Landscape Assessment (LA)
Sampling and Analysis Methods

Carl H. Key
Nathan C. Benson

REMOTE SENSING MEASURE OF SEVERITY: THE NORMALIZED BURN RATIO

If Unfamiliar With The Science Of Remote Sensing

There are many excellent references to general principles, methods and applications of remote sensing on the Internet. A recommended subject matter text is *Remote Sensing and Image Interpretation*, by T.M. Lillesand and R.W. Kiefer, 1994 (John Wiley & Sons, New York, 750 pp). If you do not have image processing software we suggest a user friendly program called MultiSpec for viewing and exploring raw Landsat data. MultiSpec is available at no cost from the LARS facility of Purdue University, and runs in PC or Macintosh environments. Documentation provides good discussion of basic remote sensing concepts: http://dynamo.ecn.purdue.edu/~biehl/MultiSpec/Index.html. MultiSpec is suggested as a general learning tool, and as a preliminary data exploring tool, but it is not capable of the processing required for burn severity mapping, as described later in this section.

Introduction to the Normalized Burn Ratio (NBR)

Raw Landsat multi-spectral data contains a wealth of information about earth features. Each spectral band responds in unique ways to surficial characteristics like water content, vegetation structure, productivity, and mineral composition. When brightness values of multiple bands are combined in mathematical algorithms, information about targeted features can be enhanced, isolated and analyzed. From available raw data, the challenge for burn severity is to develop a specific index providing an optimum measure and useful signal for fire-effects. The index we developed is called the Normalized Burn Ratio or NBR. It is similar in construct to another standard index, called the Normalized Difference Vegetation Index or NDVI. The primary difference is that NBR integrates the two bands that respond most, but in opposite ways to burning (figure LA-7). Those were determined to be TM/ETM+ Band 4 and Band 7. The NBR is calculated as follows:

Eq. LA-1 $$NBR = (R4-R7) / (R4+R7)$$

where R values are the calculated per-pixel "at satellite" reflectance quantities per band, which have been corrected for atmospheric transmittance.

Based on experience in generally forested ecosystems of the Western U.S., *R4* decreases while *R7* increases from pre-fire to post-fire TM/ETM+ acquisitions. The change is greatest in magnitude compared to other bands, and the variance within burns is greatest for *R7*. The combination of those traits, then, appears to provide the best distinction between burned and unburned areas. It also provides an optimum signal for information about variation of burn severity found within the burn. The **(***R4 - R7***)** difference is scaled by the sum of the two bands to normalize for overall brightness that is consistent across the bands. It helps remove within-scene topographic effects and between-scene solar illumination effects. This effectively isolates the real reflective differences between the bands, which enables spatial and multi-temporal comparison of the derived NBR values.

To isolate burned from unburned areas, and to provide a quantitative measure of change, the NBR dataset derived after burning is subtracted from the NBR dataset obtained from before burning, such that:

Eq. LA-2 $$dNBR = NBR_{prefire} - NBR_{postfire}$$

This measured change in NBR, delta NBR or **dNBR**, is hypothesized to be correlative in magnitude to the environmental change caused by fire, i.e. the burn severity as it relates to fire effects on previously existing vegetative communities. Assuming unburned terrain is relatively similar in phenology and moisture between the two sample dates, and the two datasets are adequately co-registered, background areas take on values near zero in **dNBR**. Likewise, burned areas assume strongly positive or negative values, depending on whether the fire distresses or actually enhances productivity on the site. The latter can occur in herbaceous communities, where severity is light and ephemeral, and burned vegetation responds quickly with renewed vigor from the release of nutrients or other factors after fire. Strongly positive **dNBR** is more typical, however, in forested and shrub-dominated areas, where fire generally creates longer-lasting conversions of biomass to less productive or earlier successional states.

In either case, burned areas can be suitably distinguished from unburned, and potential exists for a wide range of **dNBR** within the burn (depending on actual characteristics of the subject fire). This range appears to resolve the breadth of fire effects, revealing the complexity and spatial heterogeneity of the burn. It also appears to be a broader range than other radiometric indices tested, such as differenced **NDVI**. Results are constrained by the 30-meter Landsat resolution, however, which makes the index appropriate for landscape perspectives that yield whole-burn spatial data on severity.

Timing of Landsat Aquisitions

The dates from which pre- and post-fire data are acquired by the TM/ETM+ are extremely important. If not carefully considered, they may be the primary pitfall leading to unsatisfactory results.

One contributing reason is that our approach involves change detection specifically targeting burned areas. Ideally, results show only change caused by fire, with all other surface features remaining neutral so as to elucidate the unburned areas that did not change. Unfortunately, unburned features in the landscape do not remain static over time, being naturally altered by wetting and drying, and cycles of productivity (figure LA-8). As a result, the NBR difference, though apparently optimal for resolving burn variations, still may be affected by such factors. (Other band combinations tested were influenced by seasonal acquisition timing to a greater degree than NBR.)

Therefore, to better isolate and enhance the burn signal, the pre- and post-fire Landsat datasets should be chosen to represent moisture content and phenology as similarly as possible. This timing is relative to localized growing seasons, which may vary by date and location from year to year. Landsat scenes should be compared in false color for indications of seasonal differences (e.g. Bands 4, 5, 7 or 7, 4, 3 for red, green, blue, respectively). Two helpful characteristics to key on are 1) the productivity indicated by Band 4 in herbaceous and shrub communities, which typically shows strong seasonal patterns of regrowth, decline and curing out; and 2) the pattern of seasonal snowmelt or other regular change in surface moisture. Available Landsat scenes can be viewed on line to appraise these trends (see links for ordering Landsat data provided in **Appendix A**):

http://edcsns17.cr.usgs.gov/EarthExplorer/
http://edclxs2.cr.usgs.gov/

Of annual periods, early to middle growing season dates seem to yield best results. That is when unburned vegetation is green and lush (orange-red in figure LA-8), showing peak contrast with areas affected by fire. Results remain good through late in the growing season, but as areas dry and deciduous plants cure (blue-gray in figure LA-8), some resolution of the burn may be lost. In tests, NBR tended to minimize this issue, unlike NDVI, which degraded more strongly late in the season. Especially, distinction between unburned and low burn severity tends to diminish late in the growing season. By then, cured-out vegetation can mimic fire effects, and burn effects show less contrast against the background of unchanged but dry vegetation.

With these issues in mind, optimal timing of TM/ETM+ acquisitions may be difficult to achieve given cloudiness and the 16-day return interval of the satellite. If so, consider increasing options by reviewing even predominantly cloudy scenes. Cloud-free areas need only extend enough to encompass the burn(s). In addition for the pre-fire scene, data acquisition can safely occur within 2, perhaps 3 years before the burn so long as other landscape disturbances (including interim fires) are accounted for and do not interfere with the subject burn. See discussion below for options on the post-fire scene.

Two Strategies, Initial vs. Extended Assessment

With some exceptions, many severity indicators are apparent soon after fire is out. These have to do with scorching, charring and consumption of living vegetation and dead fuel, and with changes in the nature of exposed mineral soil and ash. The exceptions, though, are important. They concern initial recovery of vegetation and delayed mortality, which also contribute to near-term severity. Often those factors may not become evident until at least the following growing season, so some passage of time may be required to get the fullest assessment of the burn. Given these circumstances, we recommend two scenarios for processing and applying dNBR severity measures (figure LA-9).

Initial Assessment

The Initial Assessment bears upon the most-immediate fire effects to biophysical components that existed at the time of fire. It uses a post-fire TM/ETM+ scene from as soon after fire as possible. In order to match phenology (see discussion on Timing), the pre-fire scene generally comes from a similar period of the previous year, or if necessary, the year before that. The only exception is if the fire was very short-lived, and both scenes could be acquired within 8 to 24 days of one another. The latter also assumes there were no marked changes in moisture or productivity over the period. Initial Assessment can provide very good delineation of burned area, and preliminary estimates of severity. It may not be optimum, though, for the following reasons: 1) vegetative regeneration will likely be missing, which may lead to overestimating severity; 2) unburned vegetation may be naturally cured out by the end of the fire season, diminishing burn contrast (see Timing); and 3) when the fire season extends from late summer into fall, sun angles may be low and there may be limited time before bad weather and/or snow obscures the burn. Late-season initial assessment, though perhaps the only timing available for some emergency response applications, may show less definition of the perimeter, less data range and contrast within the burn, and poorer correlation with field data. When possible however, severity will certainly be suggested, since this timing integrates at least part of the suite of factors responding to fire. Importantly, we found that dNBR was less affected by late-season effects than other indices, such as dNDVI, and concluded that dNBR held up better and was still useful in late-season. As such, Initial Assessment can yield figures on burn size, composition and complexity within a month or two of the fire. Products may suffice for many needs concerning public information, planning and rehabilitation.

Extended assessment

The second scenario provides an Extended Assessment, which we believe is more representative of the actual severity. It postpones acquiring the post-fire TM/ETM+ scene until the next growing season, which may be as soon as a few weeks or as long as 11 months after fire, depending on the ecosystem and climate. The pre-fire scene is then usually taken from that same seasonal period but during the year of the fire, since that period typically falls before a given fire season. If necessary, a pre-fire scene can come from a year or two before fire, so long as conditions are comparable and no interim disturbances overlap the burn. By waiting until the following season, burned vegetation has had a short while to recover and demonstrate resprouting that is one factor for gauging severity. Delayed mortality may also be evident; revealing that plants green right after fire had died by the next growth period. Results of Extended Assessment, then, would be most useful for final portrayal and statistics of initial severity, or first-order fire effects. They would be suited for projects that depend on more accurate delineation of burn heterogeneity than Initial Assessment, like those comparing multiple burns over space and time, testing methods, or modeling. They also might better address long-term ecological consequences, such as impacts to sensitive communities or species, or risk factors like erosion and future fire potential.

Pre-Fire Considerations

General software requirements

There are several image processing systems that can fulfill the needs of these procedures, such as ERDAS Imagine or GRASS, which analyze geographically referenced image data. Whatever system is used, there are two capabilities that, while not unusual, may be unavailable in some systems. First, one will have to be able to write and execute algorithms that incorporate raster datasets as variables, and output raster datasets as results. Second, one will need to be able to do mathematical operations in floating-point math, which requires manipulating and storing rasters in signed 16-bit or 32-bit data formats. That is, two or four bytes per pixel, since pixels can take on positive or negative values with at least four significant digits. The software also should be able to generate topographic slope and aspect data from a Digital Elevation Model (DEM). If not part of the image processing system, a statistics package also is needed to perform regression analysis for normalization of atmospheric effects.

General Hardware Requirements

Anticipate working with a large number of big data files, in excess of 40 megabytes each. Steps can be taken to subset large scenes into smaller working regions, but some processes are best done on full scenes. Especially, if there is any possibility the data might be used for future fires or other applications. Minimum starting disk space should be in the 10-gigabyte range; minimum RAM should be at least 256 megabytes. You will find a need to eventually manage around these limitations if work expands much beyond a couple of fire years in any one-scene area. A 21-inch graphics monitor is also recommended for on-screen interpretations and digitizing.

Digital Geographic Data Needs

Besides two Landsat scenes (pre-fire and post-fire), other types of thematic GIS data are highly useful for checking registration, orienteering, or supplementing maps. A digital elevation model (DEM) of the burn surroundings is recommended. DEMs are available at no cost from the EROS Data Center, if "terrain" or "precision" corrected Landsat data is ordered. Those DEMs encompass the entire area of the scene, so the file is quite large. Some basic vector datasets for ownership, roads, trails, watersheds, lakes and streams are also helpful, as well as digital raster graphics (DRG) and digital orthophotoquads (DOQ). Finally, there may be digitized fire perimeters from the fire incident teams, which can help locate the general area of the burn. Keep in mind, these perimeters may not include all burned area, and may vary in the quality of reference data available, even over the time of one fire.

Ordering the Data

Once either or both scenarios have been decided upon, and the seasonal requirements for pre- and post-fire Landsat datasets are understood (see above), search the Landsat archives for available scenes that best meet requirements. Both Landsat 7 and Landsat 5 data may be available, so both archives should be searched. You will need to obtain the chosen scene identifiers prior to ordering. These relate to the satellite path/row (or the geographic area covered) and the date. For more information on how to preview and order data, see the FIREMON Landscape Assessment Appendix A. on Landsat Satellites.

When ordering the data, there are a number of format and processing options to choose from. For complete listings with definitions see:

http://edcdaac.usgs.gov/tutorial/daacdef.htm

For description of processing levels see:

http://landsat7.usgs.gov/l7_processlevels.html.

A sample specification is shown below for one Landsat 7 scene. Options in bold are recommended for all orders. The rest depend on your specific area, date, and working map base. If the datum is other than the default NAD83, it must be clearly specified. The highest standard processing level for ETM+ is L1G, which is geometrically corrected without ground control or relief models. We recommend, however, special-order level L1T, terrain corrected from ground control and relief models. It must be unambiguously specified, however, or processors will assume one of the lower levels. It costs about $200 extra, but it is well worth it, since terrain registrations are time consuming and may not be possible to do "in house". One additional option is to order the DEM for the scene area. This can be included at no cost, but will probably only be needed once per path/row, or not at all if a DEM is already available. As with Landsat data, map projection and datum should match what is in current use by the end user.

```
Item 001
Data granule: E1SC:L70RWRS.002:2001192063
Data set: LANDSAT-7 LEVEL-1 WRS-SCENE V002
Path/Row: 43/34
Acquisition Date: 02 July 2000
Ordering Option: E1SC:L70RWRS.002:2001192063:
L1T Product - TERRAIN CORRECTED UTM Projection
Cost:  US $800.00
Format/Media: FASTL7A: CD-ROM: ISO 9660
Additional Information: WITH TERRAIN CORRECTION USING NAD27 DATUM
ORDER Options:
   Product:     L1T
   Projection:  UTM, ZONE 11, NAD27 DATUM
   Radiometric Correction Method: CPF
   Band Combination: 1; 2; 3; 4; 5; 6L; 6H; 7; 8
   Image Orientation: NUP
   Resampling Method: CC
   Grid cell size for the pan band (8): 15.0
   Grid cell size for the reflective bands (1-5, 7): 30.0
   Grid cell size for the thermal bands (6L, 6H): 60.0
   Zone Number: 11
```

These are the recommended options, given the following factors. Further processing involves reflectance-based calculations, so terrain registration is important. Multi-temporal comparisons will be made, and pixel boundaries shift scene to scene, so the resampling method should provide the best estimate of reflectance for the geographically rectified space.

Steps to Process NBR and dNBR

The steps outlined below are intended to be somewhat generic, recognizing that differences exist between various processing systems available to users. In general, steps identify what is needed along the way, and not so much how to get there. One will need to find the proper procedures and syntax available within one's specific system. Usually systems provide analogous functions, so one should be able to adapt steps easily enough. We assume those undertaking these procedures are well grounded in remote sensing principles, and the functionality of the processing system in use. If not, one may want to consult a local expert.

I. Initial setup for operations, for each Landsat scene:

1. Review the known facts about the burn, location, start and end dates, approximate size and geographic distribution.
2. Plan a naming convention for the large number of files to be generated. A recommended sequence is provided in the FIREMON How To document.
3. Extract the Landsat scene header file and print it out for reference.
4. Import Raw Bands 4 and 7 into image processing system (or GIS) for analysis.
5. Recommended, but optional, import other bands to explore data in false color.
6. Review the burn area displayed on the pre- and post-fire imagery. Become familiar with its distribution within the surrounding landscape, and juxtaposition to geographic features and named places.
7. Ensure compatible map projections exist between all data layers, and check registration of Landsat scenes; one to another and in conjunction with basic reference data (e.g. lake boundaries). Mis-registration of more than one half pixel should be corrected.

Note: Some image processors save rasters only as integer data, while the following calculations are done in floating point math and generate datasets of real numbers. If that is a problem, reflectance and subsequent NBR calculations can be scaled by 1000 to retain positive or negative integer values with four significant digits.

II. Radiance and reflectance transformations:

The procedure is used to derive "at satellite" reflectance for Band 4 and Band 7 of each Landsat scene. To expedite calculations, if possible, an executable script can be written and calculations can be combined in one algorithm. The script can then be modified for subsequent analyses by simply replacing the scene-specific parameters.

Note: these transformations are not specific to remote sensing of burns. Rather, they are recommended for any analysis that involves quantitative comparison of different Landsat scenes. They standardize the bands, account for drift in the multi-spectral scanner, normalize daily variation in sunlight, and optionally, correct illumination differences caused by topography. Standard order for processing mathematical operators is used.

The radiance per pixel per band is calculated as:

Eq. LA-3 $$L_i = DN_i * G_b + B_b$$

where i is a particular pixel, DN_i is the per-pixel raw Landsat band brightness value (density number, or digital number), G_b is the gain and B_b is the bias for a particular band, b (in this case, Band 4 and Band 7). The G_b and B_b are reported per band in the scene header file.

The reflectance per pixel per band, is calculated as:

Eq. LA-4 $$R_i = (L_i * \pi * d^2) / (\mathrm{Esi}_b * \cos(z_s))$$

where d^2 is an eccentricity factor for earth-to-sun distance, Esi_b is the per-band exoatmospheric solar irradiance constant, and z_s is the per-scene sun zenith angle. Here, the underlying assumption is that the earth surface is flat, and only the sun zenith angle is used to calibrate Esi_b in the denominator. While topographic variables (pixel slope and aspect) have a bearing on surface reflectance, subsequent ratioing of NBR mathematically cancels those factors out, so it makes no difference whether those factors are included or not. Thus in the case of calculating NBR, we use the simplified at-satellite reflectance algorithm. For more information on variable terms, refer to the **FIREMON LA Glossary**; also, the sections on Reflectance Terms and Reflectance Incorporating Topography in the **FIREMON LA How To** document.

The reflectance algorithm yields values with a theoretical valid range of 0.0 to +1.0, or if scaled by 10^3, that is 0 to +1000. They are a per-band ratio of detected surface brightness to the incoming solar radiation available at the top of the atmosphere, which allows the bands to be compared directly, as in the NBR.

Note: Landsat scenes acquired through the Multi-Resolution Land Characteristics Program (MRLC) from the USGS EROS Data Center are already in reflectance units. They do not need the processing outlined above. However, the MRLC multi-band data is rescaled to an 8-bit byte range of 0-255. That data can be used without modification. It will have a slight impact by reducing the spectral resolution of input bands, but the range and other statistical qualities of NBR will be approximately the same as using 16-bit or 32-bit data.

III. Transmittance normalization, for Band 4 and Band 7 of one Landsat scene:

This addresses the fact that atmospheric clarity varies spatially and temporally, and the ability of light to penetrate the atmosphere (transmittance) varies per bandwidth as light is scattered by particulates and moisture. If one compares multiple dates of Landsat data, such effects should be minimized to avoid influences on surface reflectance. In some cases, multi-temporal datasets are so similar in atmospheric clarity that this normalization is not necessary, and may be skipped. Steps should still be taken, however, to first determine if that is the case. There are a number of solutions, and good literature exists on performance of different methods. Here, a relative normalization of one scene to another is undertaken, considering that results are based on band ratios, and geographic scope is typically limited to subsets of scenes. A list of procedural steps follows in brief. For those wishing more rational and detail on this method, see **Atmospheric Normalization** in the **LA How To** document.

1. Determine if transmittance is a factor, and if so, which scene is most affected by atmospheric scattering of light.

2. Using both scenes for reference, acquire a sample of pixels that represents quasi-invariant targets within low, middle, and high reflectance ranges. Targets should not be subject to seasonal changes or other disturbances that would normally affect reflectance. Sample sizes should be in the range of 4 to 8 polygons totaling 600 to 1200 pixels per level of reflectance, or more.

3. Extract pixel reflectance values, for each Band (4 and 7) from each scene, and import those to a statistics package, capable of performing linear and quadratic regression.

4. Perform curve-fitting operations on each band from both dates. Use both linear and quadratic models to determine the relationship of one scene's reflectance distribution versus the other. Assign the less-clear scene to the independent variable.

5. Select the regression model that most adequately explains the intra-band differences observed between scenes.

6. Remove widely deviant pixels before further analysis of each band, as those are most likely affected by factors other than transmittance. Use the regression model just selected and a broad confidence interval per band (e.g. 98 to 99 percent).

7. Re-run the selected regression model on the new set of pixels with deviant pixels removed.

8. Use regression coefficients just derived to transform the scene with greater atmospheric effects (the independent variable, above).

9. Finally, compare results to original band reflectance datasets. The results of regression transformation, though visually subtle, will be used in subsequent NBR calculations.

Note: Atmospheric normalization is not currently done by the USGS EROS Data Center for scenes used in National programs, such as MRLC or the NPS-USGS National Burn Severity Mapping Project. This is due to a judgment that the effort would be too complex and costly for the improvement realized, given the large number of scenes involved in such National repeat coverage.

IV. Computing NBR for each Landsat scene, prefire and postfire

At this point one has pre- and post-fire reflectance datasets for TM/ETM+ bands 4 and 7. If need be, one of the datasets has been transformed by regression to normalize for atmospheric effects. The calculation of pre- and post-fire NBR is then straightforward:

Eq. LA-5 $$NBR_s = (R4_s - R7_s) / (R4_s + R7_s)$$

Where, NBR_s is the per-pixel normalized burn ratio for scene *s*, and *R4* and *R7* are the calculated reflectance quantities, as described above, for the respective bands per scene. Note, the NBR has a theoretical range of -1.0 to +1.0, or if scaled by 10^3, -1000 to +1000 (figure LA-10).

V. Computing dNBR from the pair of NBR datasets

The NBR difference is then computed:

Eq. LA-6 $$dNBR = NBR_{prefire} - NBR_{postfire.}$$

This integrates multi-temporal NBR datasets into a single gradient, or a one-dimensional scale. The difference measures change in NBR that occurred from time before fire to time after fire. The dNBR has a theoretical range of -2.0 to +2.0, or if scaled by 10^3, -2000 to +2000 (figure LA-10).

NBR Responses

By understanding how individual TM/ETM+ bands respond, one can grasp relationships of the NBR to burn characteristics. The NBR incorporates Band 4 reflectance (*R4*), which naturally reacts positively to leaf area and plant productivity, and Band 7 reflectance (*R7*) that positively responds to drying and some non-vegetated surface characteristics. Band 7 has very low reflectance (it is absorbed) over green vegetation and moist surfaces, including wet soil and snow, just the opposite from Band 4.

Since NBR measures the difference of *R4* minus *R7*, it is positive when *R4* is greater than *R7* (figure LA-10). This is the case over most vegetated areas that are productive. When it is near zero, *R4* and *R7* are about equal, as occurs with clouds, non-productive vegetation (cured grasses), and drier soils or rock. When NBR is negative, *R7* is greater than *R4*. This suggests severe water stress in plants and the non-vegetative traits created within burns. There one finds, for example, decreased vegetation density and vigor that *R4*responds negatively to, coupled with increased exposed substrates and charred fuels, which *R7* registers positively. Charring of living and inert components, drying, and soil exposure enhance the signal registered by *R7* in comparison to *R4*. Results over recent burns, then, typically show near-zero to strongly-negative NBR.

Interpreting Results of dNBR

Continuous Data

Results are a prediction of severity. They measure the change that Landsat TM/ETM+ has been able to detect in NBR, a normalized difference of bands known to be highly sensitive to fire effects. The initial **dNBR** product is a continuous range of values that can be used directly for mapping and analysis (figure LA-10). Assign a *linear* grayscale to the range of the data, one that provides good contrast, from unburned to highest burn conditions. We use a range of -800 to +1100 **dNBR** that is assigned to the gray-level range of 0 to 255 (black to white). Unburned areas generally fall out as medium gray, and burns display a gradient of lighter grays with white at the high end. As hypothesized, the sequence of brightness corresponds to the gradient of severity.

The units of **dNBR** are dimensionless, since they are a difference of normalized ratios. We tend to speak in terms of "points" that gauge magnitude of positive or negative change in NBR. From that we infer how strongly fire has affected a site. Individual values reference conditions averaged over the whole area of a pixel. Thus, a given value may represent either uniform distribution of one severity within the pixel, or small-scale patchy distribution of multiple severities. Overall though, brightness for **dNBR** generally corresponds to a steady progression of effects, relative to the pre-fire community:

1) Increasing char and consumption of downed fuels;
2) Increasing exposure of mineral soil and ash;
3) Change to lighter colored soil and ash;
4) Decreasing moisture content;
5) Increasing scorched-then-blackened vegetation; and
6) Decreasing aboveground green biomass and vegetative cover.

Theoretically, **dNBR** (scaled by 1000) can range between -2000 and +2000, but in reality it is rare for valid data to vary much beyond -550 to +1350. (Based on the scope of disturbance factors potentially affecting natural landscapes so far encountered.)

Negative values result from post-fire NBR being greater than pre-fire NBR. This may be due to clouds in the pre-fire image, or increased plant productivity in the post-fire image. Enhanced vegetative regrowth is detected in approximately the -500 to -100 range of **dNBR**. A recent burn may exhibit this after one growing season, if severity is light and the burn is in mostly herbaceous communities that recover quickly to exceed the productivity existing before fire. Also, older burns exhibit this as they recover vegetatively from the first year post fire into subsequent years. Pixels below about -550 are likely cloud effects, or noise caused by miss registration or anomalies in original Landsat data. Extreme negative (or positive) values also appear where data from one scene overlaps missing data in the other scene, as occurs near scene edges.

Ignoring the extremes, typical unburned signals fall approximately in the range near zero, -100 to +100. This indicates relatively little or no change over the time interval. Phenological differences between pre- and post-fire scenes can shift the distribution of unburned values, sometimes as much as 50-100 points.

Positive values occur when post-fire NBR is less than pre-fire NBR (figure LA-11). These may result from clouds in the post-fire scene, or from fire effects within a burn. The latter typically occupy a range between about +100 to +1300. Values above about +1350 likely are cloud effects. Interestingly, cloud shadows do not have pronounced affect on **dNBR**. That is because NBR is a normalized ratio, and not influenced as much by brightness variation that is consistent across all bands, like that caused by shadow, as it is by inconsistent spectral differences between the bands. Cloud shadows, though, tend to boost **dNBR** slightly when in the pre-fire image, and decrease it slightly when in the post-fire scene.

Burn Perimeter

One of the first things to do with the continuous data is to interpret the burn perimeter. Software may be available to automate this, but even so, those results should be reviewed, and then manually edited if necessary. If one does not have much knowledge of the burn, it is highly recommended to consultat with others who do. Discrimination of burned area is enhanced when guided by direct field observation as much as possible. The objective of the perimeter, as we see it, is to delineate polygons that minimally encompass all burn areas. They can be used for graphic purposes, to calculate area statistics, or as a mask for isolating burn areas in GIS overlay processes (figure LA-12). A manually digitized perimeter is quick and suitably accurate for 1:24,000 mapping. It provides a good way to plan sampling, or communicate information about burn size and distribution promptly. With a linear grayscale image of **dNBR** displayed on the computer monitor, digitize on screen, following the boundary of the burned area. Zoom up to be able to faithfully follow the edge. You will need to decide on a level of generalization for the perimeter, since the actual boundary can be quite complex and convoluted. The amount of detail is a matter of scale, limited by data resolution and intended use. It is useful also to have the pre- and post-fire false-color composite images to refer to on screen.

As rules of thumb, try to retain the obvious character of the shape of the burn when displayed at a scale where individual pixels are not so obvious, e.g. 1:24,000. Do not attempt so much detail as to be outlining each individual pixel. By and large, err on the liberal side, try to stay one-half to a whole pixel outside the burn, and pass across small peninsula of unburned areas that project into the burn, if less than about three pixels wide. Also, generally do not digitize around interior unburned islands (unless some specific objective requires it). In all these cases, the burned and unburned pixels included within the digitized perimeter should be correctly identified eventually by their **dNBR** values, as when it comes time for statistical summaries of the burn. Continue digitizing all disjunct patches of the burn created by spot fires, and label all polygons with a unique identifier.

Subsequent Procedures - Using the digitized perimeters, one can extract a histogram of all **dNBR** pixel values occurring within the burn(s). This is a basic reference for comparison to other burns (figure LA-13). It supplies information on mean and variance of severity, and the frequency (or area) of values occurring across the **dNBR** gradient. It also may be used to break out the number and aerial extent of burned patches.

Next, you will be able to identify target areas for field sampling, with the perimeter, trails, roads, and other references (lakes and streams) overlaid on the grayscale **dNBR**. This allows you to find accessible areas large enough to represent the range of variability within the burn, as described in **Ground Measure of Fire Severity: The Composite Burn Index** section.

Once field plots are located and sampled, plot locations can be mapped onto the continuous **dNBR** dataset using GIS overlay functions, and the pixel values from those locations extracted. There are many creative ways of doing this, including multi-point averages or weighted averages within a local neighborhood of the plot. Due to burn heterogeneity and improved GPS locational accuracy, we look for alternatives to the commonly used 3 by 3 pixel average. We find that a straight average of 9 pixels too often interjects values in the average that are greatly dissimilar, and obviously not representative of the plot. One option is to weight the center pixel by 2 or 3 times, and throw out one pixel that has the most different value from the center pixel. Instead, however, we tend to use a five-point pixel average, where the points for sampling **dNBR** are the plot center, and theoretically plus or minus 15 meters from plot center. This results in 1 to 4 pixels being sampled per plot, depending on the juxtaposition of the plot center to the pixel center. The center pixel (the one containing the plot center) is always counted at least twice, providing extra weight to the center pixel in the average. If the plot lays dead center within a pixel, then this five-point sampling yields the value of only that one pixel.

The **dNBR** values extracted for all plots then can be imported into a database that contains corresponding plot CBI ratings, or other measures of severity determined in the field. At that time, analysis of the association between **dNBR** and observed severity can be undertaken. Fieldwork may also detect where revisions to the perimeter are needed. Each subsequent step adds a level of verification that should be specified in the metadata.

Discrete ordinal data

Continuous **dNBR** datasets can be stratified into ordinal classes, or severity levels, to simplify description and comparison of burns (figure LA-14).

The breadth and number of levels is entirely up to the user, based on requirements of the application. However, we commonly employ a seven-tiered configuration proven useful in a variety of ways (table LA-2). Value ranges of **dNBR** may vary between paired scenes. Values less than about -550, or greater than about +1350 may also occur. If they do, they are not considered burned. Rather, they are masked out as anomalies caused by mis-registration, clouds, or other factors not related to real land cover differences.

The first two severity levels (table LA-2) reflect areas where productivity increased after burning. They occur almost exclusively in herb communities where **dNBR** can be strongly negative from enhanced productivity after fire (i.e. the post-fire NBR is much greater than the pre-fire). Typical unburned pixels occupy the range near zero. The last four levels include all other burned areas where **dNBR** is distinctly positive (i.e. the post-fire NBR is much less than the pre-fire). They cover what is normally recognized as recently burned, including forest, shrub and some herb communities.

Ordinal or nominal classes such as these are useful for wide array of purposes, like reporting aerial statistics; aggregating the statistics of many burns; stratifying for study of ecological consequences or treatment; quantifying burn heterogeneity; and mapping. Ordinal or nominal classes, however, can be quite variable case to case, depending on each projects objectives and individual perceptions of burn severity.

Severity Thresholds

Unfortunately, the threshold levels reported above are not hard and fast for all **dNBR** scenarios. They are somewhat flexible. Recent experience shows shifts for some burns in the range of about ± 10 to 100 points for a given severity level. At this time, we believe the primary causes for this variation are 1) seasonality of the images, and 2) whether the timing is for initial assessment or extended assessment. Thresholds tend to elevate for early-to-middle season **dNBR** under extended assessment, as that exhibits greater range overall compared to late-season extended assessment. On the other hand, initial assessment may indicate considerably higher severity, requiring higher thresholds, when the post-fire scene comes soon after burning, as opposed to the following growing season. When the post-fire scene is drier overall than the pre-fire scene, the burn-unburned threshold tends to raise somewhat, and there is greater chance of confusion between the driest unburned pixels and lowest-severity burned pixels.

If image timing is indeed a controlling factor (relative to time since fire and time of year), then scenario-specific scales for severity may be possible to achieve in the future. In the mean time, fine calibration of thresholds can only be done individually for each **dNBR** model, using a combination of expert knowledge and correlation to ground data. We intend to add further guidance along these lines as the number of burns analyzed in different ecosystems expands.

Before ground data are analyzed, one can interactively color up ranges of the continuous **dNBR** data to determine preliminary severity level thresholds by computer (figure LA-13). We strongly advise that this procedure be done by or in consultation with someone who has direct knowledge of at least portions of the burn on the ground. That will greatly facilitate and improve the initial classification of severity levels, by introducing an ability to recognize spatial patterns as they were observed in the field.

Display the **dNBR** linear grayscale, and from the low end at about -100, progressively "color up" increasing values with one color. You will find the distribution of colored pixels being randomly scattered at first around the burn, then incrementally becoming gradually more localized, until mainly pixels near the edge of the burn are being colored. When the burn area is clearly delimited, yet not excessively to crop out potentially low severity pixels, the end value marks the approximate upper limit of unburned in terms of **dNBR**. The same procedure should be done in reverse, from the top down, to find the bottom threshold of low severity level. Then compare the burned-unburned endpoints, and revise the threshold as needed.

Do not be alarmed if some spurious "low severity" pixels show up well outside the perimeter where it did not appear to burn.. That is to be expected when setting discrete boundaries for categories based on continuous data. On the other hand, you want to reach a threshold where those spurious pixels are at minimum, and at the same time have a "burned" pixel distribution that most faithfully defines the actual burn area. The quality of severity-level discrimination can be determined later statistically, when final judgments about reliability can be weighed with field data. At this time, however, some idea of how well **dNBR** is working can be inferred by examining how well the burn-unburned threshold either includes or excludes the burn, based on the distribution of seemingly errant pixels relative to correct ones.

Most errant pixels should be scattered randomly, and relatively few in number. If there are some well-defined patches of pixels that seem to be wrong, the most likely cause is difference in moisture or phenology between the two Landsat scenes. For example, unburned meadow patches may be lighter than surroundings if the post-fire scene is particularly drier than the pre-fire. In a case study of Yellowstone National Park burns at elevations above 6000 feet, where early September and mid-October NBR values were differenced, the October scene was notably drier. Deciduous plants had become dormant in meadows and along riparian corridors, and those appeared faintly with elevated **dNBR** (figure LA-7). If location of the burn is generally known, these cases should not be cause for concern. However, thresholds may need to be shifted up or down, depending on relative scene conditions indicated by unburned background.

Note, there is always overlap in **dNBR** values between severity levels. To explore statistical properties of a particular **dNBR**, sample the population of unburned pixels. Large samples of unburned pixels (5,000 or more) typically display normal distributions. Include only areas that are not affected by extreme phenological or atmospheric differences, such as were a meadow was green in one scene and completely brown in the other, or where there was snow in only one scene. If such differences predominate over the unburned area, then the scene pair is probably not appropriate to use in the first place. The sample mean should indicate the bias from zero **dNBR**, which marks the theoretical value for no detectable change. If the unburned mean is shifted from zero, that indicates some non-fire-related difference between the scenes, and the whole dNBR histogram could be shifted by that amount to standardize multiple **dNBR** scenarios (shifts up to ± 50 points can occur). Furthermore, two standard deviations should mark an approximate 95% confidence interval for the unburned class. From our experience, a standard deviation of about 50 or less indicates a good pairing of the **NBR** scenes, with little non-fire-related difference between them.

Next, it is useful to establish a lower threshold for high severity. Change to a different color, and start at the high end of **dNBR** to progressively color in values decreasing from the highest. You will notice a point (value) where the first few pixels within the burn take on the color. This is the upper realistic limit of **dNBR** for that burn. If the burn is one that seems to include some extreme high severity, that value should be within a range of about +1100 to +1300. If that value is lower, it may be that severity did not reach maximum levels within the burn. If it is below about +600, it is possible only moderately high levels were reached, and the burn contains no high severity. If it is above about +1300, bear in mind it may result from clouds in the post-fire scene, or some other data anomaly. Examine the raw Landsat scenes to confirm that very high values are not where clouds interfere with **dNBR** in the burn.

From that highest realistic value, continue coloring values incrementally down the scale. You will notice colored pixels appearing in a number of new locations, and then progressively fewer new locations. You will see increased clumping of newly colored pixels around previously colored patches, gradually appearing to fill out and expand existing patches. Isolated uncolored pixels within colored patches may indicate you are missing some high severity, and the threshold should go lower still. At some point after patches are well formed, but before most of these patches coalesce, and before lone pixels start to frequently show up in dispersed new areas of the burn, note the **dNBR** value. If it is within a range of about +600 to +750, you can assume you are within the range of a preliminary lower threshold for high severity. Based on knowledge of the burn, aerial photos, or some other source you may want to go further down or back up the scale to try to settle on the most reasonable breakpoint. This may be influenced in part by the size of the area identified as high severity. Recall, there may be tendency for Initial Assessments to have higher thresholds for high severity, depending on the ecosystem.

If there are a few isolated pixels that are not colored within larger surrounding patches of "high severity", then that may indicate the threshold for high severity is too high, and needs to be lowered. In fire ecology terms, one might question the validity of a few isolated pixels remaining within large patches of high severity. Query the values of those pixels to determine if they should be included with high severity, as would be likely if those values are only a few points below the current high threshold. If they are quite a bit lower than the current high threshold, then those truly may be isolated areas of a lower-level severity class. In the end, one should see a distribution of high severity that is not excessively fragmented or "speckled", while at the same time, not too broadly contiguous over an unreasonably large portion of the burn.

Apply the range between unburned and high severity just determined to partition the remaining positive severity levels. Refer to table LA-2 for proportions of those levels relative to a comparative span of about 560 points between unburned and high severity. If reliable information on the burn is available, the levels can be adjusted to fit what is known. For example, one may know that certain areas burned with low severity, so thresholds can be adjusted accordingly to correctly identify those areas.

The Enhanced Regrowth levels (strongly negative values) can be determined much as described above by reversing the progression of coloring the negative values. Start by going down from unburned into the Enhanced Regrowth Low level. Focus on areas within the burn where **dNBR** appears darker than the unburned medium-gray levels outside the burn. These should correspond to meadow or grassland habitats, if present. It helps to key on the distribution and shape of familiar meadow patches evident in false-color images or aerial photographs. The range of valid negative values is basically divided by one third and two thirds to split out low and high enhanced levels, respectively, assuming the total is a span of about 420 points or more. If much less than that, retain a span of about 100 points for the Enhanced Regrowth, Low level.

Subsequent Procedures - The initial stratification of the burn can be analyzed for preliminary assessments, or taken into the field where adjustments to the model thresholds may become evident. Walk- or drive-through surveys are recommended with severity maps in hand to initially spot check for obvious agreements or discrepancies. More in-depth fieldwork is usually required for statistical validation and calibration, and the map of burn levels is useful for locating target sample areas. Refer to the **Ground Measure of Fire Severity: The Composite Burn Index** section for discussion of field protocols. The **dNBR** thresholds can then be revised as soon as consistent field observations are made. Once sufficient ground data has been analyzed, that should ultimately guide where thresholds most appropriately fall, based on statistically determined intervals.

Recognizing that ground data may never be available on some burns, thresholds for severity levels (classes) may need to be based on field results from other burns within similar ecotypes and similar timing of **dNBR**. After a number of burns have been sampled in a region, statistical confidence in thresholds for **dNBR** should increase to a point where subsequent ground data is less essential. At that time, plots can be sampled less frequently and used mainly to spot check results. That is, in fact, one goal of the whole process, so field time and expense can be minimized, without impacting availability and reliability of burn information.

Each subsequent procedure adds a level of validation that should be documented in the metadata. Once results are improved and verified as much as possible, the burn severity model can be used to compile final reports and statistics, and to address a variety of issues.