Some thoughts on sampling dNBR to derive the mean value for unchanged pixels.

The goal is to improve standardization of dNBR for more valid comparison and analysis of different dNBR datasets:

> The objective is to find a dNBR value that best represents the unchanged condition between pre- and post-fire scenes.

> Theoretically, the unchanged condition should have a dNBR mean value of zero, signifying no difference between pre- and post-fire NBR.

> When unchanged areas are sampled, however, the dNBR mean value is usually something different than zero.

> NBR is influenced by pervasive ground, solar and atmospheric conditions that vary between Landsat acquisition dates.

> We hope to mitigate these influences that generally affect the dNBR in the vicinity of burned areas.

> NBR is also influenced locally by factors such as logging, vegetation disease or stress, drought, phenology, fire and other disturbance.

> These local factors are changes that we expect to see or detect when NBR is differenced, and are not the target for this standardization.

> The reason clearly discernable local variations are avoided is because the bias will be applied to the whole image, and such variations may not represent the general unchanged area surrounding the burns; i.e. they may unduly weight the bias in inappropriate ways for the whole image.

> Correction of local variations would require a spatially explicit application of variable unburned biases, which we are not prepared to do. Moreover, this would be impossible within the burn, since fire would have negated any local variations from other causes.

> On the scale of a Landsat scene, moisture content and phenology often show local influences, but conditions may also be more pervasive.

> In sampling, we want to represent the pervasive and most typical moisture and phenology condition, while avoiding the irregularities.

> To improve comparability, we offset each dNBR by its mean unchanged value; so all corrected dNBR's have approximately the same starting point (zero) from which to measure change. We want the unchanged offset, then, to provide a baseline for gauging other changes.

> We have to sample the unchanged areas of a dNBR in order to determine the mean bias value in the vicinity of the burn.

> Remember, it is not only unburned areas to sample, but more specifically, unchanged areas that are not visibly altered by local influences.

This example starts with the 2005 Southern Nevada Complex (SNC). Many of the techniques and concepts, however, would apply generically to other areas, as well as to other change detection indices.

Suggested preferences to set up for thresholding:

The dNBR should have an appropriate grayscale to display the continuous data. In dry areas like the SNC, less range and contrast is expected in dNBR, so breakpoints of -600 to 1000 seem to work well, compared to -800 to 1100 that may be more suitable in moister areas.

Open the dNBR as a continuous grayscale image with the fire perimeter overlaid in one viewer, and side by side with a second linked-viewer containing pre- and post-fire Landsat false-color composites. False-color images also should have appropriate breakpoints so they do not appear overly saturated with high contrast. These are referenced only when needed to clarify possible causes of change detected in unburned areas.

Starting out it is good to see as much area as possible, but not so much that the data resolution is significantly degraded by the screen resolution. The Imagine™ reset button sets the zoom to the default scale, which is one data pixel for one screen pixel. For large burns, one can zoom out one click below the default and then reset as needed for detail. I would not recommend working below that scale; rather, one can pan around images to see different areas.

The following slides contain tiff images rather than jpegs, so they can be enlarged to 150% to see more detail.

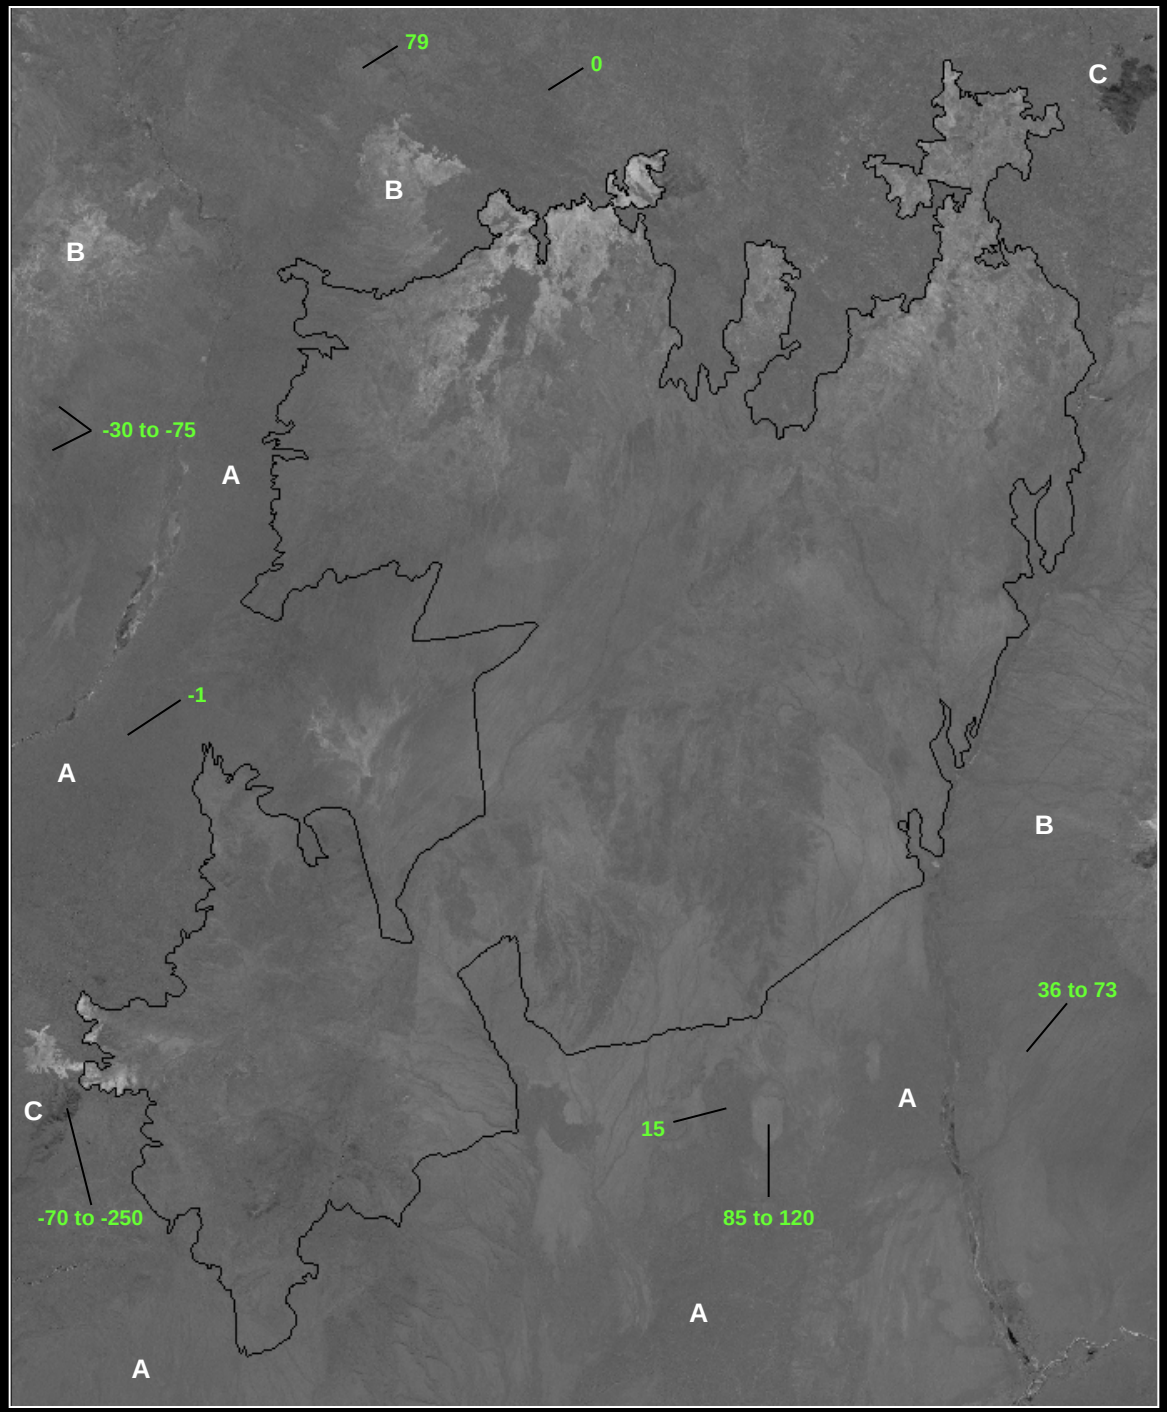

Extended Assessment 2005 Duzak fire. Subset of the Southern Nevada Complex. Grayscale dNBR (breakpoints -600 to 1000).

First, scan around outside the burn perimeter and try to find the most sizeable areas that are relatively uniform in medium gray tone (A).

This example shows a fair amount of local variation outside the perimeter, considering expected changes are more or less natural, and human influences are minimal year to year. One can expect greater variation near communities or where agriculture and logging occurs.

The unburned variation signals that scenes are not optimally matched, and more caution is needed to select suitable samples for the neutral, unchanged condition.

One expects typical unchanged dNBR values to be near zero, plus or minus around 50. In this case, dNBR at various locations often exceed that (green). Moreover, unburned areas reveal distinct spatial patterns.

The variation comes from change in NBR between the pre- and post-fire Landsat acquisition dates due to various causes. The false-color images can help you understand the temporal dynamics.

Most likely we see localized differences in greenness and moisture, which are often factors in dry ecotypes. The post-fire NBR reflects a drier, less green condition where lighter unburned pixels occur. Post-fire erosion similarly may appear, often as recognizable drainage features, including small alluvial fans. We also have some 2005 burn area that was not mapped within this fire perimeter (B).

Dark areas signal where post-fire NBR reflects a moister or greener state than pre-fire. Here, the largest dark areas (C) likely are burns that existed on the pre-fire image (May 25, 2005) and have greened up by the post-fire date (May 12, 2006). Other small dark areas in drainage-ways may be from irrigation or ephemerally wet sites.

These are all situations to be aware of and to avoid in sampling for the unchanged mean value.

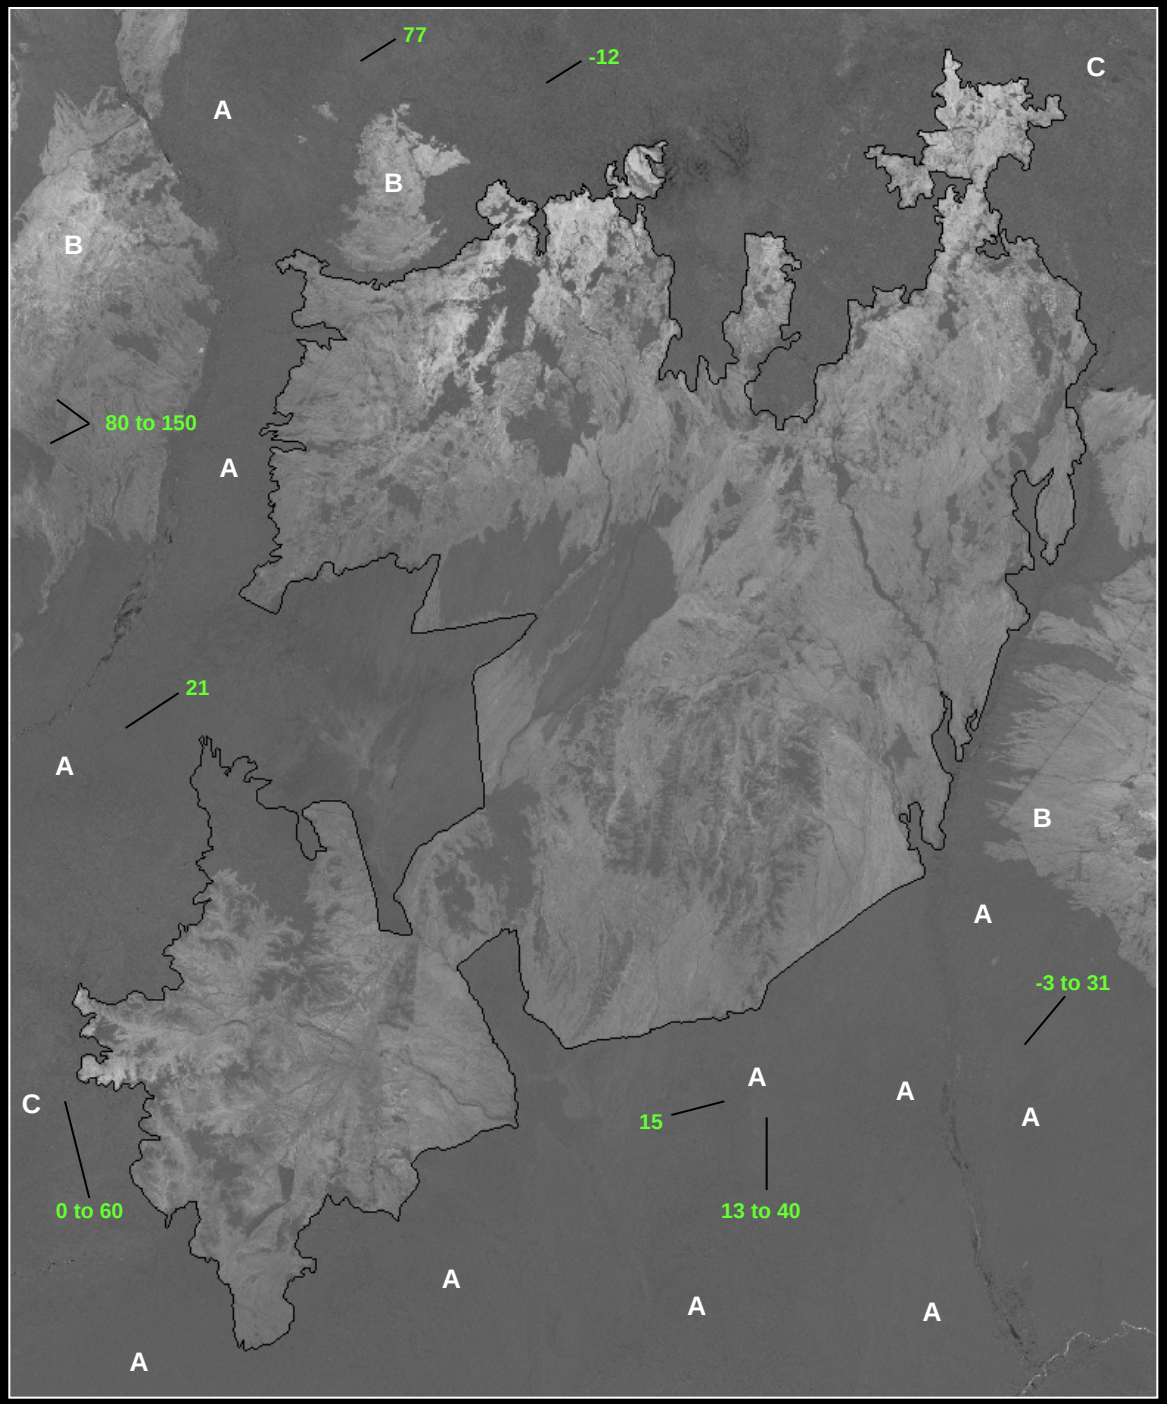

Initial Assessment 2005 Duzak fire. Subset of the Southern Nevada Complex.
Grayscale dNBR (breakpoints -600 to 1000).

Impressions of the IA dNBR are that it has better matched scenes, showing <u>much less</u> pattern variation in the unburned. Looking at other areas beyond this subset confirms this.

There is generally less numeric range of dNBR in those areas that stood out in the EA unburned (green).

We distinctly see the other 2005 burns outside the perimeter (B), while the larger dark patches on the EA do not appear here (C). Why is that? Referring to the IA false color composites, we can confirm the older burns appear on the pre-fire May 25, 2005 scene, and by the post-fire July 12, 2005 scene they look essentially the same, so little change in NBR. By contrast, the 2006 EA post-fire scene shows green-up on those areas.

Some of the small dark patches in drainage-ways also occur in this dNBR, signaling irrigation or ephemerally wet sites.

Now there are many more large, uniform areas from which to sample for the neutral, unchanged condition. There is less local variation to avoid, and we can consider more areas (A).

Go back and forth with last slide to view differences.

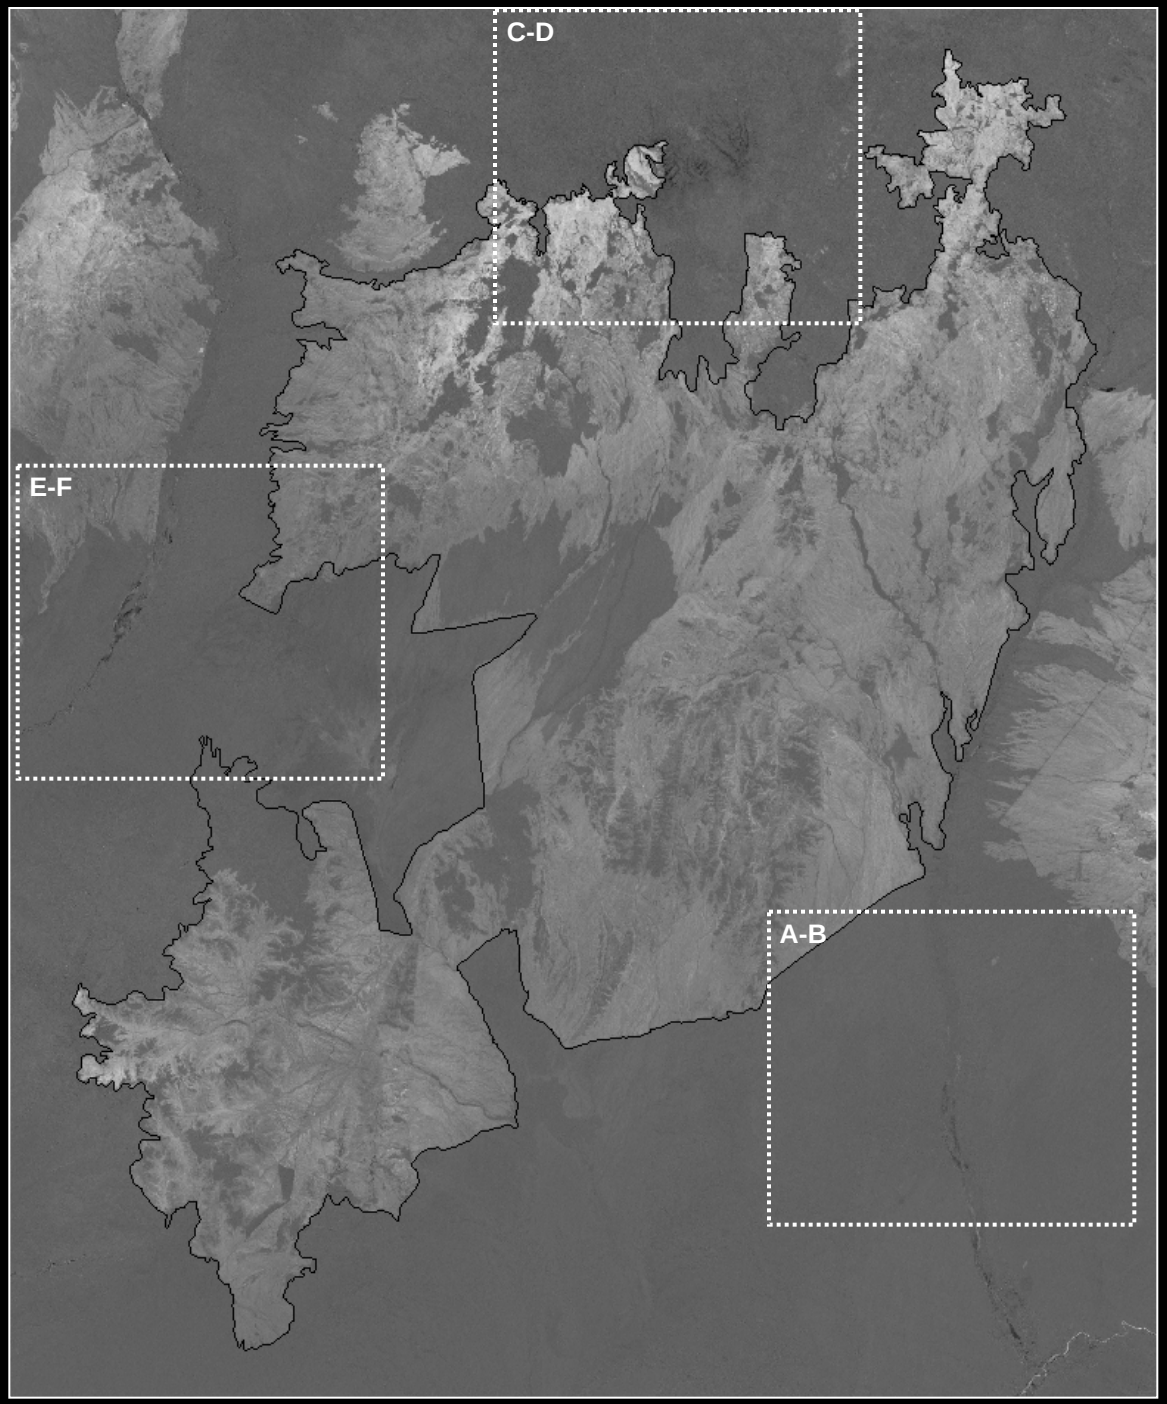

Selecting sites.

With some general considerations out of the way, here are some examples of good and bad sample areas.

The rectangles are enlarged in the following slides.

In the Imagine™ viewer with the dNBR, a new AOI is opened and AOI tools will be used to digitize sample areas on screen. Later, the final sample set of AOI's are selected, and the dNBR image statistics are re-calculated to get the mean value from the AOI population of pixels, using the image info utility.

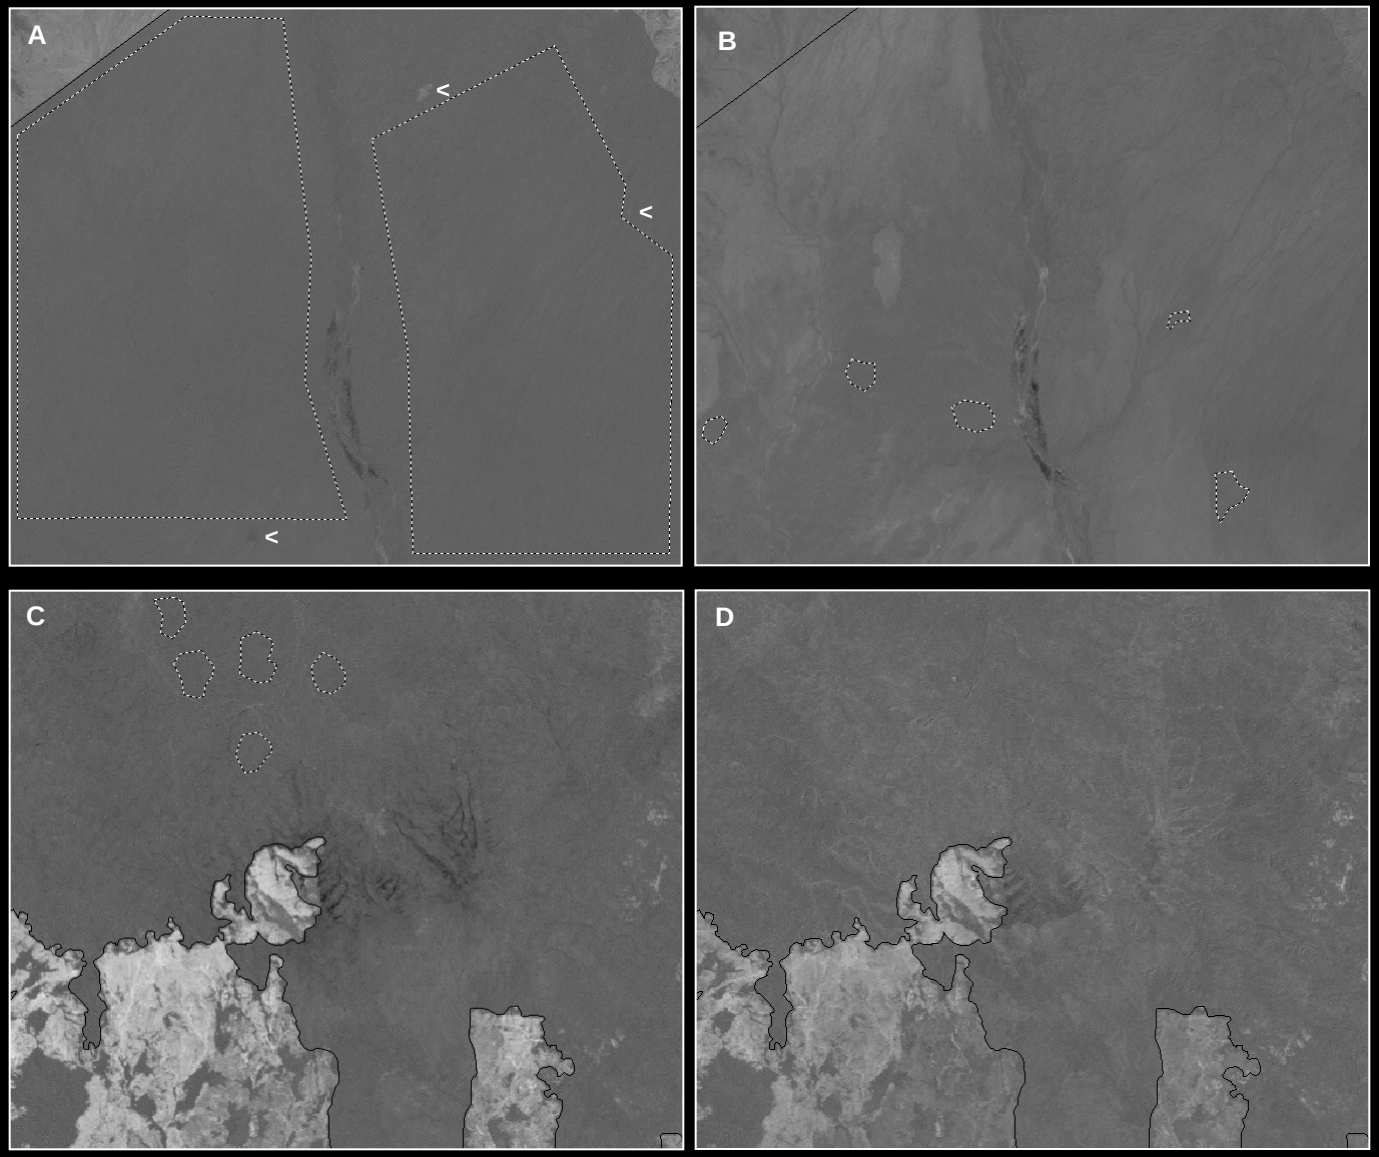

Selecting sites.

If viewed at 175%, the slide images are at similar scale as originally in the viewer. For demonstration the viewer was zoomed in slightly greater than default.

A) From the IA. Anywhere in the left AOI is very good. The right AOI is also good, but it has very faint light/dark streaking. Since that is relatively uniform and pervasive, it is not necessarily a problem, but define a large enough AOI polygon to balance this micro-variation in the sample. Avoid small spots that are much lighter or darker (<).

B) From the EA. Much of the area is unsuitable. A few acceptable spots can be found, but they must be kept small.

C) From the IA. I would probably avoid this area as well; it is a lot of work dodging bad spots to create many little AOI. Zoomed up, there is a lot of fine pattern like rippling, and random spikes of low and high values. Plenty of better areas are available.

On the other hand, it is worthwhile to place samples all around the burn, and most of this northern area looks difficult. So, because of location around the burn, some AOI like these may be worth the effort.

D) From the EA. I would hope to avoid even trying to sample this area. There is a lot of variability and pattern in the unburned area.

It is good practice to avoid areas where you have to make many decisions about what is truly unchanged.

Note, in all the AOI, try to avoid little stream beds and other features that stand out.

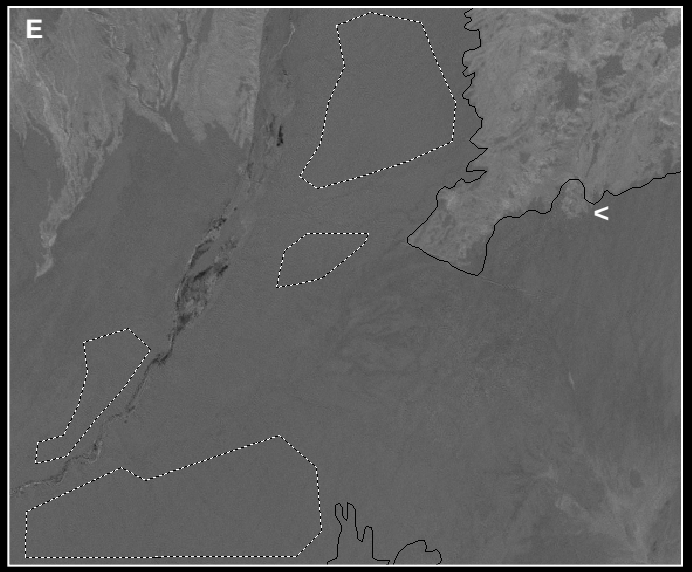

Selecting sites.

If viewed at 175%, the slide images are at similar scale as originally in the viewer. For demonstration the viewer was zoomed in slightly greater than default.

E) From the IA. Basically the same AOI can be sampled in both IA and EA dNBR. Size could be expanded in IA, however. These are good uniform areas; slight variations will have negligible effect on the mean due to the large number of pixels.

F) From the EA. One of the few regions from which sizeable AOI can be sampled. Although these AOI provide a large number of pixels (about 40,000), other areas around the burn still should be sampled. In order not to swamp the overall mean by these pixels, however, AOI size or number here may need to be cut back some, assuming the other areas can not provide such large AOI. That way there is more or less equal representation around the burn.

< Note a small exclusion of the perimeter.

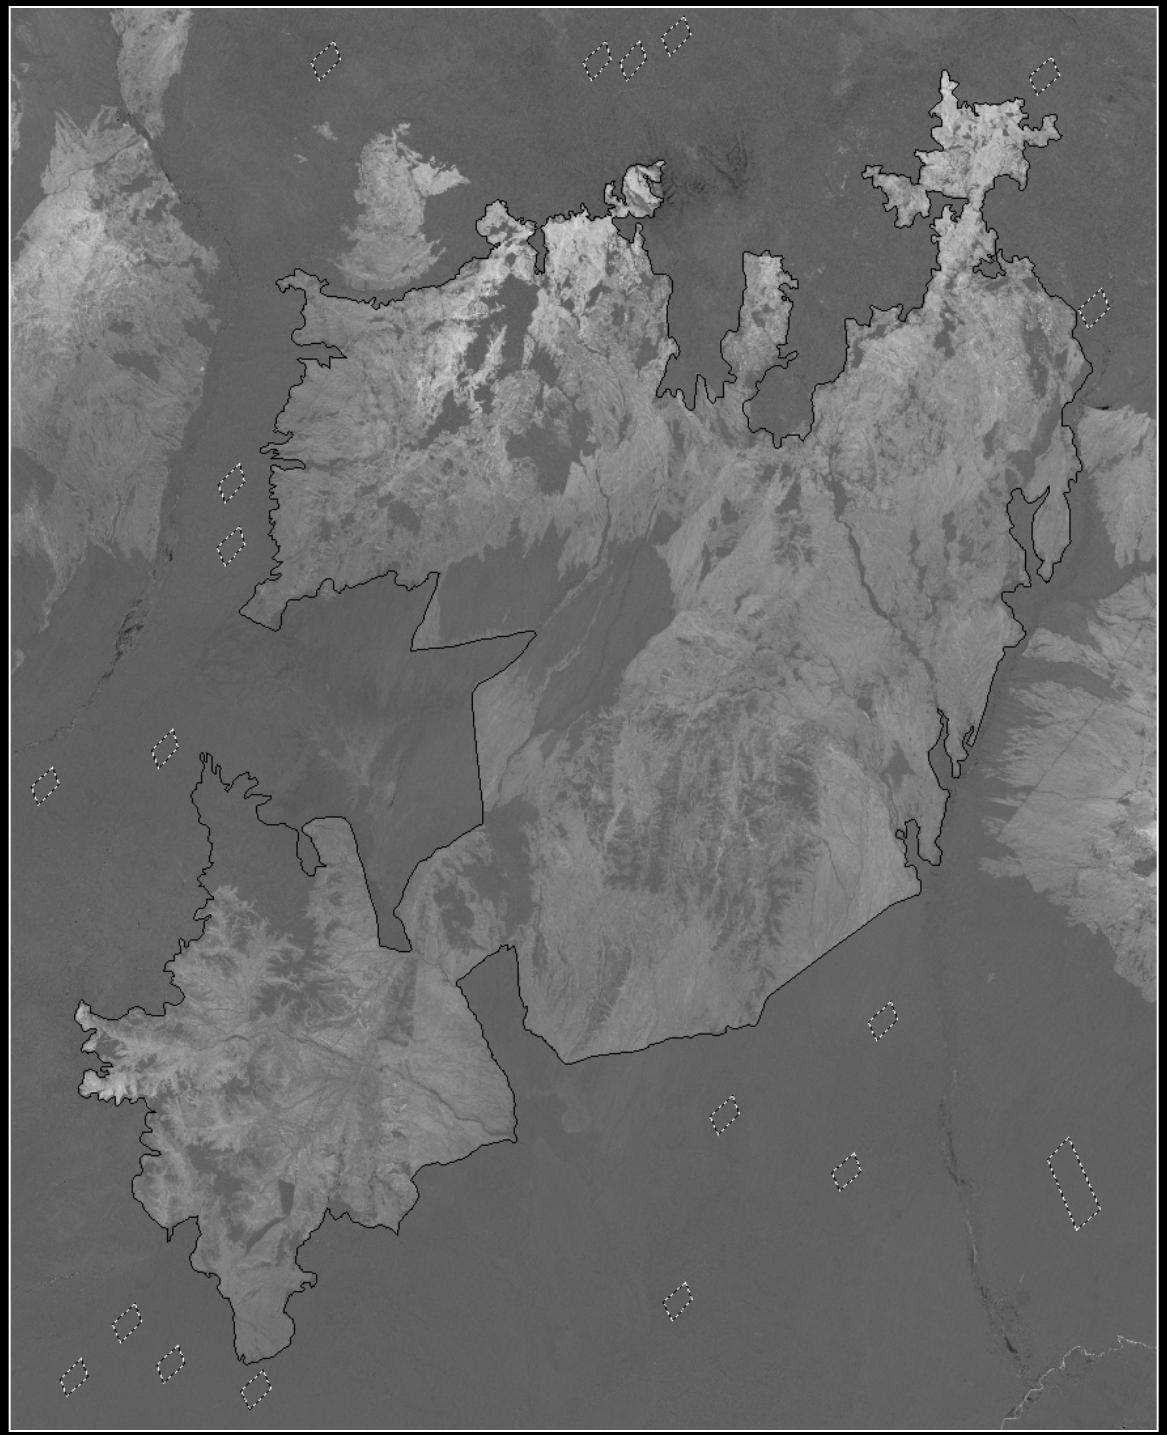

The Total Sample IA.

The previous AOI's outlined general areas where it looked suitable to sample. Now, we consider a complete lay-out of actual samples for the burn.

As far as overall sample size is concerned, at least 10,000 pixels is recommended. That sounds like a lot, but it's not such a large area. The upper AOI in the previous slide (E-F) is about 12,000 pixels.

More pixels are desirable, as well as large AOI. This tends to lessen the influence of included undesirable pixels on the mean. Larger AOI are also more practical and efficient to create. But do not sacrifice quality by including a lot of garbage in order to get larger sample sizes.

The SNC IA lends itself to ease of sampling with larger AOI. Certainly the AOI in previous slides could have been used. However, to emphasize getting samples from around the burn, and to illustrate a technique that may save time in more difficult situations, another option is illustrated.

Create one or a couple standard-shape AOI and just copy and paste those (several at once) to various localities. By pasting similar numbers of these in each region, one can try to equalize the number of pixels sampled around the burn.

The AOI object needs to be small enough so it does not include unwanted variation where the smallest target patch is located. It is helpful to find that smallest target first. I started with one of the north patches. If the AOI object fits there, the same object should fit in other locations, avoiding further adjustments.

You can convert the square-meter size of the first AOI (through AOI element properties) to pixels, and determine how many are needed. In this case, one AOI ≈ 1288 pixels. About 4 samples were taken from each region. I ended up with 22 or over 28,000 pixels total. Good enough; you could do more if you like.

I also created an AOI (lower right) which was 4 times the original to get more pixels in one. I could have used 4X versions in the other expansive areas. One could even double that for some of the larger areas; many configurations are possible.

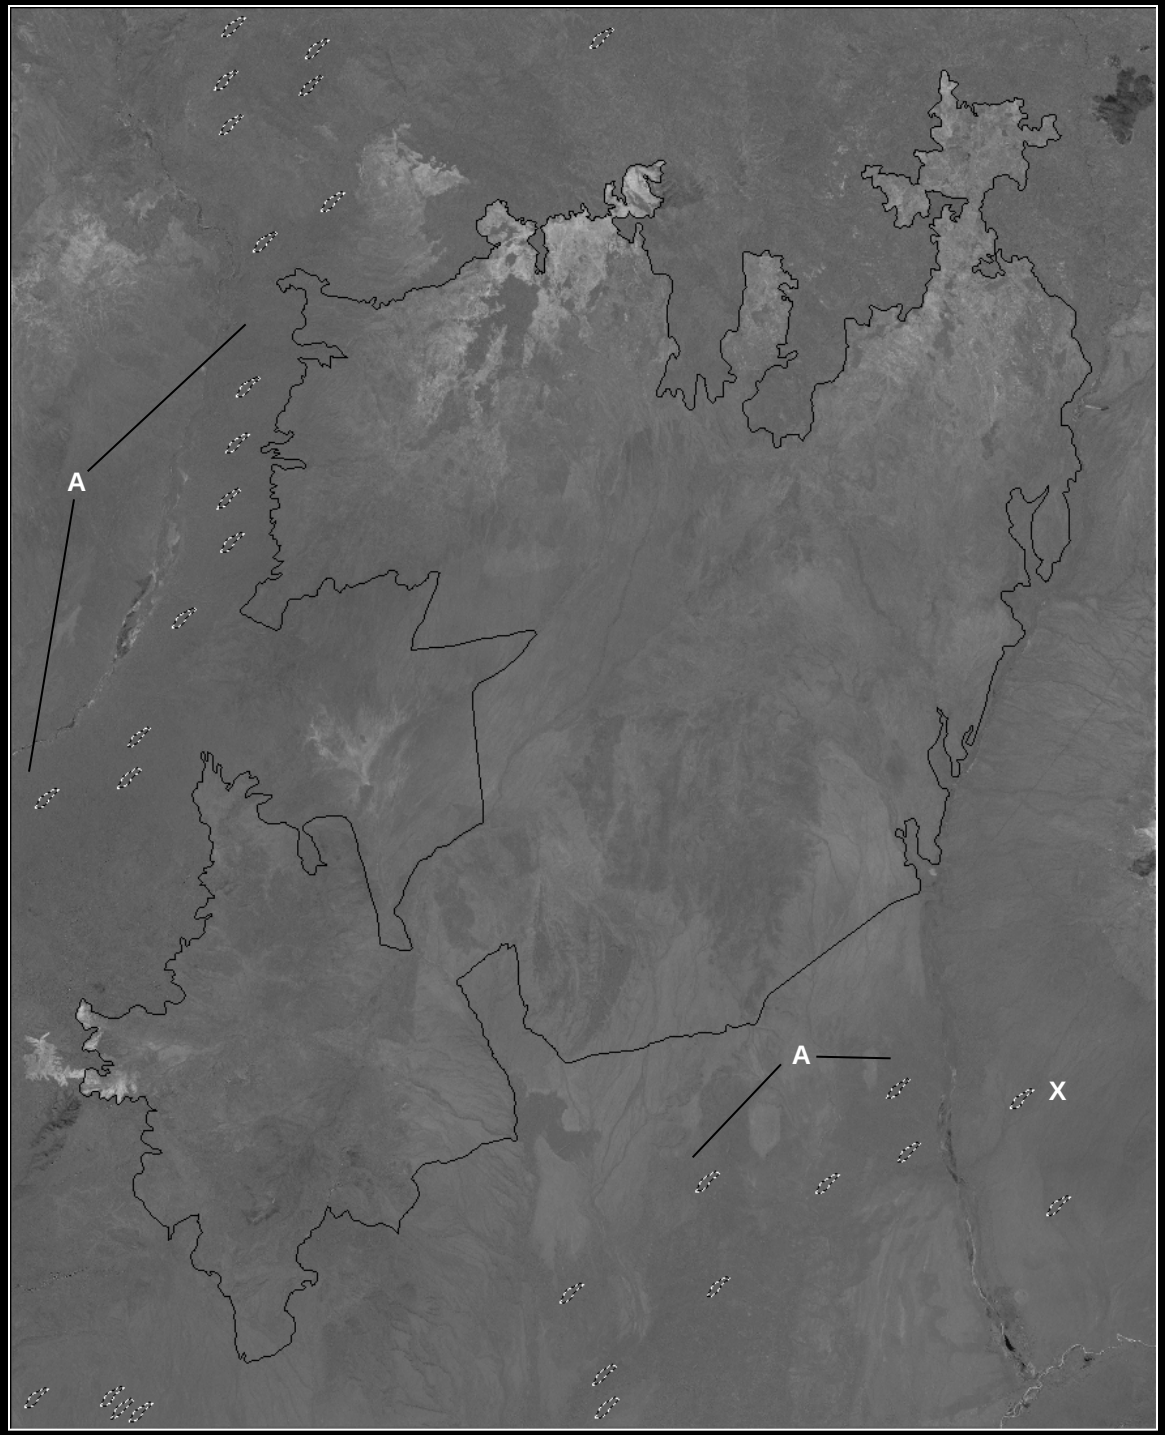

The Total Sample EA.

The SNC EA appeared more problematic, so we saved the best for last

Being able to distribute samples around the burn is more of a concern for this EA, given the difficulty finding suitably large unchanged patches.

Basically the same approach was used as for the IA, but the first AOI was quite small (≈ 388 pixels at X). More samples were needed, but that was not as much a problem as finding the fairly tight spaces in which to locate them. The small size was advantageous for that purpose. I ended with 30 AOI, or over 12,000 pixels total. Many fewer than the IA, but probably sufficient.

Pasting groups of these AOI and moving them around a zoomed-in area proved to be fairly easy. Some accommodation for larger AOI could have been made around (A) to increase sample size.

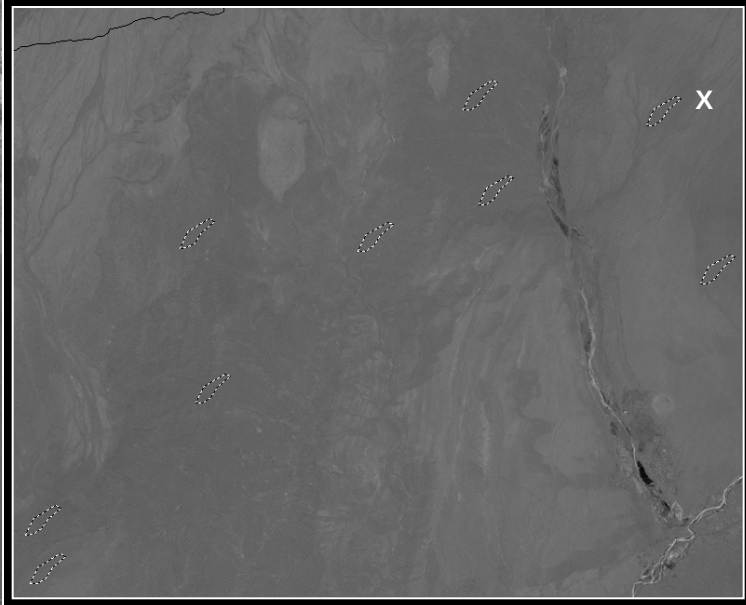

Enlargement of lower right. View at 150% for detail.

The Statistical Results.

By selecting the set of AOI for each dNBR, and recalculating image statistics (using skip factor = 1, not ignore 0, direct bin function, and use AOI), we see the long-awaited results. Well, the results are not very exciting. Both dNBR have about the same bias, +11. The Std. Dev. are good, being much less than 40, a rule of thumb borderline. For normally distributed data, which these tend to be, + or – two times the Std. Dev. yields a 95% confidence interval for any sample mean. The results are certainly not unusual, and would barely warrant any concern on their own. But if these images were compared to another dNBR that had a bias of, say -73, then the offsets may be a problem if left uncorrected. To make the correction, then, the bias is subtracted from the dNBR, so the corrected dNBR now has a mean unchanged value of zero.

The SNC IA		The SNC EA	
Min:	-113	Min:	-102
Max:	104	Max:	131
Mean:	10.8	Mean:	11.2
Std. Dev:	19.2	Std. Dev:	22.4
Median:	11	Median:	12
Mode:	0	Mode:	0

This is as far as the MTBS process needs to take the analysis just completed. There is more that can be done, however, in situations where one may want to retain the frequency distribution or pixel values of this sample population. In such cases, one can open the dNBR attribute table, and output the value and number of pixels in the sample. Select the Value and Histogram columns, use Edit Criteria to select the min and max value range, above, and then Edit Copy the data. The data can be pasted into spreadsheet or statistics datasets for further analysis. We expect the histograms to be very close to normal. Here the EA sample is slightly lopsided, but given the difficulty of finding large uniform sites and the reasonable Std. Dev., it's not too bad.

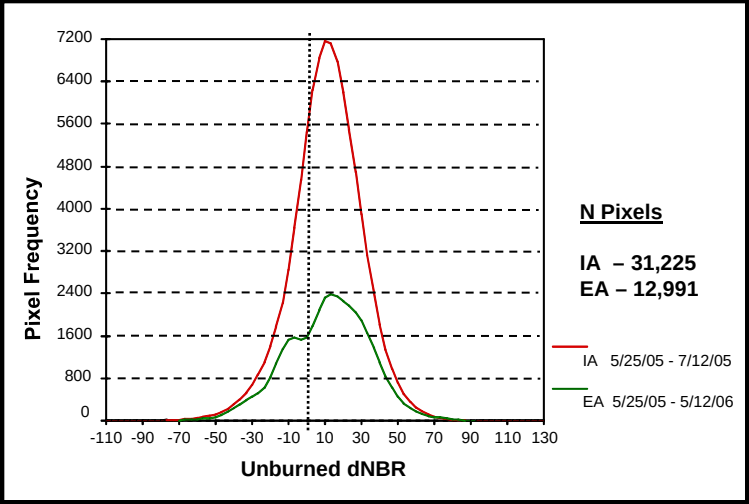

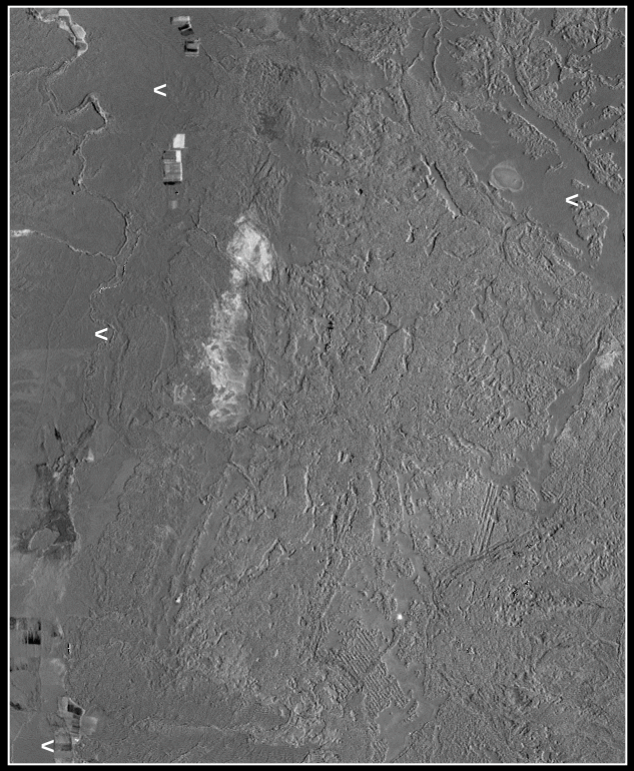

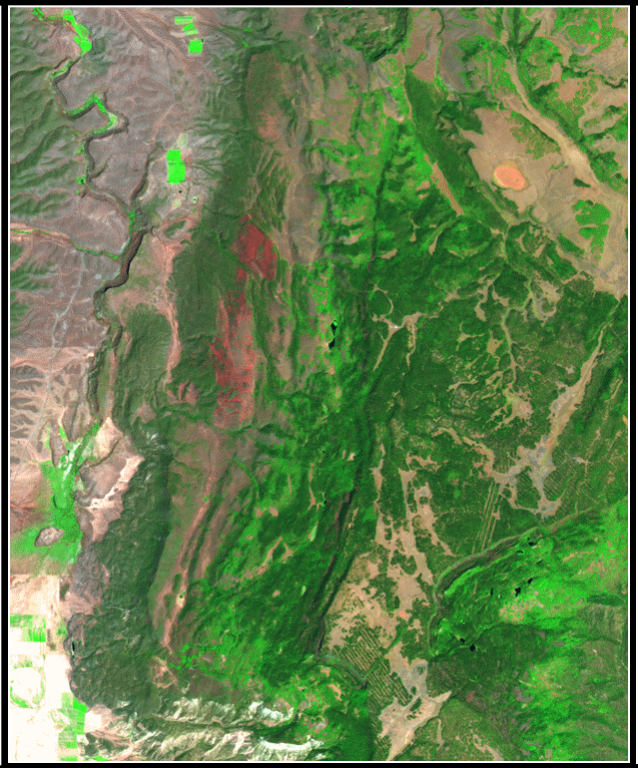

Selecting sites footnote. Miss-registered scenes.

Occasionally, you will have a dNBR that has miss-registered scenes (left). Such areas generally are not good to sample for the unchanged mean. The scenes are geographically offset from one another by more than half a pixel. Often this occurs near sides of a scene, while central portions may have good registration. This was <u>not</u> the case for the preceding SNC dNBR, but to help you recognize and avoid the condition when it occurs, this example was taken from the IA dNBR of the 2002 Sanford Fire in east-central Utah.

Virtually every edge (center) is highlighted by dark or light pixels (high – or + dNBR, left); the dNBR surface has texture, like a sunlit 3-D effect. When dissimilar patches become overlaid by faulty registration, an erroneous change is detected along the edge. Possibly some small areas around (**<**) would be okay to sample, where some fairly uniform patches are large enough to not be impacted by edge.

The situation can only be improved, however, by changing the registration of one of the scenes so that the two align better, and recalculating the dNBR (right, note, there is still some error from miss-registration to avoid). To do that takes time; iteratively flickering between the stacked scenes in a viewer, changing the top image's origin coordinates through image info, then closing and re-opening the top image until the two look co-registered. Currently, it is not practical to do this routinely for MTBS, but in cases where the procedure is feasible, it can greatly improve the appearance and quality of the dNBR product.

The End.