

EDGE-PRESERVING DATA ASSIMILATION FOR FIRE MONITORING USING OPTICAL DATA

J Gómez-Dans¹, P Lewis¹, M Disney¹, D Roy², T Quaife³, and M Wooster⁴

¹National Centre for Earth Observation & University College London, United Kingdom

²Geographic Information Science Center of Excellence South Dakota State University, United States

³National Centre for Earth Observation & University of Reading

⁴National Centre for Earth Observation & King's College London

ABSTRACT

Detecting burned area using data from moderate resolution sensors is fundamentally a change detection problem: as the fire affects the vegetation, it changes its optical properties. This change is detected and assessed to distinguish fire from other changes in the land surface (e.g. flooding, crop harvesting...). Two important challenges need to be overcome: the variation in reflectance ought to be related just the fire, and thus acquisition geometry-derived effects ought to be minimised, and a robust way of deciding whether a fire took place by considering the spectral pre- and post-fire reflectance needs to be in place. In this contribution, we propose a fire detection approach based on (i) a more advanced signal tracking method using edge-preserving data assimilation techniques to fit linear kernel models to the observations and (ii) an interpretation of the change signal using a spectral linear model.

For signal tracking, we use linear kernel models to interpret the surface reflectance observations. The kernel weights are inferred using a 4DVAR DA system with a weak constraint. We implement regularisation in a way that is "edge-preserving", i.e., abrupt changes in the signal are not smoothed over, resulting in adequate regularisation in the pre- and post-fire periods, but no regularisation over the fire itself. The proposed method is efficiently implemented as an iterative linear problem, and ultimately produces estimates of pre- and post-fire reflectance for a given geometry.

To interpret the change in reflectance caused by the fire, we assume that fire causes defoliation (hence exposed soils) and a combination of char and ash. We use a simple linear mixture model where the post-fire signal is assumed to be a mixture of the pre-fire reflectance plus a typical burn signal, typical a mixture of dry soils, char and ash. This burned signal can be spectrally modelled by a constrained quadratic function. The proposed model is also linear.

Using MODIS data, we illustrate the proposed method. The method is not dependent on sensor type, and we

review its usefulness, opportunities for inter-sensor data blending, and practical complications and limitations.

Key words: Fire; Disturbance; Remote sensing; Data assimilation.

1. INTRODUCTION

Fire is an important agent of change in the land surface. It has an important effect on vegetation and soils, injects greenhouse gases into the atmosphere, affects human activities and health, and has an impact on climate. The complex feedbacks between fire, vegetation, climate and human activities are only little understood, and therefore hard to model and to include in prognostic models. The erratic nature of fire, as well as its broad geographic distribution, coupled with the points stated earlier, call for effective monitoring efforts. Earth Observation is the only practical way to monitor fires on the required spatial and temporal scales, and EO has been widely used for this. From thermal anomalies, that signal fires when the sensor is overhead and no cloud or vegetation masks the signal, to fire radiative power (FRP), allowing an estimation of the rates of biomass combustion, thermal data has shown its utility. In many areas, due to both sensor and cloud persistence, thermal data is of limited use. Optical data, which works on interpreting the changes caused by the fire, gets around this limitation because burn scars are persistent in time. However, other dynamics on the land surface (vegetation recovery, snow, etc.) complicate the detection of fire affected pixels. Moreover, burned area products typically only provide a binary mask of where and when fires occur: they do not tell us anything about the impact of the fire on the vegetation, or tell us about the nature of the fire. Timeliness as well as geographic coverage suggest that the use of moderate resolution data, such as MODIS, MERIS or SPOT VGT, is the only practical way of global monitoring of fire.

A large number of burned area detection algorithms have been presented in the literature. A large family of approaches rely on change detection techniques applied to

vegetation indices, a way of reducing the spectral information to only fire-sensitive bands. Other methods combine the use of thermal anomalies with the optical data: the thermal anomalies can be used as a sort of local calibration for the change detection. Finally, other methods try to model the evolution of the reflectance signal, and predict a new observation. When the difference between prediction (and its associated uncertainty) and observation (plus its uncertainty) is deemed too large, a number of spectral tests are carried out to nominate the pixel as burned. This approach, most notably used in the MODIS Collection 5 burned area product, has many desirable properties: it uses semi-empirical kernel models to monitor BRDF effects, is able to propagate uncertainty throughout, and is fairly simple to implement. The drawbacks are that fitting the kernel model usually requires many observations, and hence the method breaks down in areas of high cloud cover. Additionally, the spectral tests that are carried out are fairly simplistic, and do not explore all the information available, both on the observations and on our knowledge of fire behaviour.

This contribution aims to present a new burned area detection algorithm that alleviates the problems outlined above. Additionally, in an era when many sensors are available, and the concept of constellations is finally taking hold, we want to explore ways to easily adapt this algorithm to cope with different data streams. To this end, we propose an algorithm that tries to (i) produce a complete time series of normalised surface reflectance, reducing acquisition geometry induced variability, but still capturing the abrupt brought about by fires and (ii) a spectral model that allows us to attribute a change to a fire, rather than to other processes. The first aim is met by extending ideas from variational data assimilation (DA) to the problem of tracking the surface reflectance signal, and the second aim is dealt with with a simple physical model that describe reflectance change from pre- to post-fire.

2. SIGNAL TRACKING

Linear kernel models have been used successfully to normalise changes in reflectance caused by changes in view/illumination geometry [1, 2, 3]. Provided that reasonable angular sampling is available, these semi-empirical models can be fitted to the data. The linear nature of the problem results in an efficient process, and an easy treatment of uncertainty. Using three kernels (isotropic, volumetric as in Ross and geometric as Li-Strahler) results in three parameters (kernel weights) required to model each observation, clearly an under-determinate problem. Traditional approaches to solving this have been the temporal compositing: all the observations available within a given temporal window are used to invert the kernels, using a least squares criterion. The main assumption is that within this period, the kernel weights are constant. The choice of compositing window is critical: too short, and not enough observations will be available, too long and changes in the land surface such as vegetation development etc., start making the assumption

of unchanging weights untenable. Further, robust methods should be deployed, as e.g. stray cloud could result in outliers that could bias the solution. Additionally, uncertainty estimates of the surface reflectance data ought to be available. Finally, for fire detection, extreme care must be taken so that pre- and post-fire observations are not mixed together.

The above comments suggest that this will typically be an ill-posed problem. Recently, [4] proposed Tikhonov regularisation for this problem: the inverse problem is augmented by an extra prior constraint of smoothness in the temporal trajectory of the parameters, effectively imposing temporal correlation in the problem. This acts as a way of channelling information from evidence rich regions, to other areas where the evidence might be weaker. The technique can be implemented fairly simply as an extension to the linear problem, and has been shown to work well with MODIS data. For the task at hand, the smoothness constraint will result in smoothing over the abrupt changes, hence blurring the definition of the fire. Mathematically, Tikhonov regularisation can be written as a minimisation problem. We will use it here to introduce the notation:

$$J(\mathbf{x}) = \|\mathbf{K}\mathbf{x} - \rho(\Omega, \Omega')\|_{\sigma_{obs}^2}^2 + \lambda^2 \sum_k \varphi(x_{k-1} - x_k) \quad (1)$$

In Eq. 1, we minimise $J(\mathbf{x})$. This functional is made up to two terms: a ‘‘fit to the observations’’ term (or likelihood), and a prior term. The first term represent the mismatch of the state \mathbf{x} (the kernel weights) and the observations $\rho(\Omega, \Omega')$, modulated by the observational uncertainty σ_{obs} . The prior term is made up of a combination of functionals that operate on first order differences in parameters. For Tikhonov regularisation as shown in [4], $\varphi(s) = s^2$. This term is large if there is a large difference between consecutive estimates of the state. This penalises sharp changes, and encourages smooth solutions. The amount of smoothness is controlled by the regularisation constant λ . In Bayesian terms, we extend the likelihood with a prior that states that consecutive states are distributed according to a zero-mean Gaussian distribution with a variance given by $1/(\lambda^2/2)$. The only parameter that needs to be estimated is λ , using a variety of methods, such as Generalised cross-validation (see [5], for an example with non-linear radiative transfer models). Note that, as written, this formulation is identical to a Kalman smoother.

Eq. 1 works well for tracking signals that evolve slowly. The presence of an abrupt change results in a large change between x_k and x_{k+1} , which penalises a solution that switches from one state to the next. This system memory complicates fire detection. Other assumptions on the prior term can be made that will smooth small variations, but let large and fairly rare changes through. Instead of using a Gaussian assumption, a t distribution, or a Laplacian (simply let $\varphi(s) = \|s\|$) can accommodate this. These expressions however result in numerical prob-

lems, and lead to fairly complex solutions. A more pragmatic approach is to relax the smoothness constraint in areas where large changes exist, effectively using a vector that modulates λ_2 as a function of the derivative of \mathbf{x} .

The ARTUR algorithm [6] is a simple and very efficient way of solving the problem. In this case, we choose $\varphi(s) = 2\sqrt{1+s^2} - 2$ (other functions are possible, but this one works well). The algorithm starts from the Tikhonov solution.

1. Solve the smooth problem:

$$\mathbf{x}^0 \leftarrow \min_{\mathbf{x}^{(0)}} \|\mathbf{K}\mathbf{x}^{(0)} - \rho(\Omega, \Omega')\|^2 + \lambda^2 \sum_k (x_{i-1}^{(0)} - x_i^{(0)})^2 \quad (2)$$

2. Find the “edges“:

$$\mathbf{w}^{(n+1)} = \frac{1}{\sqrt{1 + [x_{k-1}^{(n)} - x_k^{(n)}]^2}} \quad (3)$$

3. Solve relaxed problem:

$$\begin{aligned} \mathbf{x}^{(n+1)} \leftarrow \min_{\mathbf{x}} \|\mathbf{K}\mathbf{x}^{(n)} - \rho(\Omega, \Omega')\|^2 \\ + \sum_k \left[\frac{\lambda^2}{w_k^{(n+1)}} \right] \cdot (x_{k-1}^{(n)} - x_k^{(n)})^2 \end{aligned} \quad (4)$$

4. Go back to Eq. 3 and solve for new estimate of $\mathbf{w}^{(n+2)}$, $\mathbf{x}^{(n+2)}$, \dots , until convergence.

Convergence of the algorithm is guaranteed, and in our experiments, we found that it occurs in around 6-7 iterations typically, with areas that have a low number of observations taking up to 20. Use of sparse matrices results in an efficient and fast implementation. As an example, take a time series of MODIS data, containing both TERRA and AQUA observations over an area of Angola. In order to avoid edge effects at the beginning and end of the time series, we use three months of observations, but only look for fires in the central month (DoY between 152 and 183). We use the proposed scheme on MODIS band 2 (865nm), which is sensitive to fire. The results are shown in Fig. 1(a). We note how the method produces a good fit to observations immediately before and after the fire, which is indicated by the small magnitude of the “edge process“ \mathbf{w} . We use the estimate of \mathbf{w} derived from band 2 to solve the other bands (i.e., we plug \mathbf{w} into Eq. 5). The estimates of \mathbf{x} are then used to calculate white sky albedo (α_{WS} for all MODIS bands (as well as the associated uncertainties), and are shown in Fig. 1(b). This plot shows the clear step associated with the fire, as well as the significant reduction in “noise“ in comparison to the reflectance data shown in Fig. 1(a).

3. A SPECTRAL MODEL FOR BURNED MATERIAL

In the previous Section, a method to infer the pre- and post-fire reflectances for a common geometry has been presented. In order to determine whether a change can be classed as a fire, the spectral properties of the change can be used. Typically, fire results in drop of reflectance in the NIR and SWIR bands (MODIS bands 2, 4 and 5). Different algorithms have exploited this behaviour by proposing a number of vegetation indices that capitalise on this fact. However, a more physically meaningful avenue develops from the description of post-fire reflectance given in in [7]:

$$\rho_+^{\lambda_i} = fcc \cdot \rho_{burn}^{\lambda_i} + (1 - fcc)\rho_{unburned}^{\lambda_i}, \quad (6)$$

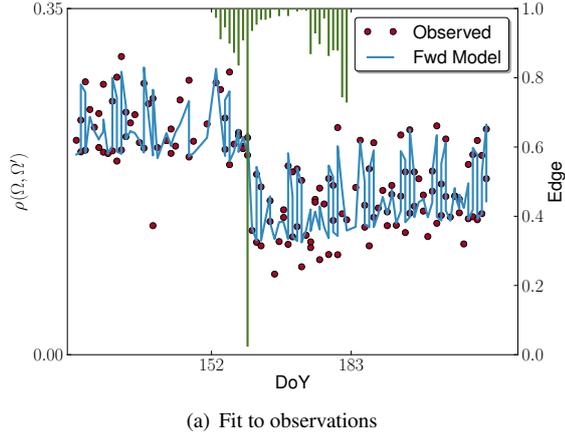
where the post-fire reflectance $\rho_+^{\lambda_i}$ is modelled as a linear combination of a typical burn signal, ρ_{burn} and an unburned material spectrum, $\rho_{unburned}$. These two components are modulated with fcc , the combined fraction of the pixel affected by the fire (f) and a radiometric combustion completeness term, cc . Assuming that $\rho_{unburned} \sim \rho_-$, the pre-fire reflectance (this assumption is useful, but it will result in complications in some scenarios), we can re-write Eq. 6 as

$$\rho_+^{\lambda_i} - \rho_-^{\lambda_i} = fcc \left(\rho_{burn}^{\lambda_i} - \rho_-^{\lambda_i} \right). \quad (7)$$

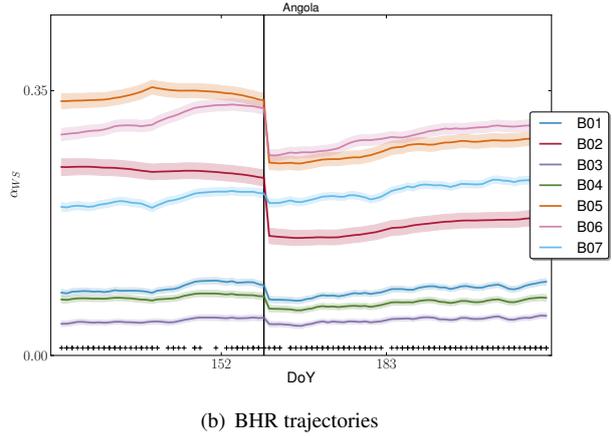
If ρ_{burn} is known, one can then calculate the value of fcc , which under this model will be related to the impact of fire in the vegetation. Note that this approach is only meaningful for understanding the fire around the time it occurred, and not on longer scales, where other changes in the land surface might complicate the model in Eq. 6. A first approximation could be that $\rho_{burn} = 0$, i.e, black char. Re-arranging terms results in fcc being equal to the relative change in reflectance, a term that has been already linked to the impact of fire. A more subtle model for ρ_{burn} can be constructed by compiling a database of materials that are likely to feature in the burned scene. We assume that the fire will result in a mixture of ash, char and exposed soil. After compiling one such database (see Grévez-Dans *et al*, in prep.), a principal component analysis shows that the first two principal components explain most of the variance (in excess of 80%) for a number of spectral configurations, as can be seen for the case of MODIS in Fig. 2. The first component is fairly similar to a char or ash spectrum (flat spectral response), whereas the second is similar to a dry soil spectrum. These two components have similar shapes for other spectral configurations, and can be modelled pragmatically with a constrained quadratic function as

$$\rho_{burn}(\lambda) = a_0 + a_1 \cdot \left[2(\lambda - \lambda_0) - \frac{(\lambda - \lambda_0)^2}{\lambda_{max}} \right], \quad (8)$$

where a_0 and a_1 control the magnitude of the projection in the char/ash and exposed dry soil axes, respectively. λ_0 is fixed to the lowest wavelength of interest



(a) Fit to observations



(b) BHR trajectories

Figure 1. (a): Using the edge-preserving DA method to fit MODIS observations (red dots). The forward modelled solution is shown as the blue line. We also depict the magnitude of \mathbf{w} as the green lines. (b): Deriving white sky albedo for all MODIS bands from the kernel weights estimated for the time series shown Fig. 1(a). Small +-signs represent MODIS observations.

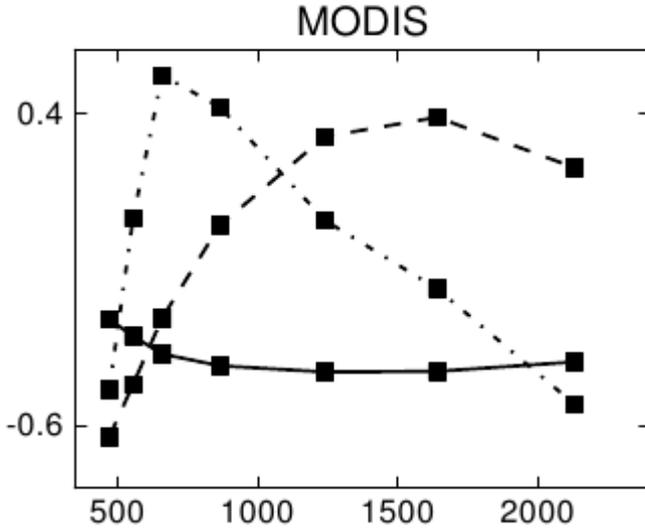


Figure 2. Three principal components resulting from the PCA analysis (x-axis: Wavelength in nm, y-axis: principal component value). Together, they represent in excess of 85% of the variance in the dataset. First component is full line, second is dashed line and third is dotted line.

(400 nm) and λ_{max} was estimated from fitting the spectra in the database, and found to be $\lambda_{max} = 2000nm$. In Fig. 3 the RMSE resulting from fitting Eq. 8 to spectra (using the MODIS bands) is shown. Note that the model provides a good fit (low RMSE) for char and soils, an intermediate fit for "brown" material (e.g. senescent leaves, woody debris, etc.), and a poor fit to green vegetation.

With the model of ρ_{burn} given by Eq. 8, we can re-write Eq. 7 as

$$\rho_{+}^{\lambda_i} - \rho_{-}^{\lambda_i} = f_{cc} \left(a_0 + a_1 f_1(\lambda_i) - \rho_{-}^{\lambda_i} \right). \quad (9)$$

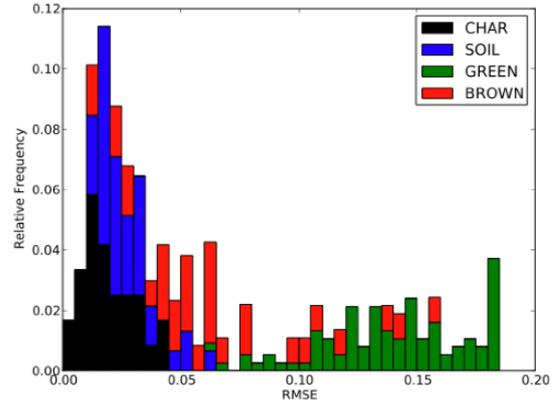


Figure 3. Fitting the model for ρ_{burn} in Eq. 8 to individual spectrum using the MODIS spectral sampling configuration.

Eq. 9 states that if pre- and post-fire reflectances are available for three or more spectral bands, one can estimate f_{cc} , a_0 and a_1 by solving the ensuing linear model. Note that in order for this equation to hold, the following assumptions have been made: the pre- and post-fire reflectances do not have biases due to BRDF effects; there are no other changes present in the scene than those attributed to fire; the burned portion of the scene can be modelled as mixture of char, ash and exposed soils; approximating $\rho_{unburned}$ by ρ_{-} is sensible. These assumptions are not controversial, provided that the pre- and post-fire data are close in time. The solution of Eq. 9 results in an estimate of the impact of the fire on the pixel from f_{cc} , as well as a description of the spectral properties of the burned material spectrum. The method uses all available evidence in an optimal way. Note that the linear nature of the problem makes propagating uncertainty

simple.

We recognise that this model resembles other linear mixture models that have been proposed in the literature to understand fire impact. However, remark that the model does not try to explain pre- or post-fire reflectance, but rather, *change in reflectance*. This simplifies the problem specification, as one does not need to deal with other "end-members", such as green vegetation: this information is assumed to be carried by $\rho_{unburned}$. Also, although the model is independent of the sensor used, different spectral configurations will result in better estimates of the parameters. In the case of sensors operating in the visible/NIR part of the spectrum (MERIS, CHRIS/PROBA), the definition of f_1 above might be changed to a simple linear term rather than a quadratic. Finally, topography, the difficulty in normalising angular acquisitions, and the potentially large period between acquisitions present a challenge to the application of these method to Landsat TM/ETM+ data.

MODIS data allows for the testing of the model. Using the official MODIS burned area algorithm, the time of burning of different pixels is used to then select all the observations in the 16 days prior to, and posterior to the fire. Linear kernel models are fitted to these observations, and estimates of nadir-looking surface reflectance (and associated uncertainties) are produced for the seven MODIS bands. These are then used to estimate the model parameters over different areas. For a single pixel, the inversion process inputs and outputs can be seen in Fig. 4. Similarly, we can look at spatial patterns of the retrieved parameters. Note the clear pattern over the large fire shown in Fig. 5, with f_{cc} close to unity in the inside (grass fires should have large combustion completeness), indicating that most of the pixel is changing to a "burn signal". On the edges of the burn scar, f_{cc} decreases to values around 0.4-0.3, suggesting that the fire has stopped and not more or less burned the entire pixel. Given that each pixel has been processed separately, this is a striking result. Finally, observe the distribution of a_0 and a_1 in Fig. 6, and note how the retrieved parameters appear bounded by the ground measurements, implying that these are physically meaningful parameters, and not just artefacts of the inversion process. Additionally, two modes can be seen in the data. The one with lower a_1 is similar to black char spectra, and can be spatially mapped to areas where trees are present, and where char deposits will be important (in this case, miombo woodlands). The other mode can be mapped to grassland fires, where very char ash and char dissipation will result in mostly exposed bare soil.

4. THE BURNED AREA ALGORITHM

In the previous sections, we have introduced a technique that inverts linear kernel models, and can thus be used to obtain estimates of e.g. bi-hemispherical reflectance. Abrupt changes are obvious in the data. The iterative algorithm introduced in Eq. 5 is used for MODIS bands 2, 4 and 5. Initially, it is used for burned area estima-

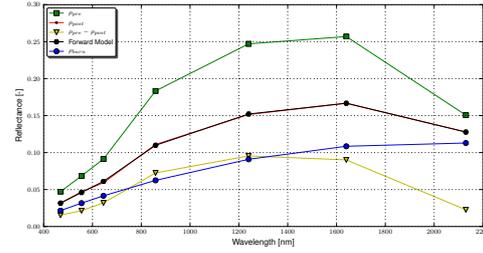


Figure 4. Spectral demonstration of the f_{cc} algorithm with MODIS data. The inversion takes as inputs ρ_- (green) and $(\rho_- - \rho_+)$, estimates the model parameters, and then ρ_{burn} (blue) and ρ_+ (red observed, black modelled) can be calculated.

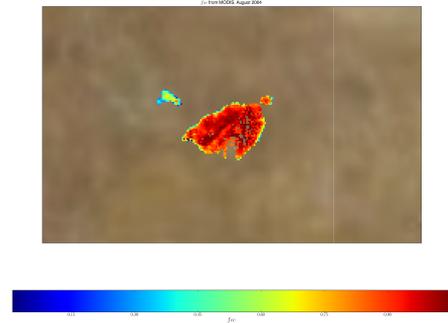


Figure 5. Spatial distribution of f_{cc} over a large grass fire. Note the high values inside the burn scar, and the tapering towards the edge.

tion for single months, but a whole year could be used if required. We extend the time series for an extra month before and after the fire so edge effects do not affect the inference. Each of these single-band processing results in a different edge process. For each, the minimum is located, and the minimum of the three edge processes is then used for inverting all MODIS 7 bands. The reason for this is to capture changes due to fire that sometimes appear more obviously in one of the three bands indicated above. Equipped with the BHR data estimates for every day, the f_{cc} algorithm is applied to consecutive days for the period of interest. Estimates of f_{cc} , a_0 , a_1 and the RMSE fit. We then select the maximum f_{cc} value of the time series, and apply a number of tests to a_0 , a_1 and the fit are assessed. Only if all the parameters have physically-meaningful values, a fire is identified. This is depicted in Fig. 7, where a fire is clearly visible as a drop in reflectance in all bands except band 7. Also shown in Fig. 7 is the temporal evolution of f_{cc} and the model RMSE error. It is obvious that f_{cc} attains its maximum value on the day of the large drop in reflectance caused by the fire, showing a small RMSE value. Given that a_0 and a_1 do have a physical meaning, their values can also be assessed to make sure that ρ_{burn} appears to look like the ash/char/exposed soil mixture that is being modelled. These tests indicate the clear presence of a fire (as well as

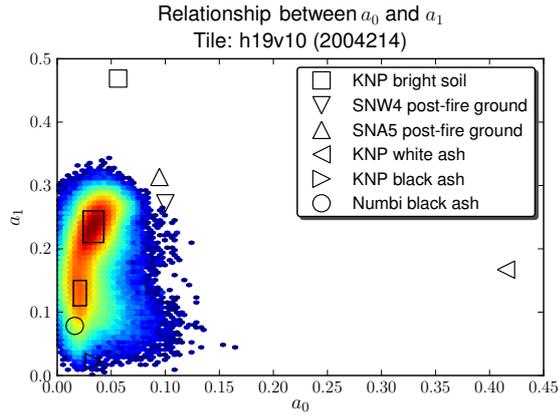


Figure 6. Scatter plot of a_0 and a_1 over Angola (colour wedge), as well as parameters retrieved from fitting the model to some ground spectra (open symbols).

the timing of the event). The actual thresholds are purely based on assuming the model is physical sound (e.g. f_{cc} should lie in $[0.1, 1.1]$, a_0 and a_1 should be consistent with spectral measurements, and RMSE fit is bounded by the noise in the original data). Note that in this particular example, one could have directly fitted the spectral model using the edge process for detection. In practice, stray clouds nearing areas with no observations can result in the edge process indicating a discontinuity, so in order to avoid just testing this effect, we run the model over the whole time series, given its speed.

In Fig. 8 we see the spatial distribution of the detected day of burn for the MCD64, MCD45 and the present algorithm. Visually, the three products show similar results for this particular area. The proposed approach appears visually similar to the MCD45 product. This is unsurprising, as the new method incorporates a way of dealing with BRDF effects that is similar to that implemented in the official MODIS product, perhaps more refined and with a more subtle set of spectral constraints. The MCD64 algorithm uses an important source of extra information to detect fires. In addition, we see in Fig. 9 that from the new method, the linear change model parameters are also estimated. This is important additional information that might give further clues to identify the effect of the fire. Note, for example, the common occurrence of f_{cc} decaying from a high value at the centre of the burn scar, to a low value towards the edge. Each pixel has been processed individually in this case, but still these remarkable spatial patterns show up (similarly for a_0 and a_1).

5. FINAL REMARKS

We propose a method for burned area detection based on tracking the surface reflectance with an edge-preserving data assimilation approach (a simple extension to the variational method with temporal regularisation introduced in [4]), plus a spectral model that explains the

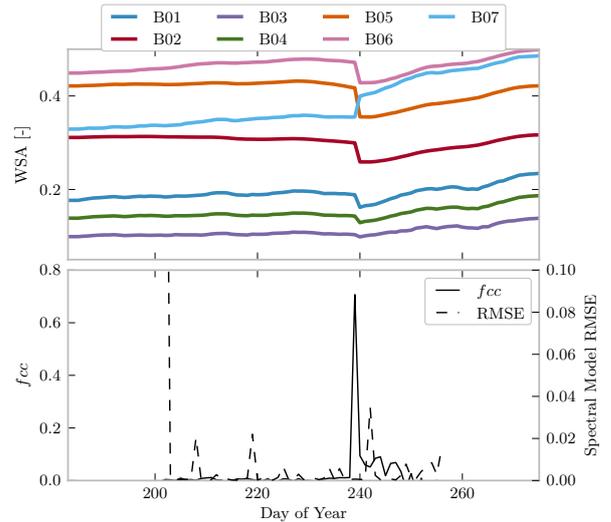


Figure 7. (Top) Result of applying the edge-preserving regularisation algorithm to a MODIS pixel in Angola. (Bottom) Time series of f_{cc} and the RMSE fit.

change in reflectance due to the change as a linear combination of the pre-fire reflectance and a typical burn signal spectral model. The use of linear operators for the signal tracking, as well as the linear nature of the spectral model results in a practical and efficient approach, with simple uncertainty propagation. We have demonstrated the behaviour of the system with MODIS data, but the algorithm can be applied to other sensors.

A number of issues can be found. The first one is the choice of regularisation parameter λ in Eq. 1 (and subsequent). In [5], a generalised cross validation (GCV) approach was taken for estimating biophysical parameters. In this case, the value of λ was estimated for the complete algorithm (*i.e.*, both regularisation, spectral model inversion and fire identification), by sweeping λ and comparing the overall accuracy of a small test dataset. It was found (not shown here) that the value of λ is very stable. In other words, fairly large changes in the estimation of λ due to the use of small training set do not result in large changes in the overall effectiveness of the proposed model. Moreover, we found that even reducing observations for the estimation of λ (under the assumption that the number of observations is an important factor in the determination of λ) still resulted in similar results. These methods are *ad hoc*, and do need to be revisited and more formally resolved. However, it does appear that stronger regularisation that would be suggested to just retrieve the surface reflectance results in a better fire detection behaviour.

Secondly, the spectral model is based on a limited set of spectra. More ground data would be useful for testing the ability of the model and maybe refine the distributions for a_0 and a_1 used to affect a candidate f_{cc} inversion as a burned pixel. In regional applications where ground data might be available, actually using the principal compo-

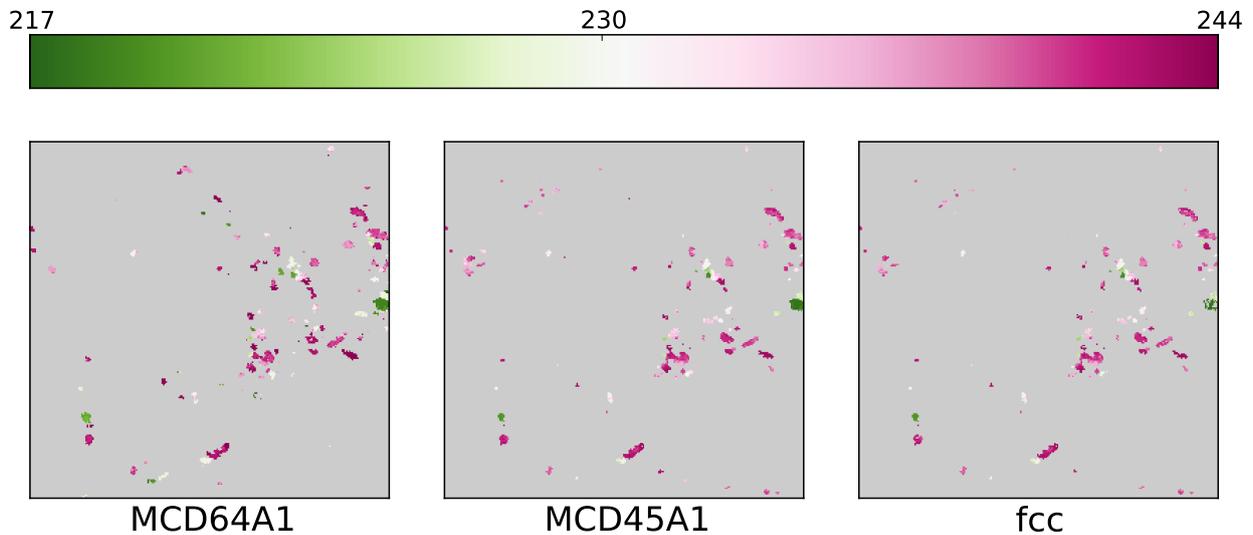


Figure 8. Day of Burn (2004) for a savannah region in Botswana, as estimated from (Left) the MCD64 product, (Center) the MCD45 product and (Right) the proposed algorithm.

nents rather than the model introduced by Eq. 8 might even be desirable. Note however, that for larger area or global application, using the model will probably result in a more general result that will adapt better to different fire types, etc. In addition to this, careful studies on some model choices based on both radiative transfer modelling [8, 9] as well as field observations is clearly required.

Thirdly, note that the model is sensor independent. Given that full uncertainty about its parameters can be derived, it opens the possibility of inter-sensor blending, where the application to one sensor can be combined with the application to a different sensor in a fairly simple Bayesian scheme. The edge process could also be used in this fashion. An important consideration will be how the spectral model deals with restricted spectral sampling of say the SWIR region. In those cases (e. g. MERIS), a simpler model (a linear reflectance change between in $[400 - 900]nm$ could be used rather than Eq. 8. Note that a method to match spatial scales would still need to be provided.

Finally, the idea of piecewise areas in time can be transposed to space, with neighbouring pixels exhibiting a strong probability of burning in the same or recent day. This suggests expanding the current Markov Random

Field (MRF) in time to one that also covers space. Again, edges would need to be dealt with in a manner equivalent to the one used for time. The spatio-temporal problem effectively provides a diffusion process for information from evidence rich areas, to areas that are more challenging (either because they have less data, or because the proposed model is less effective at describing the change brought by by the fire). This diffusion process can also be exploited for including extra fire information, such as the presence of thermal anomalies. These indicate both a heightened probability of fire at a given time and space, but also serve as a local calibration of the model. Data at different spatial resolutions can also be transparently included in this set up, which will be our next line of investigation.

REFERENCES

- [1] Wolfgang Wanner, X Li, and AH Strahler. On the derivation of kernels for kernel-driven models of bidirectional reflectance. *Journal of Geophysical Research: Atmospheres (1984–2012)*, 100(D10):21077–21089, 1995.
- [2] W Wanner, AH Strahler, B Hu, P Lewis, J-P Muller,

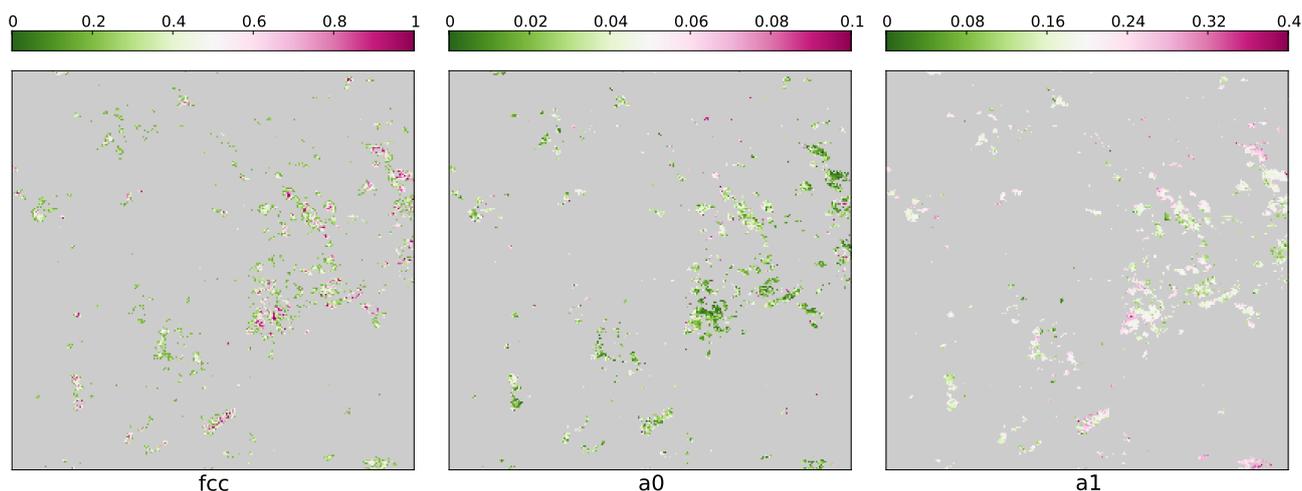


Figure 9. Same are shown in Fig. 8, but showing the additional parameters retrieved from the proposed algorithm: (Left) f_{cc} , (Center) a_0 and (Right) a_1 .

X Li, CL Schaaf, and MJ Barnsley. Global retrieval of bidirectional reflectance and albedo over land from eos modis and misr data: Theory and algorithm. *Journal of Geophysical Research: Atmospheres (1984–2012)*, 102(D14):17143–17161, 1997.

- [3] W Lucht and P Lewis. Theoretical noise sensitivity of brdf and albedo retrieval from the eos-modis and misr sensors with respect to angular sampling. *International Journal of Remote Sensing*, 21(1):81–98, 2000.
- [4] T. Quaife and P. Lewis. Temporal constraints on linear brdf model parameters. *Geoscience and Remote Sensing, IEEE Transactions on*, 48(5):2445–2450, 2010.
- [5] P Lewis, J Gómez-Dans, T Kaminski, J Settle, T Quaife, N Gobron, J Styles, and M Berger. An earth observation land data assimilation system (eoldas). *Remote Sensing of Environment*, 120:219–235, 2012.
- [6] P. Charbonnier, L. Blanc-Féraud, G. Aubert, and M. Barlaud. Deterministic edge-preserving regularization in computed imaging. *Image Processing, IEEE Transactions on*, 6(2):298–311, 1997.
- [7] DP Roy and T. Landmann. Characterizing the surface heterogeneity of fire effects using multi-temporal reflective wavelength data. *International Journal of Remote Sensing*, 26(19):4197–4218, 2005.
- [8] MI Disney, P Lewis, J Gomez-Dans, D Roy, MJ Wooster, and D Lajas. 3d radiative transfer modelling of fire impacts on a two-layer savanna system. *Remote Sensing of Environment*, 115(8):1866–1881, 2011.
- [9] P Lewis, T Quaife, JL Gomez-Dans, M Disney, M Wooster, D Roy, and B Pinty. Modelling the impact of wildfire on spectral reflectance. In *Geoscience and Remote Sensing Symposium, 2009 IEEE International, IGARSS 2009*, volume 4. IEEE, 2009.