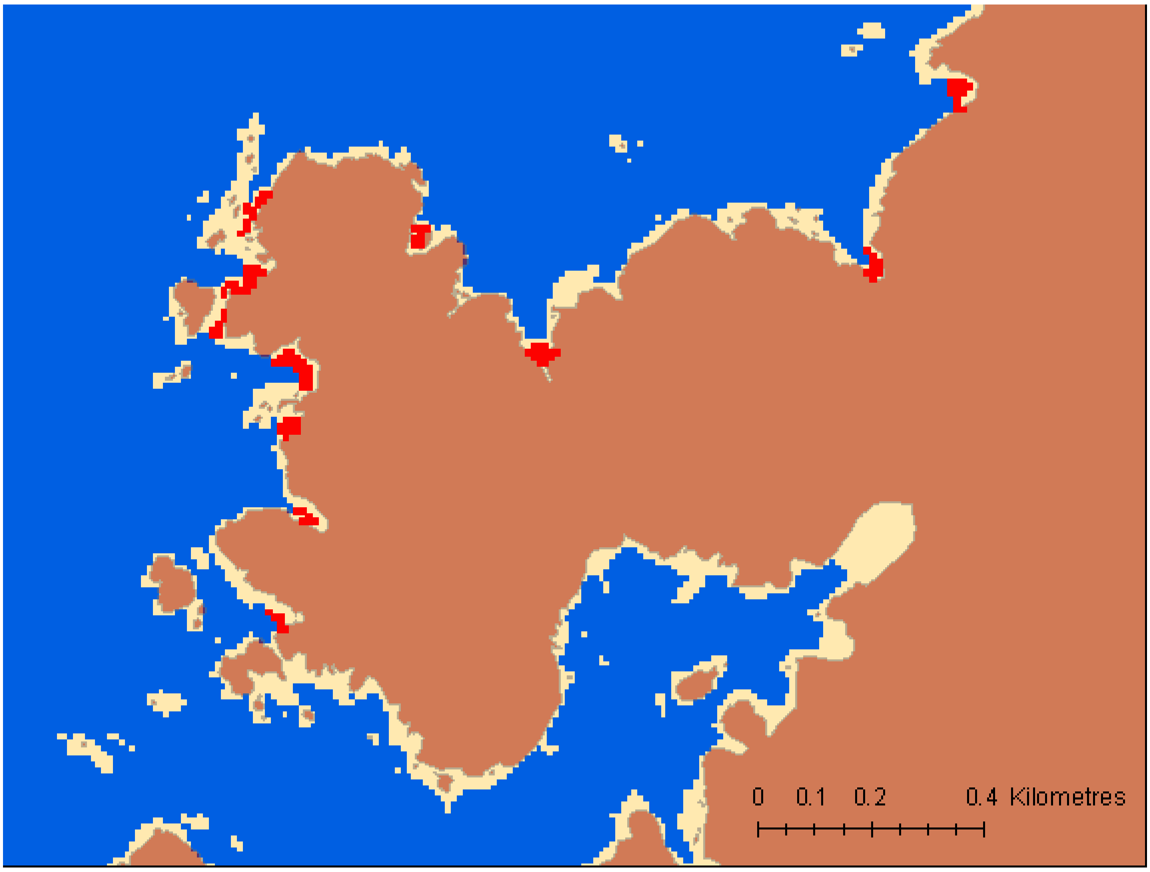

Figure 19: False positive burn scars in red in the tidal region around the island of Rona. The current land and sea mask is shown in beige and blue; the Ordnance Survey MasterMap Mean High Water Spring defined land mask is shown in darker brown. © Crown copyright [and database rights] 2021 OS 100017908

5 Recommendations for future development

Evaluation of the outputs from the local and national trials led to the following recommendations for future development. In order of priority:

1. A major shortcoming is the failure to detect burns due to the sequential nature of Sentinel-2 granule comparison whether or not there was significant cloud cover. If, after a burn has occurred, the next overpass contains cloud, there will never be a comparison between a later clear post-fire image and a clear pre-fire image. The simplest solution would be to limit the processing to granules with a lower threshold of cloud cover, though if this threshold is set too low lots of imagery will be discarded that could have been used to detect burns degrading the system's ability to both detect and date burns within as narrow a time period as possible. In the longer term a better solution should be sought. One potential option would be to split the granule into smaller segments for comparison, as demonstrated by the Coast X-Ray method developed by the Dynamic Coast project in Scotland (Fitton *et al*. 2021).

2. Currently, the earliest date granule in a processing run will not be compared with a previous image. This would not have a large impact if the code were run for the whole burn season at once. However, if it were run on a monthly basis this could lead to a number of burns being overlooked. Therefore, the code should be amended so that granules can be compared with granules processed in a previous run rather than just against images in the current run.

3. Amend the mask to be a MHWS coastline instead of MLWS to exclude any confusion in tidal water areas where errors of commission are common.

4. Amend the mask to exclude agricultural fields as ploughing can be confused with burns. Alternatively investigate using a dCSI threshold as this showed promise in the test datasets.

5. Test whether the results are better or worse for including the cloud and cloud shadow mask. The cloud mask is not very accurate and contains over- and under-predictions (as detailed on p20 of the Simple ARD User Guide). This study did not evaluate whether excluding the cloud and cloud shadow masks would affect the accuracy of burn predictions. Whether this is carried out may depend on the option pursued at point 1 and consequently how much of an issue cloud cover continues to be.

6. The current process excludes granules that are edge of orbit granules (as described in section 3.2) so some areas are being missed. Work is required to test whether these granules would work in the process or code amendment is required.

7. Simple process improvements such as combining burn areas from each image comparison into a single output vector file for that run rather than a separate file per comparison. Also configure the logfile to be more streamlined and contain only important information.

8. If the recommended changes above still don't produce an accurate enough output (i.e. too many errors to work as a semi-automated system) then re-investigate whether dynamic thresholds can be used.

9. Likewise re-investigate region growing algorithms rather than using the lower threshold approach. If efficient region growing code is found or developed the

process could be easily updated. It's possible this could result in 'cleaner' boundaries as it would act like a segmentation rather than thresholding on pixels

It is expected that implementing the enhancements outlined above in a second phase of work would result in an operational system that would enable routine analysis of S2 ARD at a national scale to meet the objectives of better understanding the scale, location and frequency of both wildfires and muirburn across the country.

6 References

Bastarrika, A., Chuvieco, E. & Martín, M.P. (2011) Mapping burned areas from Landsat TM/ETM+ data with a two-phase algorithm: Balancing omission and commission errors. *Remote Sensing of Environment*. 115 (4), 1003–1012.

Boschetti, M., Stroppiana, D. & Brivio, P.A. (2010) Mapping Burned Areas in a Mediterranean Environment Using Soft Integration of Spectral Indices from High-Resolution Satellite Images. *Earth Interactions*. 14 (17), 1–20.

Filipponi, F. (2019) Exploitation of Sentinel-2 Time Series to Map Burned Areas at the National Level: A Case Study on the 2017 Italy Wildfires. *Remote Sensing*. 11 (6), 622.

Fitton, J.M., Rennie, A.F., Hansom, J.D. & Muir, F.M.E. (2021) Remotely sensed mapping of the intertidal zone: A Sentinel-2 and Google Earth Engine methodology. *Remote Sensing Applications: Society and Environment*. 22100499.

Giglio, L., Loboda, T., Roy, D.P., Quayle, B. & Justice, C.O. (2009) An active-fire based burned area mapping algorithm for the MODIS sensor. *Remote Sensing of Environment*. 113 (2), 408–420.

Goodwin, N.R. & Collett, L.J. (2014) Development of an automated method for mapping fire history captured in Landsat TM and ETM+ time series across Queensland, Australia. *Remote Sensing of Environment*. 148206–221.

Hawbaker, T.J., Vanderhoof, M.K., Beal, Y.-J., Takacs, J.D., Schmidt, G.L., Falgout, J.T., Williams, B., Fairaux, N.M., Caldwell, M.K., Picotte, J.J., Howard, S.M., Stitt, S. & Dwyer, J.L. (2017) Mapping burned areas using dense time-series of Landsat data. *Remote Sensing of Environment*. 198504–522.

Kolden, C.A., Lutz, J.A., Key, C.H., Kane, J.T. & van Wagtendonk, J.W. (2012) Mapped versus actual burned area within wildfire perimeters: Characterizing the unburned. *Forest Ecology and Management*. 28638–47.

Nolde, M., Plank, S. & Riedlinger, T. (2020) An Adaptive and Extensible System for Satellite-Based, Large Scale Burnt Area Monitoring in Near-Real Time. *Remote Sensing*. 12 (13), 2162.

Roteta, E., Bastarrika, A., Padilla, M., Storm, T. & Chuvieco, E. (2019) Development of a Sentinel-2 burned area algorithm: Generation of a small fire database for sub-Saharan Africa. *Remote Sensing of Environment*. 2221–17.

Smiraglia, D., Filipponi, F., Mandrone, S., Tornato, A. & Taramelli, A. (2020) Agreement Index for Burned Area Mapping: Integration of Multiple Spectral Indices Using Sentinel-2 Satellite Images. *Remote Sensing*. 12 (11), 1862.

Stroppiana, D., Azar, R., Calò, F., Pepe, A., Imperatore, P., Boschetti, M., Silva, J.M.N., Brivio, P.A. & Lanari, R. (2015) Integration of Optical and SAR Data for Burned Area Mapping in Mediterranean Regions. *Remote Sensing*. 7 (2), 1320–1345.

Stroppiana, D., Bordogna, G., Carrara, P., Boschetti, M., Boschetti, L. & Brivio, P.A. (2012) A method for extracting burned areas from Landsat TM/ETM+ images by soft aggregation of multiple Spectral Indices and a region growing algorithm. *ISPRS Journal of Photogrammetry and Remote Sensing*. 6988–102.

Verhegghen, A., Eva, H., Ceccherini, G., Achard, F., Gond, V., Gourlet-Fleury, S. & Cerutti, P.O. (2016) The Potential of Sentinel Satellites for Burnt Area Mapping and Monitoring in the Congo Basin Forests. *Remote Sensing*. 8 (12), 986.

Appendix 1: Full box plot outputs

East Cairngorms difference classes and difference indices boxplots

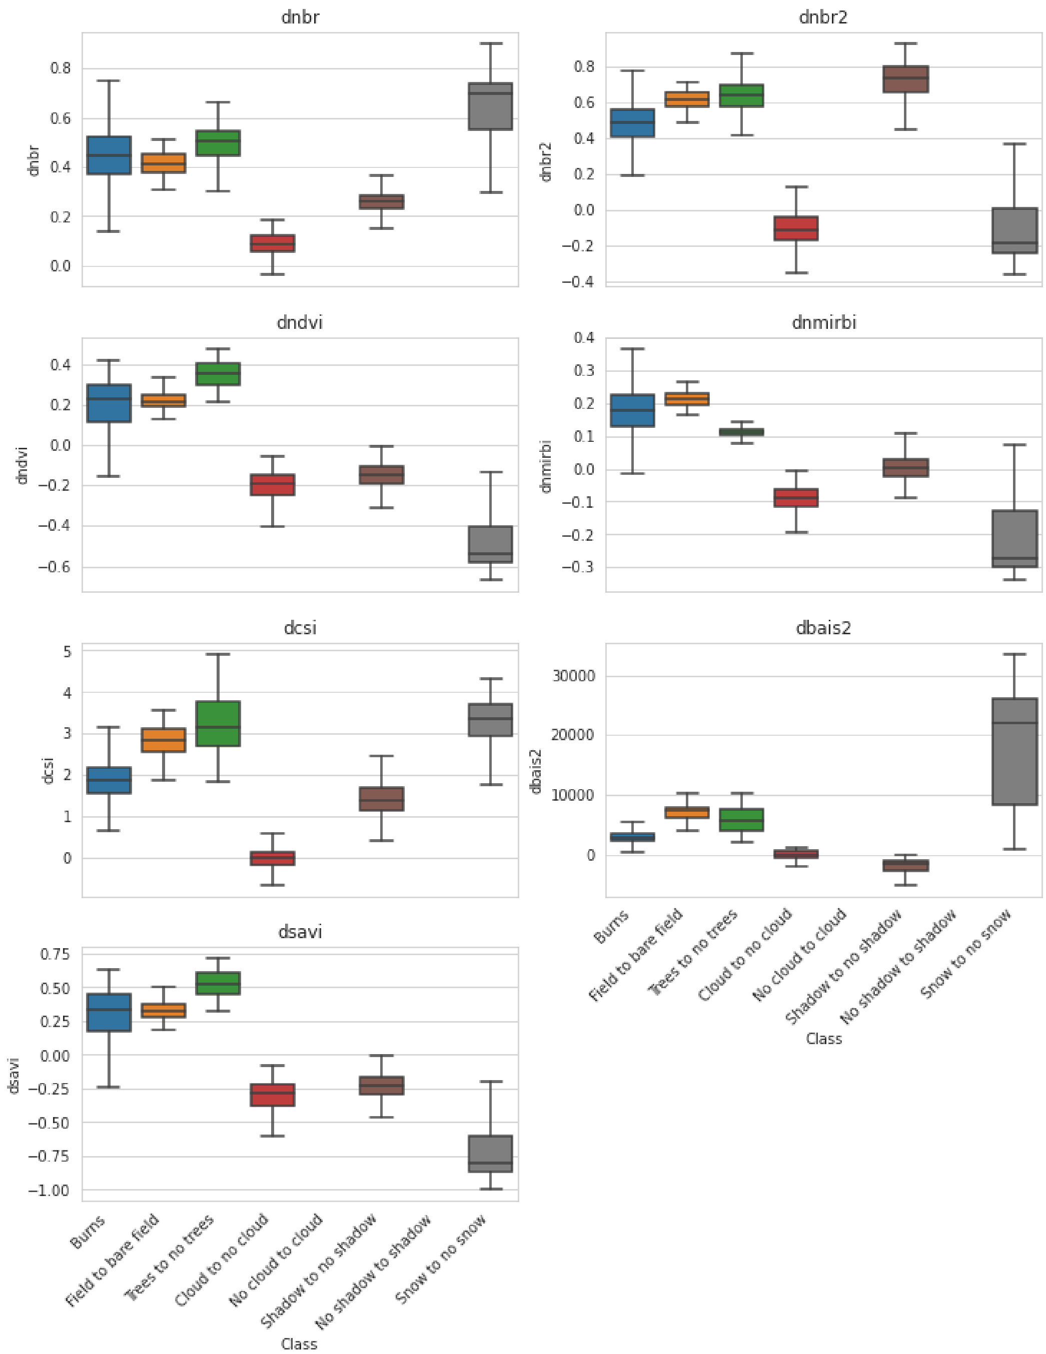

Observations:

- dSAVI and dNDVI are the best at separating burns from the ‘image difference’ classes. However, there is still confusion with other land use change that involves removal of vegetation.
- dCSI shows the best separability between burns and these land cover changes.

East Cairngorms post burn image indices and land cover class boxplots

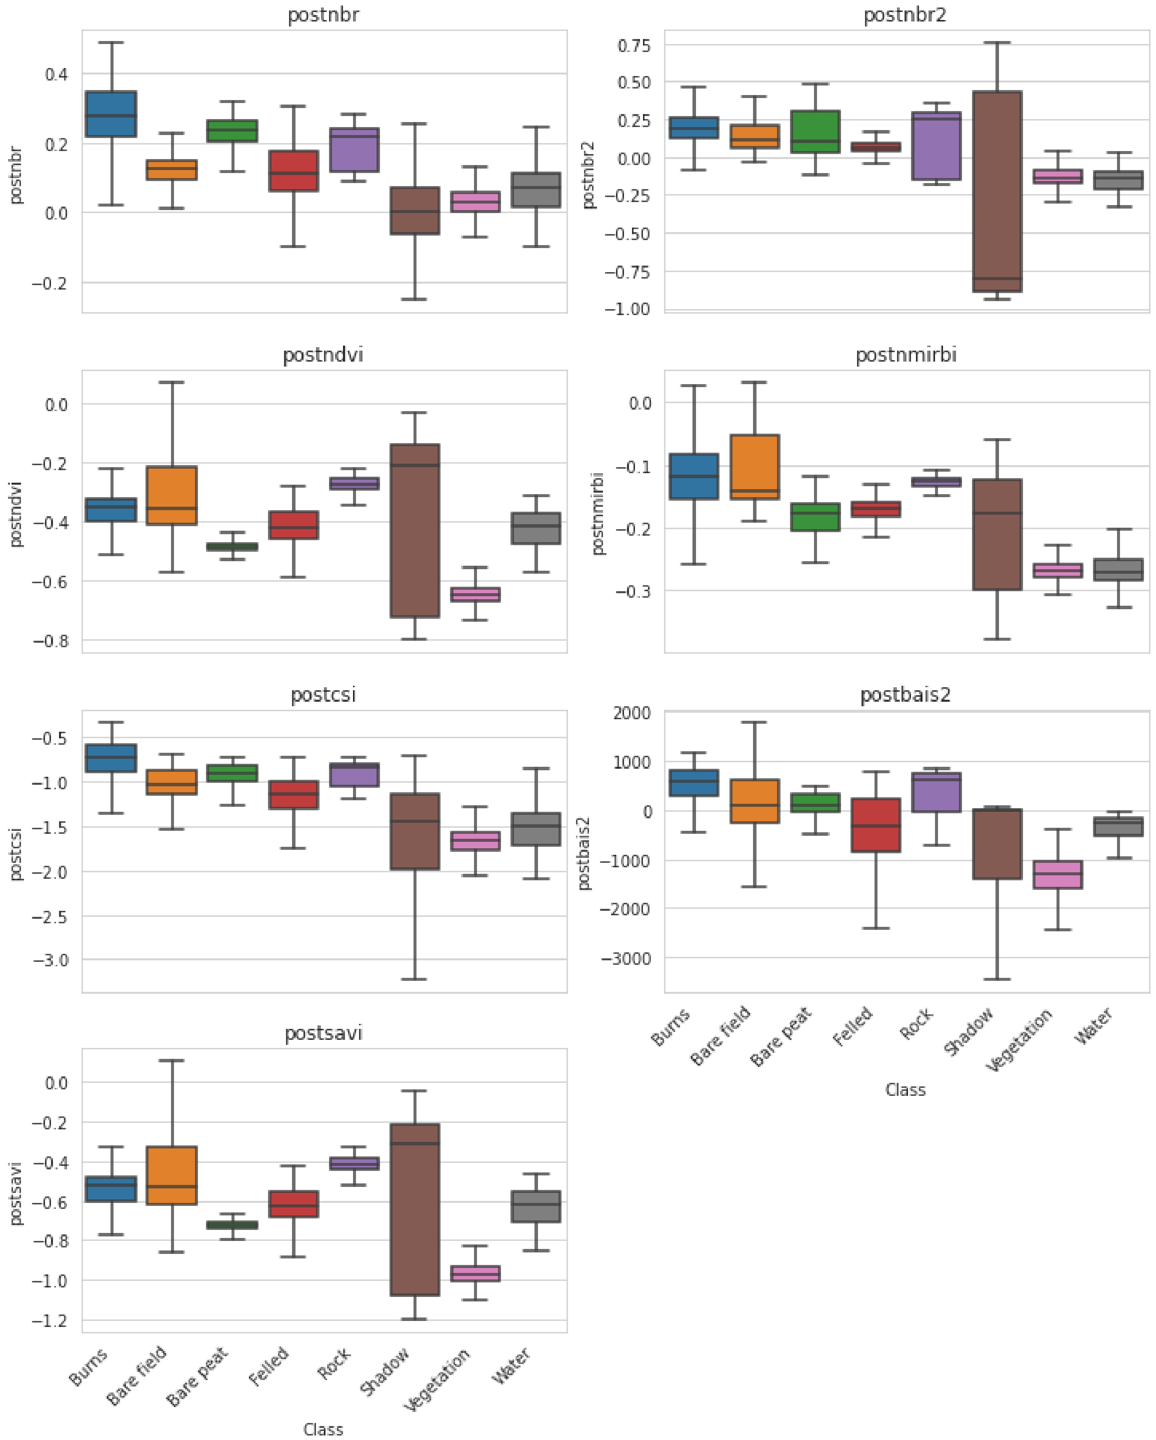

Observations:

- In the post fire image NBR shows the best separability between burns and the other classes, though still some overlap with bare peat. CSI also shows good separability with other classes but there is more overlap between burns and rock and bare field classes.

Skye difference indices and difference classes boxplots

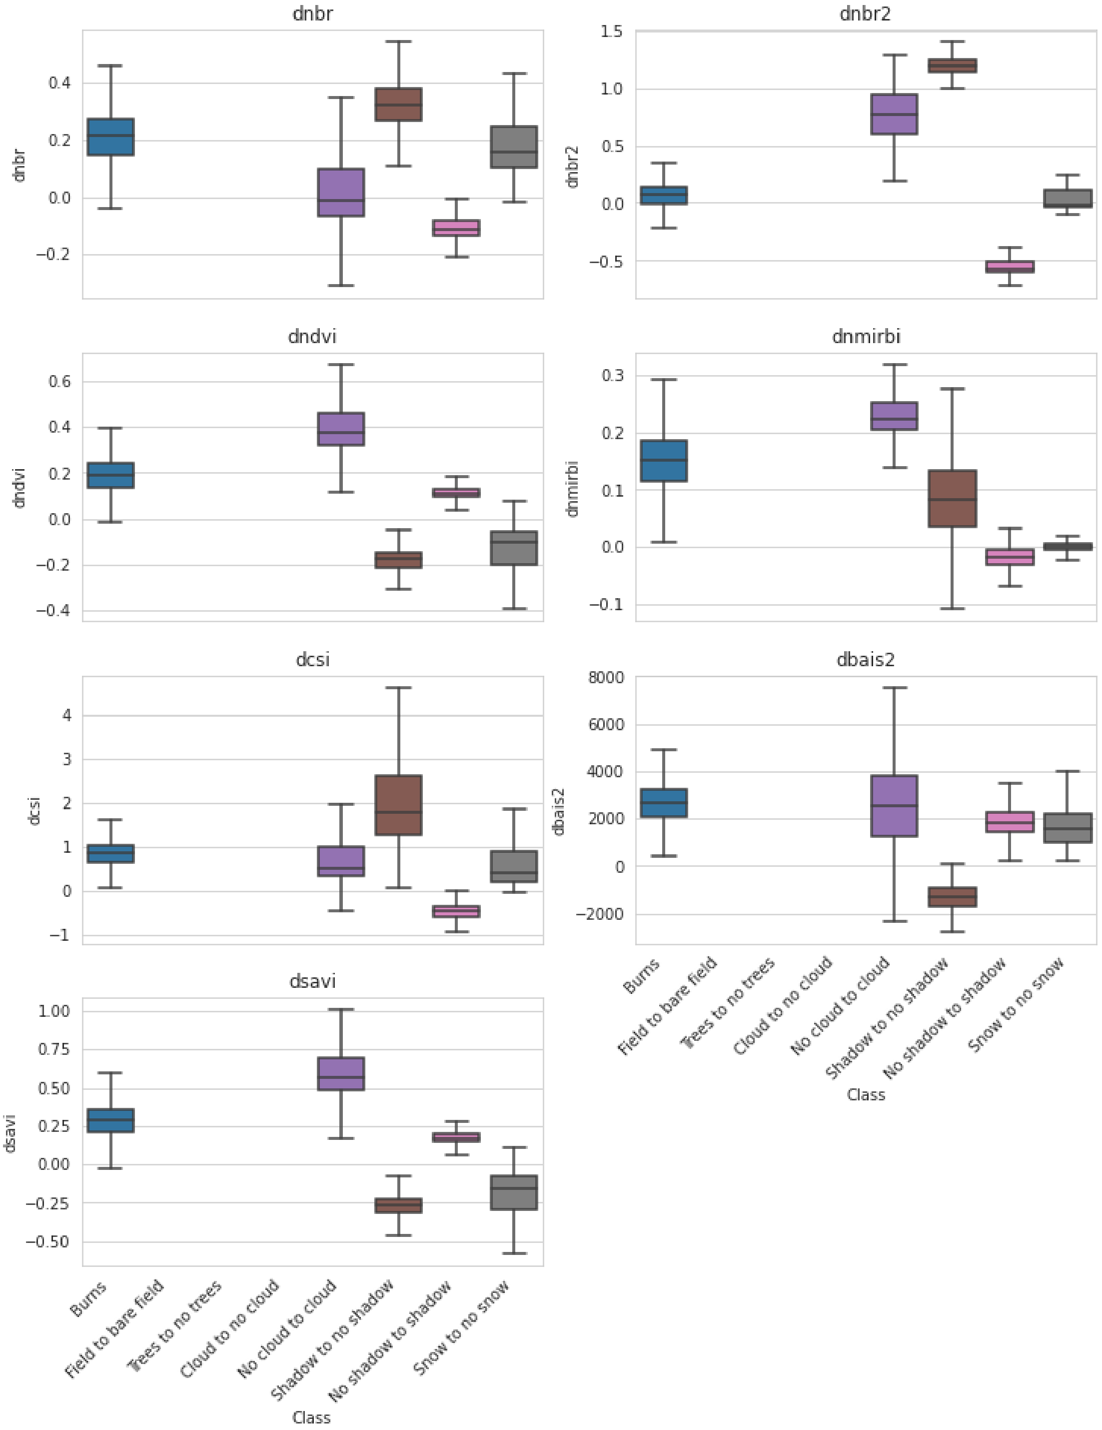

Observations:

- dNBR for burns can get confused with snow to no snow differences.
- The class dNBR2 shows the best separability apart from similar snow to no snow class confusion.

Skye post burn image indices and land cover class boxplots

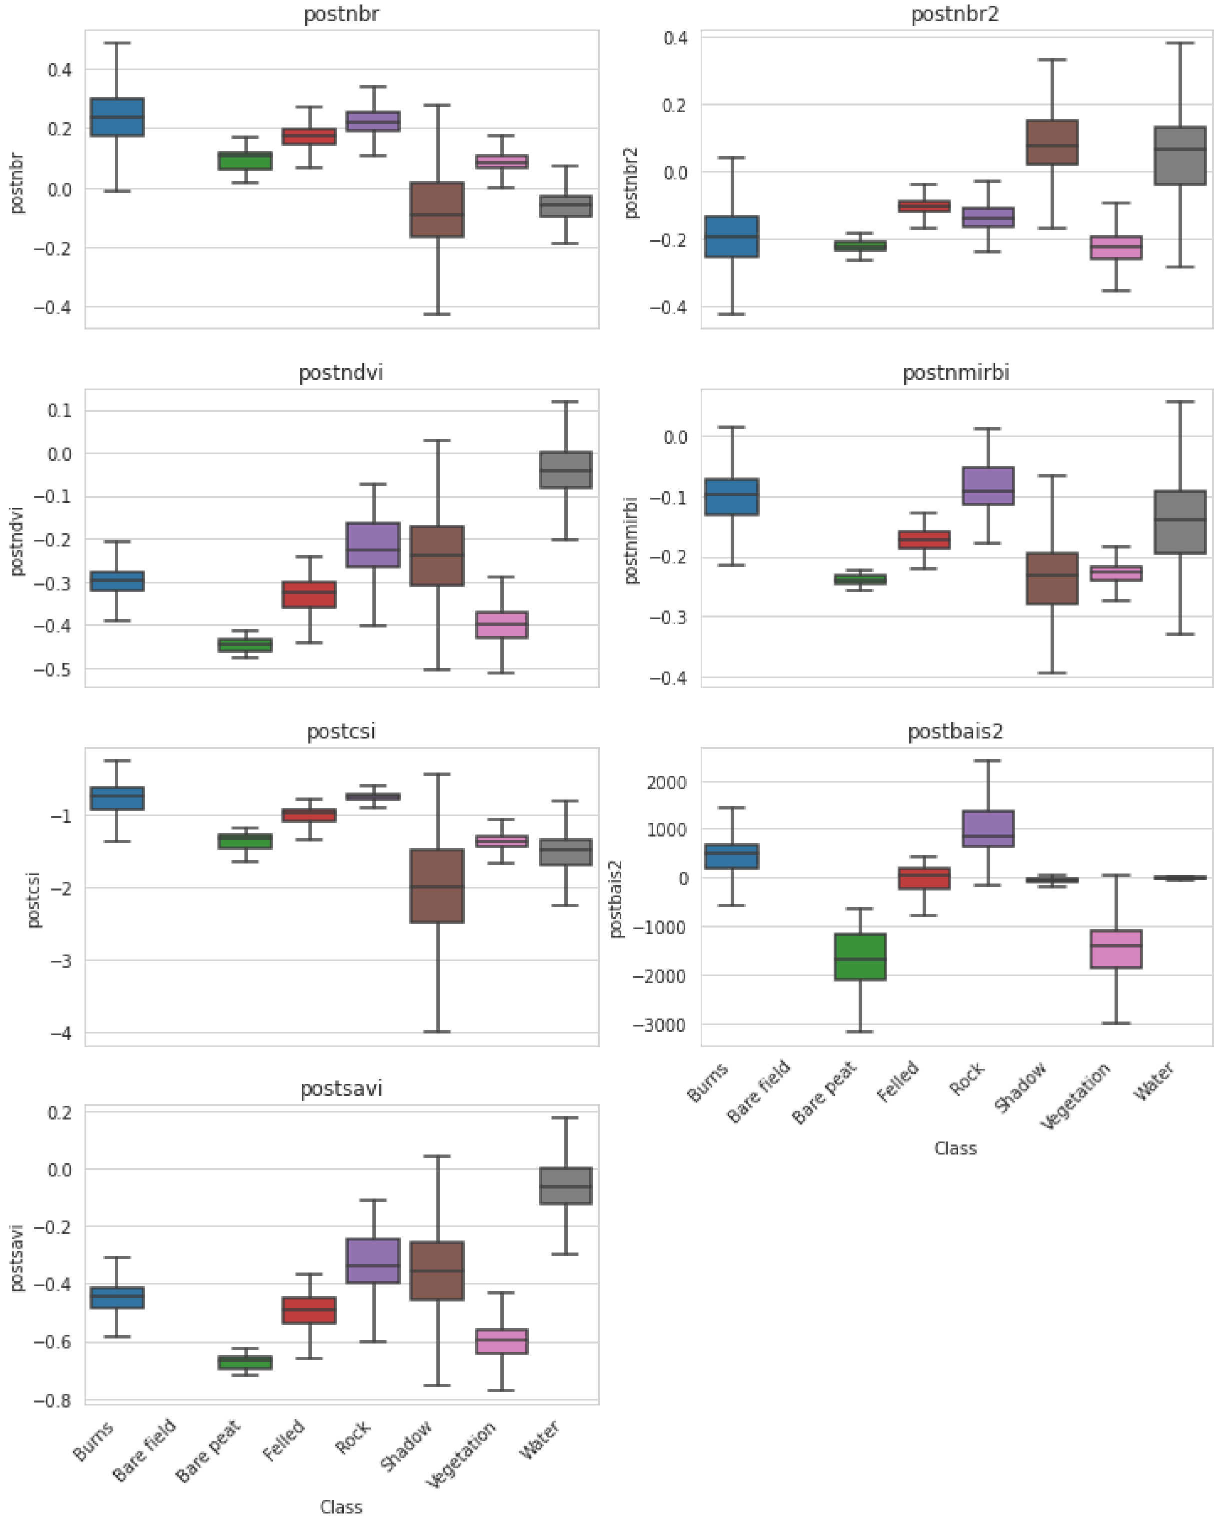

Observations:

- In the post fire image NBR still shows the best separability between burns and the other classes, though there is some overlap with rock and felled areas.

Combined test sites - difference indices and difference classes boxplots

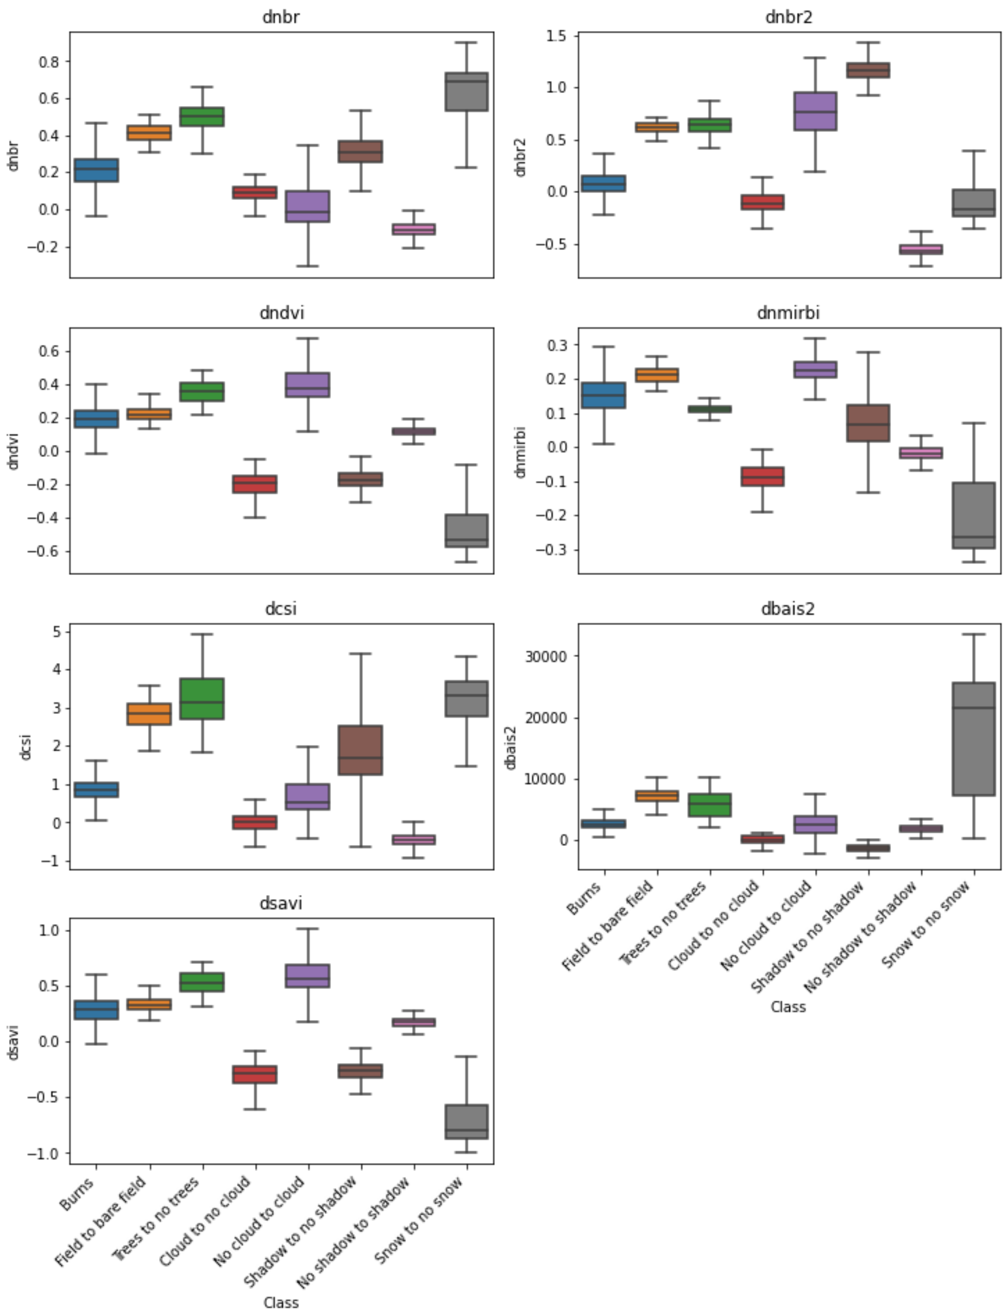

Observations:

- Combining the two datasets will not change the difference with other land cover changes as there were no bare fields or felled areas in the three-week time period of the Skye imagery. dNBR and dNBR2 appear to show increased separability of burns with felled areas and cleared fields compared to the Cairngorms boxplot but this is due to the large areas of burns in Skye bringing down the dNBR values.
- dSAVI still shows reasonable separability with the other classes (especially snow to no snow) though with a little overlap with no cloud to cloud (which did not exist in the Cairngorms image pair).

Combined data post burn image indices and land cover class boxplots

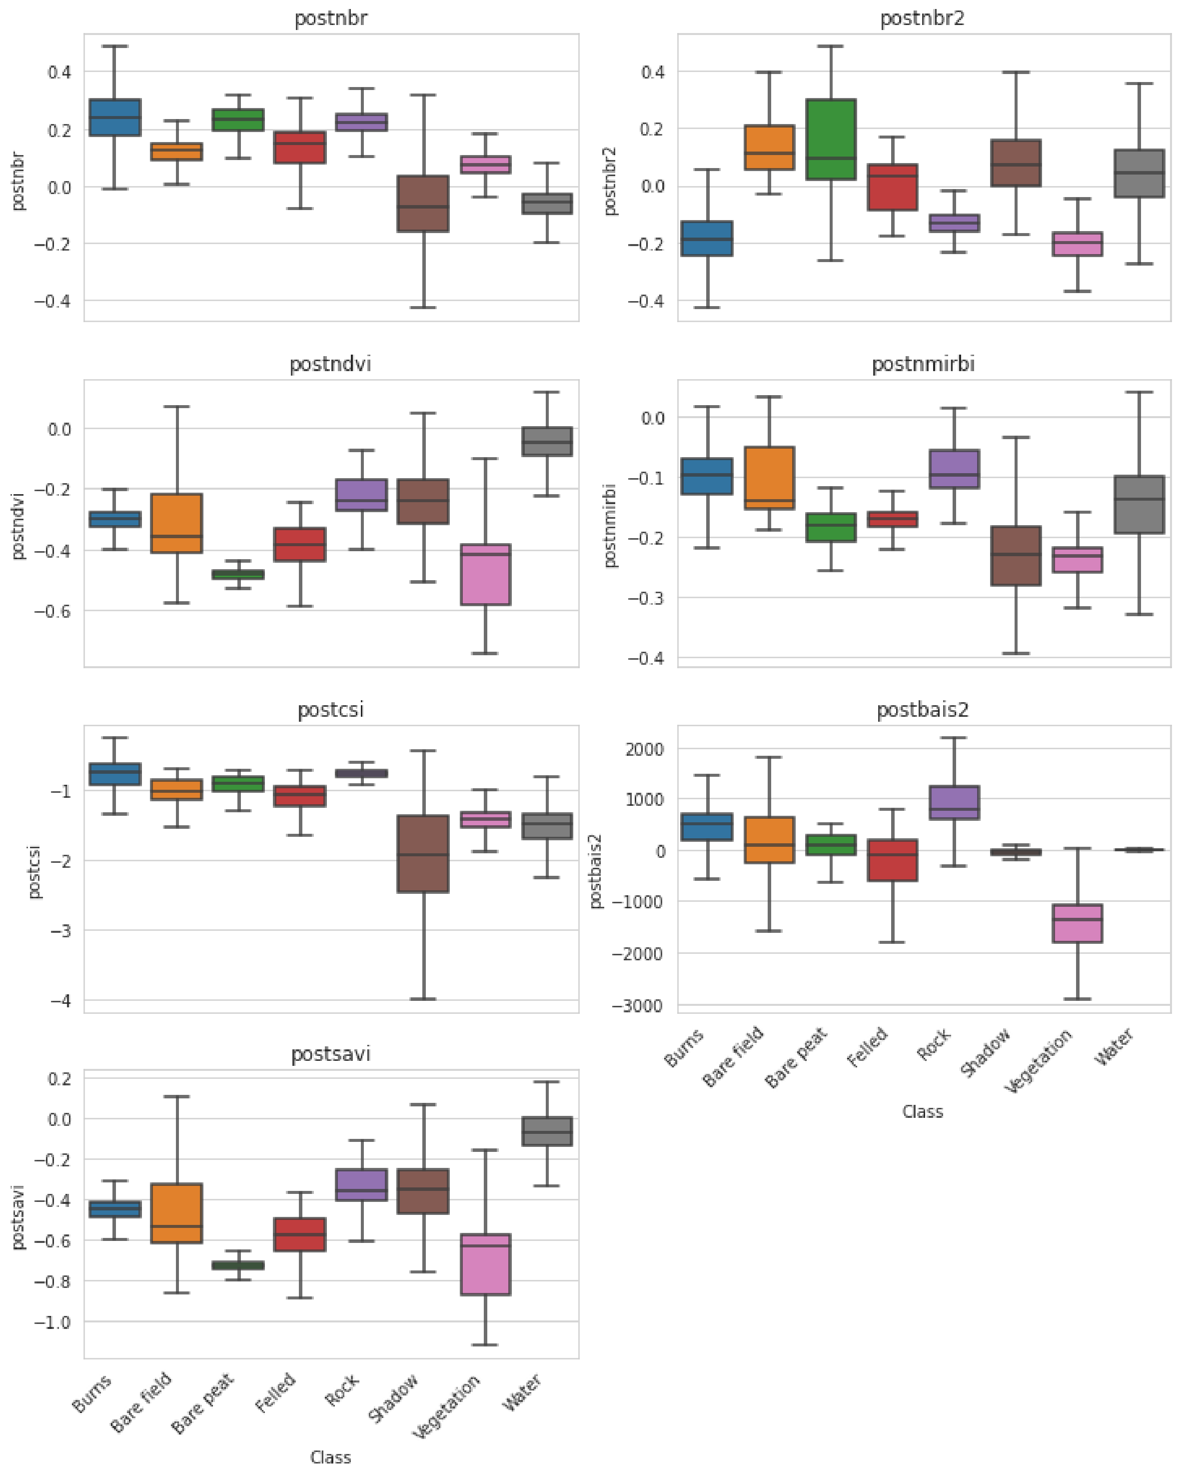

Observations:

- With the two test areas combined the results are a bit confused but NBR or NBR2 show the best overall separability from other classes. However, there is clear overlap with bare peat and rock for NBR and rock and vegetation classes for NBR2.

Appendix 2: Scaling up the process – running in JASMIN

Setting up and running Python code

Virtual machines and a project workspace were set up on JASMIN.

In order to access JASMIN from Windows:

- MobaXterm was downloaded and installed – a new terminal window is opened.
- To login to JASMIN **ssh -A login2.jasmin.ac.uk**
- To login to the JNCC part of JASMIN **ssh -A login.jncc-analysis-sci-m.jasmin.ac.uk**
- To login to the muirburn VM: **ssh 192.168.3.6**
- To activate an environment containing the "current" common software packages (including a modern Python): **module load jaspy/3.7**

<u>On first running the following set-up steps in italics are required</u>
Rasterio and Fiona are not included in Jaspy. The way to enable new packages is to create a virtual environment. Use the following steps:

- *Create a virtual environment to hold packages for this and transfer system site packages:* ***python -m venv --system-site-packages virtenv***
- *Activate the virtual environment -* ***source virtenv/bin/activate***
- *Pip will likely be out of date so update that first, and then install the packages you need e.g. rasterio:* ***pip install --upgrade pip*** *then* ***pip install rasterio*** *and* ***pip install fiona***

- Activate the virtual environment (if not first-time setup) - **source virtenv/bin/activate**

- To create the copy of the code you want to run from the home folder
    - **mkdir code** and **cd code**. If removing a previous version use **rm –r JNCCBurnCalcs**
    - **git clone** https://github.com/Scottish-Natural-Heritage/GIG-JNCC-Muirburn.git
    - Or to clone a branch that is being tested: **git clone --branch duncan_v1 --single-branch** https://github.com/Scottish-Natural-Heritage/GIG-JNCC-Muirburn.git

- To run the code:

    - **cd code/JNCCBurnCalcs**
    - Look at and change any parameters in config.py by using **vim config.py**. Press **i** to insert and make changes, esc to quit inserting, and **:x** to quit vim. In here set the input and output directories for processing.
    - **python operationalcode.py** to run the code.

- By default, the working directory is set to the home directory. Change directory to the group workspace: **cd /gws/nopw/j04/jncc_muirburn**

- Once a session is finished run:
    - **deactivate**
    - **module unload jaspy/3.7**
    - **exit** (3 times)

Other JASMIN information

Using screens to avoid disconnection from MobaXterm

On a Linux terminal, e.g. JASMIN VM or SCI server you can run multiple terminal sessions concurrently, each running independent tasks. You can then switch between these different sessions to see how each job is getting on - great for multi-tasking. You can disconnect from screens and they continue to run, even if you log out of JASMIN entirely.

screen -S my-processing-job -t my-tab-title

This creates a 'screen' with name 'my-processing-job' and title 'my-tab-title' that will appear on the MobaXterm tab

The screen is a multiplexed terminal window. You can create several concurrent ones, each of which can run different jobs in parallel. The key advantage is you can log off the VM and the jobs continue to run.

Once you've created the screen, start the job as normal.

To detach from the screen type **Ctrl+a+d** (hold down Ctrl key and while holding it type 'a' and then 'd').
You can now log off and the session keeps running.
Log back into JASMIN and to the VM you started the job on and type **screen -ls**

This lists currently running screens. The screen name given at the beginning should make the session easy to identify (in case there were several). If the session is still going then type

screen -r
If there is only one session it will re-attach. If there are several then type:

screen -r <session id>

Data storage

Each home directory has a default quota of 100 GB. You can find your current usage by running the following Linux command:

- **pdu -sh /home/users/<username>**

You are only allowed to exceed this limit for a very brief period of time but if you continue to exceed the limit, you will be unable to add any more files or run jobs and will be required to reduce your usage. The quota limit control of 100GB is enforced on the user home directory.

There is a daily incremental and weekly full backup of your home directory. Your home directory is the ONLY storage which is automatically backed up.

- **cd ~** takes you to your home directory

Group workspace structure

We recommend that a sensible directory structure is set up within your GWS and that the following conventions are used within your GWS:

```
<your_gws>/
   users/
      <userid>/   # each user can create their own directory here
   public/        # required if you want to share data via HTTP
   data/
      internal/   # internal/intermediate data
      incoming/   # third-party data brought to the GWS
      output/     # output data generated by project
```

See the GWS etiquette article for more details about GWSs and the [GWS data sharing via HTTP article] for information about use of the public directory.

Appendix 3: Literature review of dynamic thresholding methods

Thresholding was investigated by JNCC with the following conclusion:

In general, where papers talk about the idea of thresholding differently for different areas, they agree that it is necessary to some degree. However, the degree to which thresholds are adapted seems to vary in practice - whether at country level, ecosystem level, scene level or habitat level.

The approaches that seem to re-occur are:

(a) Using static thresholds across all areas (Kolden *et al*. 2012; Verhegghen *et al*. 2016; Bastarrika *et al.* 2011; Hawbaker *et al*. 2017; Goodwin & Collett 2014)
(b) Using thresholds that are static over areas that are ecologically or geographically similar (ecosystems, countries) (Smiraglia *et al*. 2020), or
(c) Using a method that generates thresholds on the fly, meaning they're customised for each situation (Nolde *et al.* 2020; Giglio *et al.* 2009; Roteta *et al*. 2019)

Of these, (b) seems to be the most common approach. There's arguably a halfway point between (a) and (b) as well, consisting of fuzzy-thresholds and multi-spectral agreement indices (Boschetti *et al.* 2010; Stroppiana *et al*. 2012, 2015; Filipponi 2019).

In terms of this project and future work, the different approaches to thresholds need to be tried, and the final method is going to depend on how well they perform. Option (c) seems like the ideal case, but it could take a lot of time to develop. Likewise, (b) is popular and seems to perform well, but would still take time, and there would be an additional step to identify the different areas that need different thresholds. The fuzzy-threshold approach shows promise.

In terms of identifying the values used for thresholds, the approach mostly seems to be data-driven expert opinion (i.e. plotting the distributions of values where fires are known or strongly suspected, choosing appropriate thresholds based on these distributions, trying them, and adjusting accordingly). Even where thresholds are dynamic, there's usually an opinion-driven step to determine the parameters within which these dynamic thresholds should fall.