





# Burn extent and severity mapping by spectral anomaly detection in the Landsat data cube

**L.J. Renzullo**<sup>a</sup> , **S. Tian**<sup>a</sup> , **A.I.J.M. Van Dijk**<sup>a</sup>, **P. Rozas Larraondo**, **M. Yebra**<sup>a</sup> , **F. Yuan**<sup>b</sup> and **N. Mueller**<sup>b</sup> 

<sup>a</sup> Centre for Water and Landscape Dynamics, Fenner School of Environment and Society, The Australian National University, ACT, <sup>b</sup> Digital Earth Australia, Environmental Geoscience Division, Geoscience Australia, ACT

Email: [Luigi.Renzullo@anu.edu.au](mailto:Luigi.Renzullo@anu.edu.au)

**Abstract:** Mapping the extent and severity of bushfires is an important part of post-fire damage assessment and contributes to the fire history of a region. This information in turn is used in estimating the gradual increase in fuel load after the fire and hence is a key variable in anticipating future fire risk. Spectral indices, i.e., linear combinations of multi-spectral bands in remote sensing imagery, are conventionally used to determine burn extent and severity by comparing differences in imagery obtained pre- and post-fire events. For example, the Normalised Burn Ratio (NBR) exploits differences in the relative spectral response in shortwave and near-infrared wavebands to identify areas of burnt vegetation in satellite imagery. Well-known limitations of this differencing approach include its limited consistency and applicability over large areas, the requirement of *a priori* knowledge of where and when a fire occurred, and the need for imagery acquired within a reasonable time before and after the burn so that seasonal changes and recovery effects are minimal.

To address these challenges, we developed a combined burn extent and severity mapping approach that uses the full spectral information in time series of Landsat satellite observations available through the Digital Earth Australia archive. The method is primarily designed for perennial vegetation that does not burn frequently. The principle is to identify spectral anomalies in space and time, i.e., spectra that stand out significantly from the time series. The method quantifies the average spectral response of a pixel using the robust geometric median, which is relatively insensitive to residual atmospheric effects. Deviation of the spectral response from the geometric median is quantified through the cosine distance, a measure based on spectral similarity. Pixels with a distance greater than the equivalent of three standard deviations from the mean (i.e., statistical outliers) are identified as having changed. Subsequently, absolute and relative NBR and cosine distance changes are calculated to identify burns from other possible landcover changes. A subsequent region-growing step improves the classification by contracting pixels with below-threshold evidence of burning. In an optional post-processing step, corroborating data such as fire detections from thermal remote sensing (e.g., Geoscience Australia's Sentinel Hotspots fire detection system) and other ancillary data can be used to improve classification further.

We evaluated the method for several case studies in southern Australia through a comparison with independently derived burn extent maps provided by government agencies. The results show that the fully-automated algorithm developed produces classification results that are commensurate with conventional supervised image classification methods, but with the benefit of being repeatable and fully automated.

**Keywords:** *Burn mapping, spectral anomaly, Landsat, Digital Earth Australia, change detection*

## 1. INTRODUCTION

Mapping the spatial extent and severity of a bushfire event is a common aspect of post-fire damage assessment. It also helps to extend the so-called fire history for an area, which is key information in assessing future fire risk, e.g. through predicting fuel load build-up since the fire. Spectral indices, i.e., linear combinations of multi-spectral bands in remote sensing imagery, have long been used to highlight features of interest in the landscape and to monitor changes in these features through time. For example, the Normalised Burn Ratio (NBR) (Key and Benson, 1999) exploits differences in the relative spectral response in shortwave and near-infrared wavebands to identify areas of burnt vegetation in satellite imagery. The NBR is often used along with vegetation indices (e.g., NDVI) in imagery acquired pre- and post-fire to map burn extent. However, a well-known limitation of the use of spectral indices to map burn extent is their limited consistency and applicability over large areas. For example, contrasting soil types or land uses can have a similar index value to burnt vegetation, thus leading to misclassification. Moreover, index-based approaches to burnt area mapping require *a priori* knowledge of where and when the fire occurred, and imagery acquired within a reasonable time before and after the burn so that the effects of seasonal vegetation change and post-fire recovery are minimal. These all pose challenges when using infrequent, higher-resolution imagery, such as that available from the Landsat series. Conventional remote sensing methods to map burn extent and severity are at least partially supervised and typically require at least some knowledge of the approximate locality and date of fire occurrence, followed by the visual selection of a sequence of imagery preceding and immediately following the event. Satellite data characteristics such as spatial resolution, areal extent, spectral and temporal coverage all play important roles in determining the accuracy with which a burn can be mapped by such an approach (Chuvieco *et al.*, 2019). For example, some remote sensing systems offer a high resolving capability of  $\sim 1$  m, but the data are collected infrequently (e.g., monthly), giving rise to the temporal sampling issues already mentioned. Conversely, MODIS provides more frequent imagery but at lower resolution. This is suitable for large-area burn mapping, and indeed MODIS is used to good effect in the North Australia and Rangelands Fire Information System (NAFI, <http://www.firenorth.org.au/nafi3/>), although the degree of operator supervision and output quality control and associated resources requirements are unknown to us. In any case, the 500-m resolution of MODIS imagery is often insufficient for the detail of mapping required in southern states. In those cases, burn mapping is only done where the occasion calls for it, using a variety of methods (e.g., airborne or on-ground visual surveys, or supervised imagery classification), with accuracy and detail that are extremely variable, and an overall result that is not reproducible. This also means that the resulting fire history is incomplete and of generally unknown accuracy. Overall, Landsat observations appear to provide a suitable balance between spatial resolution, spectral coverage and temporal sampling frequency for most burn area mapping applications (Goodwin and Collett, 2014; Chuvieco *et al.*, 2019).

Our objective was to develop a fully automated algorithm and workflow that could be used to provide a comprehensive and reproducible fire history by making use of the very large amount of Landsat-series imagery contained in Geoscience Australia's Digital Earth Australia database. While the algorithm can be and was applied across the continent, the primary application was for the construction of an up to date fire history for forested environments in southern Australia. Requirements for the workflow to be developed were that it would be (1) *scalable*, so it could be implemented anywhere and everywhere in Australia without operator input; (2) *sensor-agnostic*, so that it could take advantage of the full Landsat series as well as more recent Sentinel-2 observations; (3) *freely available* under an Apache 2 license (i.e. no restrictions on the use and modification of the code); and (4) *documented* in the form of self-documented code, case studies and user guidance.

## 2. DATA AND METHODS

### 2.1. Digital Earth Australia Landsat archive

The data used by the workflow are derived from Geoscience Australia's Digital Earth Australia (DEA) archive of Landsat imagery (Lewis *et al.*, 2016). The DEA holdings span the Landsat series of thematic mapper (TM, including the enhanced ETM+) and operational land imager (OLI) imagery. The data was georeferenced, cloud-masked, atmospherically corrected and resampled to 25-m resolution tiles in Australian Albers projection by Geoscience Australia. The imagery is currently accessible through the National Computational Infrastructure using an instance of the Open Data Cube (<https://www.opendatacube.org/>); a Python-based library that provides a programmatic interface to the data and supports various data selection, manipulation and analysis operations.

## 2.2. Spectral anomaly detection

Our approach to identifying burn events in time series of multispectral Landsat imagery is to examine deviations of the reflectance spectra from the typical or ‘normal’ spectrum for a given pixel. When the deviation of a given spectrum is above a predefined distance (in spectral space) from the norm it is called a *spectral anomaly*, in the statistical sense of the word. The approach therefore requires the calculation of the normal spectrum and the choice of distance measure. While it has been applied in the following it may be used to identify any landcover change. We illustrate the approach with example spectra in Figure 1.

Figure 1 shows the Landsat reflectance spectra for a vegetated area before and after a bush fire. Before the fire (blue curve) we see a spectrum that has the spectral shape typical of vegetation, namely the blue (0.4  $\mu\text{m}$ ) green (0.55  $\mu\text{m}$ ) and red (0.65  $\mu\text{m}$ ) bands delineating the *green peak*; the shape rise in reflectance along the red-edge to the near-infrared (NIR) plateau (0.8  $\mu\text{m}$ ), and a drop in reflectance in the shortwave bands 1 (1.6  $\mu\text{m}$ ) and 2 (2.2  $\mu\text{m}$ ). While the shape of the spectrum is consistent with that expected for a vegetated target, the values are somewhat higher than expected. In this case, the vegetated pixel was contaminated with light haze (thin cloud or smoke). After the fire (orange curve) the characteristic spectral response of vegetation has been replaced by a monotonically increasing reflectance curve that is typical of bare ground. Land clearing rather than burning produces a similar spectral response. A normal spectral curve for this location was calculated using the *geometric median* (Roberts *et al.*, 2017). The geometric median (GM) is displayed as the grey curve in Fig. 1. This multivariate statistic is representative of the reflectance spectra for the typical state of the target, uninfluenced by potential cloud contamination, sensor artefacts, or infrequent cover disturbances. As with univariate statistics, the multivariate GM is less sensitive to outliers than the arithmetic mean. In this work, the GM, denoted  $\mathbf{x}^{GM}$ , served as the normal reflectance spectrum for our time series.

The next step in our approach is to estimate deviation of the spectrum for each date in the time series from the GM. Standard distance metrics used in multivariate statistics for classification, or discriminant analysis, are the Euclidean or Mahalanobis distances (Campbell, 1980). Both provide sum-of-squared differences between a multivariate variable and a reference (usually the multivariate mean). However, these measures do not consider the similarity between the variable and the reference. In our case, this means that the two very different looking spectra can have very similar Euclidean distances despite appearing very different. For example, the pre- and post-burn reflectance spectra in Fig. 1 have Euclidean distances of 0.140 and 0.165 from  $\mathbf{x}^{GM}$ , respectively. A consequence would be that the two spectra would be misclassified as unburnt vegetation.


An alternative measure of distance is based on the cosine similarity of reflectance spectra (Roberts *et al.*, 2018). The similarity is expressed by the angle,  $\theta$ , between two vectors in multi-dimensional space. For a sample spectrum,  $\mathbf{x}_s$ , and the GM,  $\mathbf{x}^{GM}$ , it is calculated as

$$\cos(\theta) = \frac{\mathbf{x}_s \cdot \mathbf{x}^{GM}}{\|\mathbf{x}_s\|_2 \|\mathbf{x}^{GM}\|_2}, \quad (1)$$

where  $\|\cdot\|_2$  is L2-norm, or the square root of the sum of square elements of  $\mathbf{x}$ , and the cosine distance between  $\mathbf{x}_s$  and  $\mathbf{x}^{GM}$  is defined by

$$d_{\cos}(\mathbf{x}_s, \mathbf{x}^{GM}) = 1 - \cos(\theta) \quad (2)$$


A cosine distance of 0 indicates perfect similarity (i.e., identity) between two spectra, while anything else indicates a dissimilarity of some degree. In the case of the spectra in Fig. 1, cosine distance values between pre- and post-fire spectra from the GM are  $d_{\cos} = 0.005$  and 0.241, respectively.



**Figure 1.** Landsat reflectance spectra for vegetation for a date before (blue) and shortly after (orange) a burn event, along with the Geometric Median (grey) of the spectra whole time series.

The GM and cosine distance calculation are applied per pixel using a multi-year time series for a defined geographic extent extracted from the Data Cube. For each pixel, a data matrix,  $\mathbf{X}$ , is constructed from  $N_s$  cloud-free  $p$ -dimensional multispectral observations  $\mathbf{x}_s$ , where in our case,  $p = 5$  corresponding to the 5 solar reflective Landsat TM bands. We calculate the GM,  $\mathbf{x}^{GM}$  for the time series and the cosine distance using Eq. (2) for each date  $t_k$ ,  $k = 1, \dots, N_s$ . From the time series of  $d_{\cos}(t_k)$  we calculate the 75<sup>th</sup> percentile ( $Q3$ ) and interquartile range (IQR) of values to define the critical distance threshold,  $d_0 = Q3 + 1.5 \text{ IQR}$ .

A spectral anomaly is identified for any date  $t_k$  in the time series where  $d_{\cos}(t_k) > d_0$ . A cosine distance above the critical threshold represents a statistically significant deviation from  $\mathbf{x}^{GM}$ . However, due to an occasionally high likelihood of cloud contamination (despite preliminary screening), we include the additional condition that a spectrum  $\mathbf{x}_s(t_k)$  is an anomaly if and only if its nearest neighbours  $\mathbf{x}_s(t_{k-1})$  and  $\mathbf{x}_s(t_{k+1})$  are also identified as anomalies in the time series. An example time series of cosine distances for a pixel is shown in Figure 2.



**Figure 2.** Time series of cosine distance calculated from Landsat reflectance spectra and the geometric median,  $\mathbf{x}^{GM}$ , for a single pixel. Anomalies in the time series identified as points above the critical distance threshold  $d_0$ .

An anomalous period is defined by a temporally persistent statistically significant deviation of the spectra from the GM. Such a period is highlighted in Fig. 2. The start and end times,  $t_{start}$  and  $t_{end}$ , define the duration of the anomalous period as,  $duration = t_{end} - t_{start}$ . This represents the passage of

time from when the disturbance (in our case fire) was first detected, to the time when the affected area appears to have recovered, and the spectral response of the land cover has returned to typical conditions and no longer is considered a statistical outlier. A related measure is what is defined here as burn *severity* as,

$$severity = \int_{t_{start}}^{t_{end}} d_{\cos}(t) - d_0 dt, \quad (3)$$

which we use as a measure of the degree of disturbance.

Using only cosine distances would often identify seasonal variation in grassland areas. An effective filter to eliminate these patterns in the cosine distance maps was the introduction of the NBR indices. The Euclidean distance in NBR space ( $d_{nbr}$ ) between the NBR time series and the median NBR ( $NBR^{GM}$ ) calculated from the geometric median was calculated as,

$$d_{nbr} = \sqrt{\sum_{k=1}^N (NBR_{tk} - NBR^{GM})^2} \quad (4)$$

where  $NBR_{tk}$  is the NBR indices calculated in date  $t_k$ . The threshold of the outliers is determined as the same way as cosine distance using  $Q3$  and IQR of  $d_{nbr}$  values. Thus, the anomalies in the time series are identified as points above the critical distance threshold for both cosine distance and NBR distance. The final estimate of *severity* is calculated from the overlap in anomalous periods from both Eq (3) and Eq (4).

### 2.3. Data augmentation with ancillary data and analysis

The spectral anomaly detection step above will identify any land cover disturbance that results in a statistically significant change in the spectral reflectance response of land cover beyond the typical state characterised by the geometric median spectrum,  $\mathbf{x}^{GM}$ . To determine whether the change in spectral response is likely to be fire-related, we include a step to test for consistency with the expected change in the NBR. Erroneous small-scale mapping errors can occur in the results, e.g., as a consequence of remaining cloud edge shadows. Conversely, burn extent may visually extend beyond the area classified as ‘severely’ burned (i.e., statistical outliers beyond  $d_0$  threshold). Region-growing, i.e. expanding the burn area extent by lowering the threshold,  $0.67 d_0$  produced good results, augments the initial ‘severe’ burnt area mapping to produce ‘moderate’ burn extent. These ‘moderately’ burnt pixels will include further areas that do not qualify as outliers but do show a substantial decrease in NBR and are adjoining pixels detected as burns. In an optional post-processing step, corroborating data such as fire detections from GA’s thermal remote sensing fire detecting system Sentinel Hotspots (Geoscience Australia, 2014) or from agency incident report archives can be used to improve classification and attribution further.

## 2.4. Verification data

To evaluate the performance of the spectral anomaly detection method, we compiled digitised maps of burn extent and sometimes severity from various sources. Generally, the mapping was done by state agencies involved in fire management and derived by the variety of means, including surveys and supervised image classification or manual digitisation.

## 2.5. Workflow software and Outputs

The entire workflow described in this paper was developed in Python. A ‘validation toolbox’ was also developed, containing a library of functions to import reference burn mapping data (e.g., as vector files) and to calculate statistics of agreement based on true and false positives and negatives. Worked examples of both mapping and validation were developed in Python notebooks, explaining and illustrating the code as well as the underlying theory. All code is available under a free and open Apache-2.0 license via <https://github.com/GeoscienceAustralia/burn-mapping>.

All output variables are stored in a single, self-described NetCDF-formatted file. Variables include:

- *StartDate* – the detected start-date of severe and moderate burns
- *Duration* – the duration of surface change due to fire
- *Severity* – the severity of surface change due to the bushfire
- *Severe* – binary map of severely burnt pixels
- *Moderate* – binary map of both the moderately and severely burnt pixels
- *Corroborate* – binary map representing corroborating evidence (by default the location of fires detected by GA’s Sentinel Hotspots system using MODIS data only with a 4-km buffer)
- *Cleaned* – starting month of detected burns after post-filtering using corroborate evidence

## 3. RESULTS

We applied the spectral anomaly detection method to a number of scenes in the Landsat archive guided by the verification data we had available at the time of our investigation. We show the mapped burn area extent through maps of severity (Eq. 4) and duration for select case studies in Section 3.1. A comparison of mapped with verification data are presented in in Section 3.2.

### 3.1. Burn area mapping

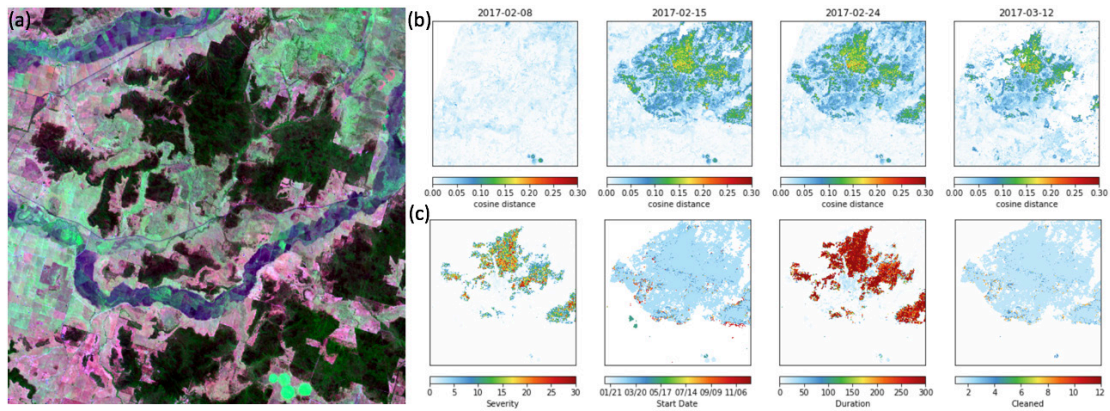
The application of the spectral anomaly detection approach for burn area mapping is illustrated in Figure 3 for 25 km x 25 km areas around the Sir Ivan fire, NSW. Using Landsat-8 OLI spectra from DEA we calculated the geometric median,  $\mathbf{x}^{GM}$ , for data in the period 2013-2016. This is displayed in Fig. 3a as a false colour composite (rgb=b5,4,2). The next step was to calculate the cosine distances,  $d_{cos}$  (Eq. 2), for each of the dates in the extended time series 2013 – 2017. These are displayed in Fig. 3b for 1 date before the bushfire and 3 dates after.

The severity (Eq. 3), start date, duration and cleaned burn extent is displayed in the panels of Fig. 3c. The severity suggests that much of the forested areas (dark green area to the north of the scene, Fig. 3a) was burned in the fire. The Start Date plot shows that much of the scene, including the grassland areas was burnt in Feb 2017. The duration (in days) suggests that much of the grassland areas recovered in ~month, while the forest areas were still affected > 1 year from the fire start date. To minimize the possible attribution of non-fire related change to the estimated burn extent, we applied a post-filtering procedure to the moderate burn extent based on corroborating evidence from GA Hotspots data. A one month buffer before the mapping year was used when extracting the Hotspots data to avoid missing burns at the end of previous year.

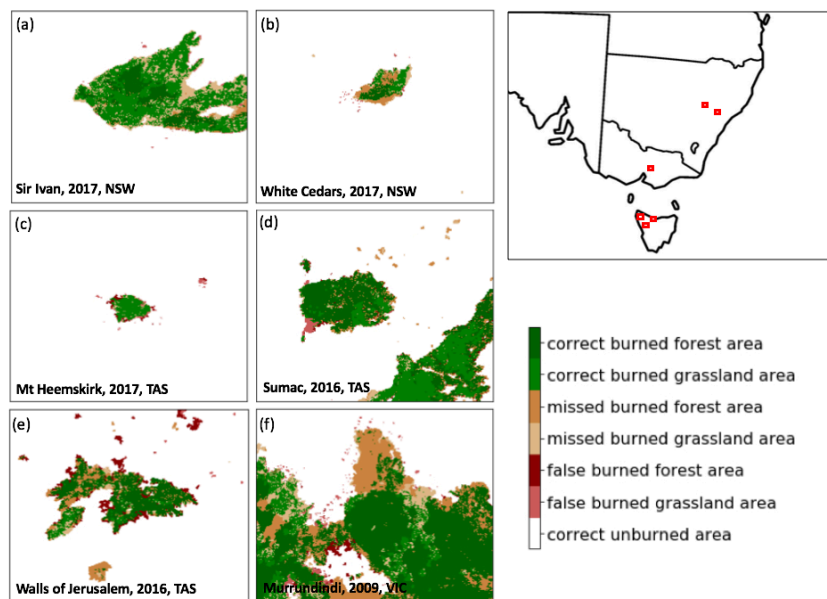
### 3.2. Burn area evaluation

Our evaluation of the method was based on comparisons between estimated burn area and those obtained from the compiled verification data (Section 2.3). Categorical statistics were derived from the method’s ability to identify burned area for forest and grassland across the case studies (Figure 4, Table 1). For most studies sites, burned forest areas were identified with more than 73% accuracy, and as high as 97% for Sir Ivan fires (Fig. 4a). The exception, with 59% accuracy, was the Black Saturday fires of 2009 (Fig. 4f) where the method appears to have missed forested areas around Murrundindi, Vic. However, a visual inspection of pre- and post-fire imagery did not indicate the area to have been affected, despite the verification data (likely from on-ground assessment) indicating that it had. The method was able to correctly identify grassland fires with a better than

67% success rate overall. Accuracy of the method’s mapping of grassland fires for the Walls of Jerusalem, Tas. (Fig. 4e) was 34%. However inspection of the verification data in this instance suggested far more area was burned than we could corroborate through manual visual assessment. In terms of false detection the method performed very well, with errors generally less than 1%. The exception was the Walls of Jerusalem where the method indicated burnt forest area that was not indicated in the verification data.



**Figure 3.** Example burn area mapping using spectral anomaly detection for Sir Ivan, NSW: (a) geometric median calculated for the period 2013-2016 and displayed as a false colour rgb; panel (b) shows the cosine distances calculated for dates before and after bushfire in Feb 2017; panel (c) shows the severity, start date duration and cleaned burn extent (as defined in Section 2.2)



**Figure 4.** Categorical evaluation of the burned areas for case studies areas against verification data.

#### 4. DISCUSSION AND CONCLUSION

Our objective was to develop a fully automated algorithm and workflow that could be used to provide a comprehensive and reproducible fire history by making use of the very large amount of Landsat-series imagery contained in Geoscience Australia’s Digital Earth Australia database.

Evaluation so far has demonstrated that our fully automated workflow produces burn extent and severity mapping results that agree with visual assessment of the imagery. Calculating formal classification performance statistics using the compiled fire extent mapping from various state agencies as a reference was hamstrung by data quality and inconsistent conceptual definitions used. Visual assessment indicated that the accuracy and spatial detail of available mapping were extremely variable, presumably reflecting the requirements for the intended use. Conceptual issues occurred in some cases where the mapping otherwise appeared of high quality. For example, commonly the reference mapping corresponded to the outer perimeter

**Table 1.** Evaluation summary of the case studies of this investigation. Year of burn is provided as well as false positive (FP) and true positive (TP) detection rates for forest and non-forest areas.

Region name	State	latitude	longitude	year	FP/TP forest	FP/TP Non-forest
Walls Of Jerusalem	TAS	-41.69	146.25	2016	0.010, 0.759	0.000, 0.342
White Cedars	NSW	-32.47	149.89	2017	0.001, 0.726	0.005, 0.683
Sumac	TAS	-41.14	144.85	2016	0.004, 0.770	0.000, 0.797
Sir Ivan	NSW	-32.06	149.66	2017	0.000, 0.968	0.007, 0.673
Murrundindi	VIC	-37.38	145.59	2009	0.004, 0.589	0.001, 0.730
Mt Heemskirk	TAS	-41.87	145.26	2017	0.001, 0.758	0.000, 0.918

of the largest contiguous area burned, whereas our automated analysis can indicate unburnt areas within that perimeter, as well as unconnected fire areas outside it (e.g., as the result of spotting by ember attack). All of these caveats made the interpretation of formal classification statistics challenging. Nonetheless, visual interpretation indicates that typically, our mapping results were of a quality that is commensurate or better than that available from mapping agencies based on current operational methods. Furthermore, the workflow developed allows for the augmentation of the mapping by integrating corroborating evidence from, e.g., satellite-based fire detection systems and agency incident reports.

The main advantage of the developed method over current operational mapping techniques is the opportunity for automated burn mapping over large areas. Code performance has been tested at national scale, and this demonstrated that routine (e.g., annual) national-scale mapping at 25-m resolution is entirely feasible. Such a national product should help produce a more comprehensive fire history for use in fire management but also for government and private industry to assess fire risk. For mapping of individual, high-impact fire events, at the very least, the workflow developed can be used as a starting point for further improvement and augmentation, potentially presenting considerable productivity gains.

## ACKNOWLEDGMENTS

This research was funded by Geoscience Australia through Digital Earth Australia, and the ANU authors gratefully acknowledge their support. The Australasian Fire and Emergency Service Authorities Council (AFAC) facilitated the communication with individuals who assisted in the collation of reference burn mapping data: Warwick Hehir (NSW Rural Fire Service), Kate Waghorn (ACT Emergency Services Agency), Simeon Telfer (Dept. Environment and Water, SA), and Carolyn McMillan (Landgate, WA), Lindsay Mitchel (Dept. of Primary Industries, Parks, Water and Environment, Tasmania), Naomi Withers (Dept. Environment, Land, Water & Planning, Victoria). The assistance of all involved is gratefully acknowledged.

## REFERENCES

- Campbell, N.A. (1980), Robust procedures in multivariate analysis I: Robust covariance estimation, *Applied Statistics*, 29(3), pp. 231-237.
- Chuvieco, E., Mouillot, F., van der Werf, G. R., San Miguel, J., Tanasse, M., Koutsias, N., Garcia, M, Yebra, M., Padilla, M., Gitas, I., Heil, A., Hawbaker, T.J., Giglio, L. (2019). Historical background and current developments for mapping burned area from satellite Earth observation. *Remote Sensing of Environment*, 225(March), 45–64. <https://doi.org/10.1016/j.rse.2019.02.013>
- Geoscience Australia 2014. Sentinel Hotspot Product Description Document V1.2, code D2014-145826, Geocat Reference 70869
- Goodwin, N.R. and 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*, 148, 206-221.
- Key., C.H. and Benson, N.C. (1999). The normalized Burn Ration (NBR): A Landsat TM radiometric measure of burn severity. United States Geological Survey, Northern Rock Mountain Science Center. (Bozeman, MT)
- Lewis, A., Lymburner, L., Purss, M. B. J., Brooke, B., et al. (2016). Rapid, high-resolution detection of environmental change over continental scales from satellite data – the Earth Observation Data Cube. *International Journal of Digital Earth*, 9(1), 106–111. <https://doi.org/10.1080/17538947.2015.1111952>
- Roberts, D., Mueller, N., McIntyre, A. (2017). High-dimensional pixel composites from earth observation time series. *IEEE Transaction on Geoscience and Remote Sensing* 55(11), 6254-6264, doi: 10.1109/TGRS.2017.2723896
- Roberts, D., Dunn, B., Mueller, N. (2018). Open data cube products using high-dimensional statistics of time series. In *proc.*, IEEE Geoscience and Remote Sensing Symposium, Valencia, Spain July 2018.