Integration of Remote Sensing Data with Bushfire Prediction Models

C. Miller [a], **J. Hilton** [a], **V. Lemiale** [a], **C. Huston** [a], **A.L. Sullivan** [b] **and M. Prakash** [a]

[a]*Digital Productivity Flagship CSIRO, Clayton, Victoria, 3169*
[b]*Land and Water Flagship CSIRO, Australian Capital Territory, 2601*
Email: *claire.miller@csiro.au*

Abstract: The ability to predict the spread of a bushfire is important for operational management during bushfire events as well as risk assessment in vulnerable regions. Models for the rate of spread of a bushfire typically require a number of input parameters to operate, including topographic information, meteorological conditions and information on fuel conditions and variability. The predictive ability of a bushfire spread model is limited by the accuracy of these input parameters. Accurate national topographic information exists and detailed meteorological information is readily forecast. However, fuel condition information is often specific to a particular fire behaviour model for a particular vegetation type, such as forest or grassland, and generally requires detailed surveys of disparate sources to prepare suitable spatially explicit maps. This preparation can be both time-consuming and cost-prohibitive. Furthermore, these fuel parameters often require a pre-processing step before they can be used in a predictive model.

We demonstrate an approach to determining vegetation and fuel classification from remote sensing data such that the data can be directly integrated with a computational bushfire perimeter propagation model. We use freely available Landsat 8 spectral data for the input remotely sensed data set. The fuel classification uses an unsupervised K-Means algorithm allowing forest, grassland, and water regions to be classified. Two extensions to the method are also trialled for potential improvement to results; the addition of a topography data set as an additional input to the spectral data sets, and a principal component analysis run prior to the K-Means classification.

The bushfire perimeter propagation model uses the CSIRO Spark framework for predicting the spread of fires across the landscape and is already integrated with the open-source OpenCV image analysis tool-kit. The K-Means algorithm is also implemented in OpenCV such that the data can be processed within the same framework. We demonstrate the applicability of the approach to a bushfire prone region in South East Australia. Our system can be used for any region within Australia and potentially worldwide, allowing up-to-date fuel classification and input data to be automatically generated and used in predictive models.

Using these methods, it was determined that, out of all the spectral band combinations, using the coastal aerosol, green, and SWIR 1 produced the best classifications. The addition of topography information did not appear to improve the method, and neither did the principal component analysis on its own. However, using the first component of the principal component analysis in combination with an input band containing the local variance of this component gave a slight increase in the calculated accuracy.

There is the potential for fuel information to be similarly generated from the spectral data, as well as full colour images of the terrain. Input data layers produced using this process would all share the same spatial characteristics, removing the need for any pre-processing such as cleaning, re-projection or alignments steps. Furthermore, the layers can be generated immediately from the most up-to-date data sets available, ensuring currency in the modelling process.

Keywords: Remote sensing, satellite imagery, bushfire modelling, k-means, unsupervised classification

1 INTRODUCTION

The prevailing fuel properties of vegetation are fundamental to the behaviour and spread of a bushfire (Sullivan et al., 2012). Consequently, it is necessary to have high quality, spatially explicit information about those fuel properties in order to accurately predict bushfire spread over the landscape. However, such information is difficult and costly to collect as it generally requires *in situ* physical assessment of fuels (e.g. Gould et al. (2011)). In Australia, particularly, it requires considerable effort given the large areas at risk combined with the highly dynamical nature of the vegetation and use of the land. Where spatially-explicit fuel information does exist, it is often collected sporadically due to the expense and thus out of date when it is needed. Furthermore, continued advances in bushfire behaviour models can mean that previously collected fuel information is no longer sufficient or appropriate to run a new or revised bushfire behaviour model (Cruz et al., 2015).

As satellite data becomes more accessible at higher qualities and resolutions, acquiring up-to-date detailed fuel information quickly and efficiently is increasingly feasible (e.g. Keane et al. (2001)). This is chiefly due to the rapid development in sensing platforms and high-speed delivery of raw data. However, converting this raw data into usable information requires a high level of processing. In the field of bushfire management, much work has been done, for example, in developing specific data processing methods for determining grassland curing state (e.g. Newnham et al. (2011)). Frequently the products of these methods are static maps of fuel state applicable for a given period, but often the map is much older than the most recently available data due to the time and effort required to process it. Consequently, predictions of likely fire spread can be erroneous because they are not using up-to-date information. A method of providing rapid processing of raw remotely-sensed data into timely fuel information for provision to bushfire spread prediction is therefore required.

Before specific fuel attributes can be calculated it is necessary to know vegetation type, as the attributes required are vegetation-dependent. The first step towards providing fuel information directly through remote sensing is therefore the classification of the data into vegetation types. Image classification has recently started allowing the expansion of study areas for hazard research for both pre-event risk and post-event impact analysis. Very high resolution satellite imagery allowed Thouret et al. (2014) to obtain infrastructure information for a large area in Peru in order to assess a city's vulnerability to flash flooding and earthquakes. The use of image classification to provide fuel type classifications has been tried for specific regions in countries internationally, such as Riaño et al. (2002) in Spain and Varga and Asner (2008) in Hawaii. We are yet to see this directly incorporated into a fire model however, particularly in an Australian environment.

Various methods can be used to perform image classification, with the optimal method highly application dependent. Features to consider when choosing a method include allowable computation time, available training samples, contrast between classes, data resolution, available spectral bands, and existing auxiliary data. Following Lu and Weng (2007) we categorize the methods as per-pixel, contextual-based, knowledge-based, sub-pixel or per-field. It is also possible to combine methods in what is termed a combinative approach.

The most straightforward and frequently used classification methods are those in the per-pixel category. Per-pixel algorithms assume no spatial relationship between pixels. The more commonly used algorithms in the per-pixel category include maximum likelihood (ML), neural network (NN), decision tree (DT) and support vector machines (SVM). Of these four, ML is the simplest to implement, however a study comparing ML, NN, and DT algorithms for land cover classification in England by Pal and Mather (2003) clearly found ML to be the worst performing and DT comparable to NN in quality but better in terms of speed. Alonso-Benito et al. (2013) performed a similar comparison between SVM, ML and NN for mapping forest fuels on Tenerife Island, Spain and found SVM methods produced the highest accuracies.

The above per-pixel methods can be referred to as supervised algorithms. This means the algorithm is trained on a set of known points and then run using the set of parameters determined by this training set. It is also possible to have unsupervised classification algorithms which create classes purely based on the distributions of the input data. Two such algorithms are ISODATA and K-Means (KM) clustering. These methods are provided with a set of input parameters and iteratively partition data into clusters by determining appropriate cluster centroids. Although these algorithms do not tend to appear in the fuel classification literature, they are well established methods (Wu et al., 2007) for situations with insufficient training data.

Contextual- and knowledge- based methods use additional levels of information to the spectral data. In the case of contextual-based this is further information about the data itself, such as local variation in values surrounding the pixel. In the case of knowledge-based it is taken from auxiliary data that could influence classes, such as topography. As an example, Riaño et al. (2002) found the addition of a texture map and topography improved their supervised ML classifications when generating fuel maps for Spain from Landsat.
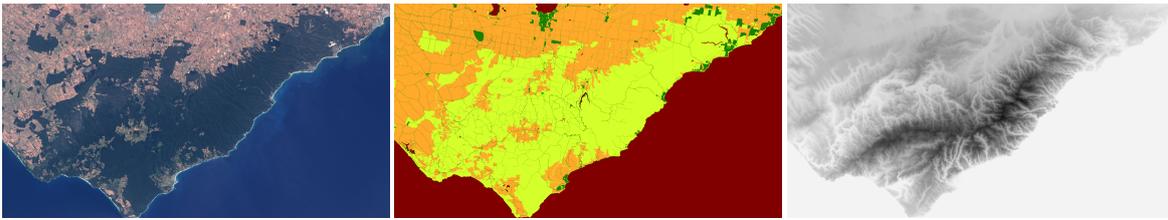
Figure 1.   RGB satellite image of Otways (left), the recorded land use (centre), and topography (right). The land use categories are: water (red), forested (yellow-green), urban (dark green) and grassland (orange). Topography ranges from 0 m (white) to 663 m (black) above sea level.

The last categories are sub-pixel and per-field. Sub-pixel algorithms allow pixels to be split into multiple classifications. This is useful in situations such as when the data resolution is too large to capture small land features. Per-field approaches consider the image as a series of objects or fields. The image is divided into objects, sets of spatially-connected pixels, and classification is done on the objects rather than individual pixels.

There are currently several limitations in the application of remote sensing data for fuel state classification. The satellite cannot see through forest canopy, and hence cannot directly provide information on forest surface fuels. Additionally, current resolutions are still too low to provide information on small scale features, such as roads. However, satellite resolution is constantly increasing and in the near future it is possible this latter issue will become obsolete. Regardless of these limitations, the benefits of automation and data recency mean that there is enormous potential for remote sensing data for calculation of fuel and coverage data in open terrain.

This paper provides an outline of methods for determining the fuel information required for Australian fire behaviour models to integrate with the fire spread prediction model, Spark (Hilton et al., 2015). It details how we used remote sensing methods to provide high resolution, up-to-date vegetation classifications from Landsat 8 data to Spark. This is the first step towards providing fuel information using remote sensing techniques such that fire simulations can be run with the most up-to-date fuel information and with minimum time.

2   CLASSIFICATION TECHNIQUE

2.1   Data and study area

Spectral data was taken from the Landsat 8 satellite, courtesy of the U.S. Geological Survey (USGS). Landsat satellites have been in use since 1972 with the purpose of monitoring land use. Landsat 8 is the most recent satellite, launched in 2013. There were 7 bands of interest: coastal aerosol (1), blue (2), green (3), red (4), near infrared (5), and short-wave infrared, SWIR, 1 (6) and 2 (7). The satellite provides coverage at a frequency of 16 days at a resolution of 30 metres for these bands. USGS also provides a 'quality of data' band which details the quality of the data with respect to possible defects such as cloud, snow or terrain occlusion. For further information, see U.S. Geological Survey (2015).

The area of interest for this paper was around the Great Otway National Park in Victoria, Australia. This area consists of a variety of land types and uses; coastal, forested national park, grassland national park, rural, grassland, and farming areas. A check of the quality band indicated no defects were occuring within the data. Auxiliary datasets available for the area were a topography map from Geoscience Australia (2013) and a land usage map from The State of Victoria, Department of Economic Development, Jobs, Transport and Resources (2015). Images of the Otways region and its recorded land use and topography are shown in Figure 1.

2.2   Classification algorithms

For this paper we used the unsupervised K-Means (KM) algorithm due to the lack of training data available to run a supervised method. Potential parameters were tested on single band classifications to determine the most appropriate. Once these were chosen, additional levels of information were added to test for improvements to output quality and computational time. The different combinations of data used as input to the KM were:

1. Multiple spectral bands; up to all seven available.
2. The first component of a principal component analysis (PCA) for each combination of the bands.
3. The first three principal components from a PCA performed on all possible combinations of the bands.
4. The ground height, slope, or aspect with the first principal component from the PCA.

Required Parameter	Tested Values	Selected Value
Number of clusters	4, 10, 50, 100	50
Termination criteria	Max. iterations=100; Final $\delta < 10$	-
Execution attempts	1, 2, 5	1
Centre initialisation	Arthur and Vassilvitskii (2007) (AV), Random	AV

Table 1. Function parameters tested and selected for the OpenCV implementation of the K-Means algorithms. Termination criteria had only a single set of values and hence the selected set was the tested set.

5. Local variance, using a 5x5 moving window, of the first component from the PCA along with the component itself.

An explanation of the implementation of both the KM and PCA algorithms is given in Section 2.3. The terrain related data sets were used to see if the addition of vegetation-influencing information could improve the classifications determined. These input data sets were set to the same resolution as the spectral bands and treated as an additional input band to the KM algorithm.

It was necessary to automatically assign a category to the computed cluster output in order to effectively analyse the classification. This was done using the land use data set by finding the median value of the land use occurring within each of the computed clusters.

2.3 OpenCV implementation of fuel classification

In order to provide fast implementation of the fuel classification processes, the OpenCV libraries of Bradski (2000) were used. OpenCV was already a key component of the Spark framework and hence the data can be directly processed into the simulation framework. These libraries provided the functions for both the KM and PCA algorithms. The KM function is moderately customisable through a number of parameters. The parameters required, tried, and selected for the OpenCV implementation of this algorithm are given in Table 1. The PCA algorithm in OpenCV calculates the covariance matrix with the input data pixels as the sample units, and the spectral readings as the variables measured on these sample units.

2.4 Accuracy evaluation of each classification method

The comparison data available for this classification was the land use data set. In order to determine the success of the classification the pixel amount by which the final classification deviated from this data set was analysed. This was done in the following stages:

1. When determining the appropriate values for the parameters in the single band analysis, the accuracy of water classification was determined. As water bodies are a generally constant landmark, except in periods of severe and extended drought, their locations are known to a high degree of confidence even if data was collected several years in the past. This was considered a good measure of accuracy for the purposes of choosing the best values. Two values were calculated; the number of pixels correctly classified as water, and the number of pixels which were classified as water but should not be water.

2. In the determination of the appropriate parameters the compactness of the classifications was also compared. The compactness is the sum of the squares of the difference between the pixel values and the determined cluster (classification group) centres. This is calculated within the OpenCV function.

3. For the consequent classifications the resulting images were firstly judged by eye. This involved checking that all the water bodies were in the appropriate places, and that the location of the forested and grasslands was intuitively correct and producing smooth rather than highly alternating regions.

4. Once the classifications of interest were selected, these were quantitatively compared to the land use image by determining the percentage of pixels in the image that were classified differently in the computed classification to the land use data.

3 INTEGRATION WITH A BUSHFIRE SPREAD SIMULATION MODEL

3.1 Fuel attributes for fire spread prediction

Empirical rate of spread models used in Australia are developed for a specific type of vegetation. As a result the primary fuel attribute required to determine the most appropriate rate of spread model for simulation is

vegetation type. This then dictates subsequent fuel attributes necessary to run the model. Subsequent fuel attributes may include: age since last fire, horizontal and vertical continuity, height, amount, moisture content, ratio of live to dead material, fuel particle size distribution, and ratio of fuel cover to no fuel.

For many years formally developed models existed for only two fuel types: dry eucalypt forests and open grasslands (Luke and McArthur, 1978). While this number has increased over the years (see Cruz et al. (2015)), for the purposes of this paper we restrict our analysis to identification of these two fuel types.

3.2 Integration with the Spark bushfire model

Spark (Hilton et al., 2015; Miller et al., 2015) is a simulation framework that models the spread of a bushfire perimeter across a landscape. A bushfire spread simulator traditionally consists of two different sub-models: a propagation model that expands the two dimensional fire perimeter in time according to some speed, and a rate of fire spread model that determines the one dimensional speed in the direction of the prevailing wind given a set of fuel and weather conditions (Sullivan, 2009b). Though physical models for fire spread do exist, in order to allow swift computation, empirical models are used for the rate of spread (Sullivan, 2009a).

Spark uses the level set method for the propagation model and enables the use of any rate of spread model. A key advantage of using a level set method is its ability to inherently deal with discontinuities in the spatial data. Therefore, there is no need to consider this potential complication when processing data for classification. As the propagation method is independent of the selected rate of spread model, rate of spread can be user defined and therefore potentially formally published or conceptual in nature. Spark allows definition of multiple spread models and the locations for which each model is applicable are dictated by a 'classification grid' input which references the appropriate model by numerical index. As vegetation dictates the particular rate of spread model used, the methodology we demonstrate is used to produce this classification grid input to Spark.

There are also a number of alternate bushfire spread simulation models that exist both in Australia and internationally. As OpenCV is an existing component of Spark, it is simple to directly integrate the classification process into the framework. The methodology, however, is also applicable to other fire simulation tools.

4 CLASSIFICATION RESULTS

The KM algorithm was run with each Landsat band individually for all test parameter values given in Table 1 in order to determine useful values for consequent classifications. All classifications were run on a standard desktop machine with Intel Core i5 processors. On average, computational time was linearly proportional to the number of clusters; taking 9s for 4 clusters and 165s for 100 clusters. Times were similar across the different execution attempt counts and centre initialisation techniques. The accuracy of water classification was also similar over the execution attempt counts. Using 4 clusters gave higher water classification accuracies, 97% compared to 92% for 10 or more, while the different centre initialisations varied by just 1%; with 93% for AV and 94% for random. However, using AV centres and using 10 or more clusters produced lower amounts of misclassification of water. In the case of the higher cluster counts this was quite a significant difference from 38% misclassification to 16%, and using AV centres decreased misclassification to 18% from 22% for random. Compactness had little dependence on execution attempts or centre initialisation, however it was significantly better as cluster count increased from 4 to 10, and then again from 10 to 50. The selected values, shown in the right column of Table 1, were chosen as a compromise between computational time and accuracy.

Five different workflows, as specified in Section 2.2, were tested. All possible combinations of bands were run as both direct input into the KM algorithm, and also through a PCA with both the first and top three principal components input into the KM. Directly into the KM the computational times ranged from 92s for two bands up to 129s for all seven. With the PCA times were similar over the different bands, around 63s for one principal component and 92s for three. Ground height, slope, and aspect were each used individually alongside the first principal component, as was the local variation of the principal component. Computational times increased to 89s with the ground height, aspect, or slope, and 111s with local variation, which included computing the variation for each pixel.

By qualitative assessment no band combination was considered sufficiently accurate in all workflows, but the six combinations that were sufficient for all but one were: bands 3/6, bands 1/3/6, bands 2/3/6, bands 2/4/6, bands 2/3/6/7, and bands 3/4/6/7. In terms of the topography the slope and aspect data sets produced the most useful classifications. Although urban areas were defined within the land use data set, there was not enough urban pixels within the image for any of the classifications to find as a majority. Hence only the three classes of water, forest, and grass were output for these classifications.
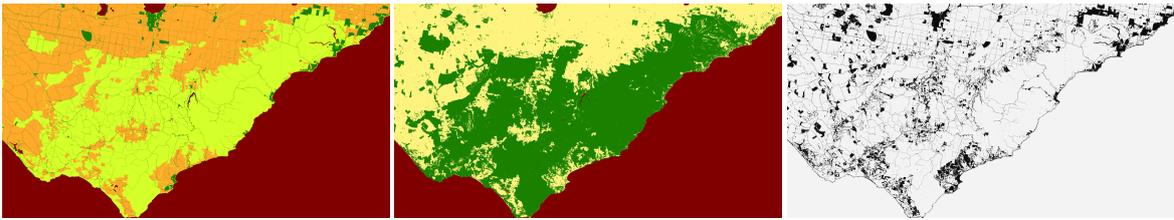
Figure 2. Results from the classification using the first principal component of a PCA on bands 1, 3, and 6 combined with the variation of the component. Shown here are: the original land use image (left), the categorised output with red as water, yellow as grass and green as forest (centre), and the difference between the original land use and the categorised output (right) where white indicates the classes were the same and black that they were different.

Once the consistently effective bands were detected, the difference between the classified image and the land use image was determined. This was calculated as the percentage of pixels that were classified differently to the land use image. The best combination was found to be bands 1, 3, and 6; the coastal aerosol, green, and SWIR 1 bands. This produced the overall lowest difference of 8.6%, with variance as an input, as well as the lowest on average using the different workflows at 9.0%. We note that 3.1% of the image was urban and hence accounts for around a third of the difference. The results of this classification are shown in Figure 2 along with an image showing the regions for which the classes were different. The largest difference value found from the selected classifications was 9.2%. The best workflow was found to be the variation in combination with the PCA, producing the lowest differences of all the different combinations and workflows, with 8.7% on average. The standard KM was next at 8.9%, followed by aspect at 9.1%, and then slope and the two PCA workflows all sitting between 9.1% and 9.2%.

5   CONCLUSION

The classification methods demonstrated in this paper are the first step towards the goal of using satellite imagery data as input to fire simulations. We have validated that it has the potential to provide a model with information on the vegetation type occurring within a region. The methods were developed within the same framework as Spark such that they can be easily incorporated into the modelling framework. The optimal inputs to the K-Means algorithm were; the first component of a PCA analysis on the coastal aerosol, green, and SWIR 1 bands and the local variation of this component. The calculated accuracy of this classification was 91% for this region. A third of the misclassification was due to the algorithm not having enough urban data available to define an urban class.

The next steps to providing the most useful simulation inputs are to produce appropriate vegetation information to provide a training set for a supervised algorithm and compare this approach to the unsupervised method. Once the optimal classification method has been determined, it will be possible to produce more detailed fuel properties appropriate to the relative vegetation classes. With more work, this methodology has potential to provide immediate, up to date fuel information on any region in Australia to bushfire models. This will result in a capability to run simulations immediately from the satellite data, and hence with almost no manual effort.

REFERENCES

Alonso-Benito, A., L. A. Arroyo, M. Arbelo, P. Hernndez-Leal, and A. Gonzlez-Calvo (2013). Pixel and object-based classification approaches for mapping forest fuel types in Tenerife Island from ASTER data. *International Journal of Wildland Fire 22*(3), 306–317.

Arthur, D. and S. Vassilvitskii (2007). k-means++: the advantages of careful seeding. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*.

Bradski, G. (2000). The OpenCV library. *Dr. Dobb's Journal of Software Tools*.

Cruz, M. G., J. S. Gould, M. E. Alexander, A. L. Sullivan, W. L. McCaw, and S. Matthews (2015). Empirical-based models for predicting head-fire rate of spread in Australian fuel types. *Australian Forestry 78*(3), 118–158.

Geoscience Australia (2013). 1 second strm derived digital elevation models version 1.0.

Gould, J. S., W. L. McCaw, and N. P. Cheney (2011). Quantifying fine fuel dynamics and structure in dry eucalypt forest (Eucalyptus marginata) in Western Australia for fire management. *Forest Ecology and Management 262*(3), 531–546.

Hilton, J. E., C. Miller, A. L. Sullivan, and C. Rucinski (2015). Effects of spatial and temporal variation in environmental conditions on simulation of wildfire spread. *Environmental Modelling & Software 67*(0), 118 – 127.

Keane, R. E., R. Burgan, and J. van Wagtendonk (2001). Mapping wildland fuels for fire management across multiple scales: Integrating remote sensing, GIS, and biophysical modeling. *International Journal of Wildland Fire 10*(4), 301–319.

Lu, D. and Q. Weng (2007). A survey of image classification methods and techniques for improving classification performance. *International Journal of Remote Sensing 28*, 823–870.

Luke, R. H. and A. G. McArthur (1978). *Bushfires in Australia*. Canberra: Australian Government Publishing Service.

Miller, C., J. Hilton, A. Sullivan, and M. Prakash (2015). Spark – a bushfire spread prediction tool. In R. Denzer, R. Argent, G. Schimak, and J. Hřebíček (Eds.), *Environmental Software Systems. Infrastructures, Services and Applications*, Volume 448 of *IFIP Advances in Information and Communication Technology*, pp. 262–271. Springer International Publishing.

Newnham, G. J., J. Verbesselt, I. F. Grant, and S. A. Anderson (2011). Relative greenness index for assessing curing of grassland fuel. *Remote Sensing of Environment 115*(6), 1456 – 1463.

Pal, M. and P. M. Mather (2003). An assessment of the effectiveness of decision tree methods for land cover classification. *Remote Sensing of Environment 86*, 554–565.

Riaño, D., E. Chuvieco, J. Salas, A. Palacios-Orueta, and A. Bastarrika (2002). Generation of fuel type maps from Landsat TM images and ancillary data in mediterranean ecosystems. *Canadian Journal of Forest Research 32*(8), 1301–1315.

Sullivan, A. L. (2009a). Wildland surface fire spread modelling, 1990-2007. 2: Empirical and quasi-empirical models. *International Journal of Wildland Fire 18*(4), 369–386.

Sullivan, A. L. (2009b). Wildland surface fire spread modelling, 1990-2007. 3: Simulation and mathematical analogue models. *International Journal of Wildland Fire 18*(4), 387–403.

Sullivan, A. L., W. L. McCaw, M. G. Cruz, S. Matthews, and P. F. Ellis (2012). Fuel, fire weather and fire behaviour in Australian ecosystems. In R. A. Bradstock, A. M. Gill, and R. D. Williams (Eds.), *Flammable Australia: Fire Regimes, Biodiversity and Ecosystems in a Changing World* (2nd ed.)., Chapter 3, pp. 51–77. Collingwood, Vic: CSIRO Publishing.

The State of Victoria, Department of Economic Development, Jobs, Transport and Resources (2015). Victorian land use information system 2010.

Thouret, J.-C., S. Ettinger, M. Cuitton, O. Santoni, C. Magill, K. Martelli, B. Zuccaro, V. Revilla, J. A. Charca, and A. Arguedas (2014). Assessing physical vulnerability in large cities exposed to flash floods and debris flows: the case of Arequipa (Peru). *Natural Hazards 73*(3), 1771–1815.

U.S. Geological Survey (2015). Landsat missions.

Varga, T. A. and G. P. Asner (2008). Hyperspectral and lidar remote sensing of fire fuels in Hawaii Volcanoes National Park. *Ecological Applications 18*(3), 613–623.

Wu, X., V. Kumar, J. R. Quinlan, J. Ghosh, Q. Yang, H. Motoda, G. J. McLachlan, A. Ng, B. Liu, P. S. Yu, Z.-H. Zhou, M. Steinbach, D. J. Hand, and D. Steinberg (2007). Top 10 algorithms in data mining. *Knowledge and Information Systems 14*(1), 1–37.