Evolving land cover classification algorithms for multi-spectral and multi-temporal imagery

Steven P. Brumby*, James Theiler, Jeffrey J. Bloch, Neal R. Harvey, Simon Perkins, John J. Szymanski, and A. Cody Young

Space and Remote Sensing Sciences, Los Alamos National Laboratory, Mail Stop D436, Los Alamos, New Mexico 87545, U.S.A.

ABSTRACT

The Cerro Grande/Los Alamos forest fire devastated over 43,000 acres (17,500 ha) of forested land, and destroyed over 200 structures in the town of Los Alamos and the adjoining Los Alamos National Laboratory. The need to measure the continuing impact of the fire on the local environment has led to the application of a number of remote sensing technologies. During and after the fire, remote-sensing data was acquired from a variety of aircraft- and satellite-based sensors, including Landsat 7 Enhanced Thematic Mapper (ETM+). We now report on the application of a machine learning technique to the automated classification of land cover using multi-spectral and multi-temporal imagery.

We apply a hybrid genetic programming/supervised classification technique to evolve automatic feature extraction algorithms. We use a software package we have developed at Los Alamos National Laboratory, called GENIE, to carry out this evolution. We use multispectral imagery from the Landsat 7 ETM+ instrument from before and after the wildfire. Using an existing land cover classification based on a 1992 Landsat 5 TM scene for our training data, we evolve algorithms that distinguish a range of land cover categories, and an algorithm to mask out clouds and cloud shadows. We report preliminary results of combining individual classification results using a K-means clustering approach. The details of our evolved classification are compared to the manually produced land-cover classification.

Keywords: Feature Extraction, Genetic programming, Supervised classification, K-means clustering, Multi-spectral imagery, Land cover, Wildfire.

1. INTRODUCTION: REMOTE SENSING OF FOREST FIRES AND LAND COVER

Between May 6 and May 18, 2000, the Cerro Grande/Los Alamos wildfire burned approximately 43,000 acres (17,500 ha) of forest and 235 residences in the town of Los Alamos, New Mexico (USA). Restoration efforts following the fire were complicated by the large scale of the fire, and by the presence of extensive natural and man-made hazards. These conditions forced a reliance on remote sensing techniques for mapping and classifying the burn region and surrounding vegetation. During and after the fire, remote-sensing data was acquired from a variety of aircraft- and satellite-based sensors, including Landsat 7, to evaluate the impact of the fire and begin to monitor the rehabilitation of the ecosystem.

Remote sensing of forest fires has traditionally involved human interpretation of visible wavelength and/or infrared photography. Since the introduction of aircraft and satellite mounted multi-spectral imaging instruments, e.g., the Advanced Very High Resolution Radiometer[1] (AVHRR) on the NOAA Polar-orbiting Operational Environmental Satellite (POES) series, and the Thematic Mapper (TM) and Enhanced Thematic Mapper (ETM+) instruments on the Landsat[2] series of Earth observation satellites, several physics-based and empirical algorithms for detecting forest fires have appeared in the literature. Two general approaches exist: detection of "hot-spots" and fire fronts, e.g., using thresholds on brightness temperature[3,4,5,6,7] in AVHRR band 3 (3.7μm), and mapping of post-fire burn scars. A number of researchers have investigated the use of Landsat TM imagery for measuring wildfire impact by mapping of the burn scar. For example, Lobo et al[8] apply a combination of spectral image segmentation and hierarchical clustering to the mapping and analysis of fires in Mediterranean forests. Kushla and Ripple[9] use Landsat imagery to map forest survival following a wildfire in western Oregon (USA), and investigate linear combinations of post-fire and multi-temporal TM band ratios and differences.

* Further author information: (Send correspondence to S.P.B.) Email: brumby@lanl.gov

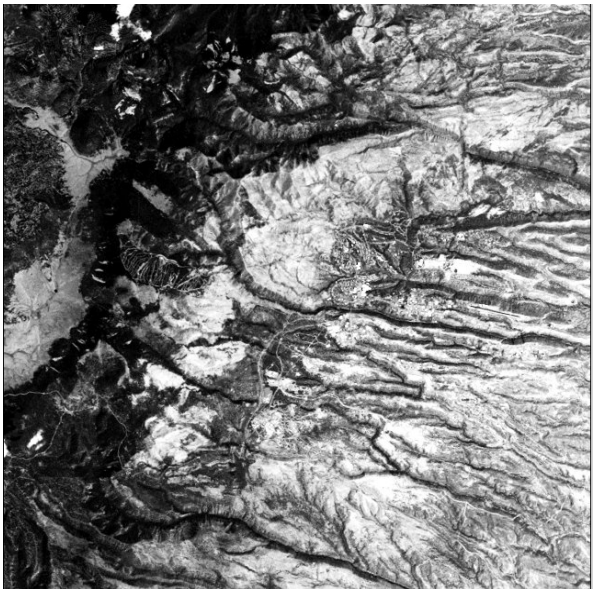

Figure 1. Post-fire, July 19, 2000: Bright region in center of image is the burn scar. Los Alamos town lies against the underside of the burn scar. Topography changes from forested mountains (left) to bare mesas.

Figure 2. BAER Team burn-severity map over topographic map: Medium gray region marks high severity burn, pale gray region marks low severity/un-burned region. This image taken from the official BAER team web site: http://www.baerteam.org/cerrogrande

Beyond classification and mapping of the wildfire burn scars, rehabilitation efforts require up-to-date forest inventories and land-cover maps. These can be used to plan rehabilitation efforts, and to estimate remaining forest fuels and hence the risk of further significant wildfires. These mapping products need to be revised on a time scale of years, as destroyed forest gives way to new plantings or as erosion sets in. Land-cover map makers have used Landsat TM and ETM+ data for many years, and more or less automated algorithms for land cover feature extraction are the subject of an extensive literature[10,11,12]. Such techniques generally require some parameter setting for any given scene, and we interested in exploring how a machine may learn to set these parameters or find new algorithms.

We have reported previously on the application of a machine learning technique to the classification of forest fire burn severity using Landsat 7 ETM+ multispectral imagery[13]. For the present work, we are interested in classifying and mapping the post-fire burn scar and the background land-cover classes that form its context. We report results of combining individual feature classification results using a K-means clustering approach (e.g., Ref. 11). We describe the behavior of our evolved classifiers for burn scar and for one particular land-cover case, a generic forest finder. Our multi-feature land-cover map is compared qualitatively to a manually produced land-cover map based on field data collected before the wildfire.

2. MACHINE LEARNING: GENIE

GENIE[14,15,16] is an evolutionary computation (EC) software system which uses a genetic algorithm[17,18,19] (GA) to assemble image-processing algorithms from a collection of low-level ("primitive") image processing operators (e.g., edge detectors, texture measures, spectral operations, and various morphological filters). This system has been shown to be effective in looking for complex terrain features, e.g., golf courses[20]. GENIE can sequentially extract multiple features for the same scene to produce terrain classifications[21], which we will describe in greater detail, below. The implementation details of the GENIE software have been described at length elsewhere[14,16], so we will only present a brief description of the system here.

GENIE follows the classic evolutionary paradigm: a population of candidate image-processing algorithms is randomly generated, and the fitness of each individual assessed from its performance in its environment, which for our case is a user-provided training scene. After fitness has been assigned, reproduction with modification of the most fit members of the population follows via the evolutionary operators of selection, crossover, and mutation. The process of fitness evaluation and reproduction with modification is iterated until some stopping condition is satisfied (e.g., a candidate solution with sufficiently high score is found).

The algorithms assembled by GENIE will generally combine spatial and spectral processing, and the system was in fact designed to enable experimentation with spatio-spectral image processing of multi-spectral and hyper-spectral imagery. Each candidate algorithm in the population consists of a fixed-length string of primitive image processing operations. We now briefly describe our method of providing training data, our encoding of image-processing algorithms as "chromosomes" for manipulation by the GA, and our method for evaluating the fitness of individuals in the population.

2.1. Training Data

The environment for the population consists of one or a number of training scenes. Each training scene contains a raw multi-spectral image data cube, together with a weight plane and a truth plane. The weight plane identifies the pixels to be used in training, and the truth plane locates the features of interest in the training data. Providing sufficient quantities of good training data is a crucial to the success of any machine learning technique. In principle, the weight and truth planes may be derived from an actual ground campaign (i.e., collected on the ground at the time the image was taken), may be the result of applying some existing algorithm, and/or may be marked-up by hand using the best judgement of an analyst looking at the data. We have developed a graphical user interface (GUI), called ALADDIN, for the manual mark-up of raw imagery. Using ALADDIN, the analyst can view a multi-spectral image in a variety of ways, and can mark up training data by painting directly on the image using the mouse. Training data is ternary-valued, with the possible values being "true", "false", and "unknown". True defines areas where the analyst is confident that the feature of interest does exist. False defines areas where the analyst is confident that the feature of interest does not exist. Unknown pixels do not influence the fitness of a candidate algorithm.

2.2. Representation of Image-Processing Algorithms

Traditional genetic programming[22] (GP) uses a variable sized tree representation for algorithms. Our representation differs in that the total number of image-processing operator nodes is fixed (although not all of these need be used in the final graph), crossover is carried out directly on the linear representation, and we allow reuse of values computed by sub-trees through a set of temporary memory "scratch planes".

We have restricted our "gene pool" to a set of useful primitive image processing operators ("genes"). These include spectral, spatial, logical and thresholding operators. The set of morphological operators is restricted to function-set processing morphological operators, i.e., gray-scale morphological operators with a flat structuring element. The sizes and shapes of the structuring elements used by these operators are also restricted to a pre-defined set of primitive shapes, which includes the square, circle, diamond, horizontal cross and diagonal cross, and horizontal, diagonal, and vertical lines. The shape and size of the structuring element are defined by operator parameters. Other local neighborhood/windowing operators such as mean, median, etc., specify their kernels/windows in a similar way. The spectral operators have been chosen to permit weighted sums, differences and ratios of data and/or "scratch" planes, where a scratch plane is a block of memory for storing intermediate calculations within a candidate image-processing algorithm.

A single gene consists of an operator, plus a variable number of input arguments specifying from where input is read, output arguments specifying where output is to be written, and any additional parameters that might be required to specify how the specific operator works (e.g., the diameter and shape of a structuring element used in a morphological filter). The operators used in GENIE take one or more distinct image planes as input, and generally produce a single image plane as output. Input can be taken from any data plane in the training data image cube. Output is written to one of a number of scratch planes, temporary workspaces where an image plane can be stored. Genes can also take input from scratch planes, but only if that scratch plane has been written to by another gene positioned earlier in the chromosome sequence. We use a notation for genes[13] that is most easily illustrated by an example: the gene [ADDP rD0 rS1 wS2] applies pixel-by-pixel addition to two input planes, read from data plane 0 and from scratch plane 1, and writes its output to scratch plane 2. Any additional required operator parameters are listed after the output arguments.

Note that although all chromosomes have the same fixed number of genes, the effective length of the resulting algorithm may be smaller than this. For instance, an operator may write to a scratch plane that is then overwritten by another gene before anything reads from it. GENIE performs an analysis of chromosome graphs when they are created and only carries out those processing steps that actually affect the final result. Therefore, the fixed length of the chromosome acts as a maximum effective length.

2.3. Supervised Classification and Fitness Evaluation

Each candidate image-processing algorithm generates a number of intermediate feature planes (or "signature" planes), which are then combined to generate a Boolean-valued mask for the feature of interest. This combination is achieved using a standard supervised classifier (we use the Fisher linear discriminant[23]), and an optimal threshold function.

Complete (or "hard") classification requires that the image-processing algorithm produce a binary-valued output plane for any given scene. It is possible to treat, e.g., the contents of the first scratch plane as the final output for that candidate image-processing algorithm (thresholding would generally be required to obtain a binary result, though GENIE can choose to apply its own Boolean thresholding functions). However, we have found it useful to perform the combination of the data and scratch planes using a non-evolutionary method, and have implemented a supervised classifier backend. To do this, we first select a subset of the scratch planes and data planes to be "signature" planes. For the present experiments, this subset consists of just the scratch planes. We then use the provided training data and the contents of the signature planes to derive the Fisher Discriminant, which is the linear combination of the signature planes that maximizes the mean separation in spectral terms between those pixels marked up as "true "and those pixels marked up as "false", normalized by the total variance in the projection defined by the linear combination. The output of the discriminant-finding phase is a real-valued single-plane "answer" image. This is reduced to a binary image by exhaustive search over all the training pixels to find the threshold value that minimizes the total number of misclassifications (false positives plus false negatives) on the training data.

The fitness of a candidate solution is given by the degree of agreement between the final binary output plane and the training data. This degree of agreement is determined by the Hamming distance between the final binary output of the algorithm and the training data, with only pixels marked as true or false (as recorded in the weight plane) contributing towards the metric. The Hamming distance is then normalized so that a perfect score is 1000.

3. RESULTS

3.1. Training Data

The remotely-sensed images used in this paper are Landsat 7 ETM+ 30 meter multi-spectral data (ETM+ bands 1–5 and 7). These scenes are Level 1G radiance corrected and georeferenced standard data products obtained via the U.S. Geological Survey (USGS) EarthExplorer[24] web site. We used a post-fire Landsat scene from July 17, 2000, Path 34 and Row 35 for our burn scar mapping, and a pre-fire Landsat scene from July 1, 1999, Path 34 and Row 35, for training our land-cover classifiers at approximately the same stage of senescence as the post-fire imageThese multi-temporal images were registered to each other using the standard co-registration tools of RSI's ENVI[25] remote sensing software package. The image displayed in Fig. 1 is a false-color image, which has then been converted to gray-scale and has had its contrast enhanced for the printing process. As we are interested in mapping vegetation and burn scars, we generally view the data using a Visible/Infrared/Thermal pattern of a thermal IR band (ETM + band 7, 2.2μm) for the red component, a near IR band for the green component (band 5, 1.65μm), and a visible red band for the blue component (band 3, 0.66μm). A Landsat 7 Path/Row swath has an across-track field-of-view of approximately 185 km, with similar along-track length, resulting in a field-of-view of approximately 34,000 sq.km, which is much larger than needed for this study, and presents memory problems for our software if we attempt to ingest the whole scene. Hence, we spatially crop the image to a 1000 x 1000 pixel region centered on the Los Alamos National Laboratory for training purposes, and apply our algorithms to an 80km wide (East–West) by 60 km long (North–South) region that encompasses the area of interest (Jemez Mountains and Pajarito Plateau). We chose not to use the 60m thermal or 15m panchromatic data to avoid the added complication of re-sampling of data.

We did not have any atmospheric measurements available for the scene, so we did not attempt to carry out any corrections for haze or atmosphere. The topography of Los Alamos is complex, consisting of a dormant volcano (the Jemez Mountains) rising to approximately 10,000 feet (3.3km), surrounded by a radiating network of mesas at 7,000 – 8,000 feet (Pajarito Plateau), falling off to the Rio Grande river valley at approximately 6,500 feet elevation. Traditionally, illumination effects due to complex topography can be approximately "factored out" by using band ratios, or removed using principal components analysis (see, e.g., Ref. 10,11). Here, we are interested in the GENIE software's ability to derive results based on the raw imagery, and do not add any additional band ratio or band difference planes.

Our training data was based on two existing map products. The burn scar training data was derived from the official Cerro Grande Burned-Area Emergency Rehabilitation (BAER) Team's burn severity map, Fig. 2, which was produced by trained observers flying over the fire, and visual inspection of high-resolution (~1 meter) aerial color/infrared photography collected during and immediately after the fire. Using this map as a guide, we marked up several regions of the Landsat image as "burn", and several regions as "non-burn" (Fig. 3). The BAER Team assign "burn severity" on the basis of tree mortality –

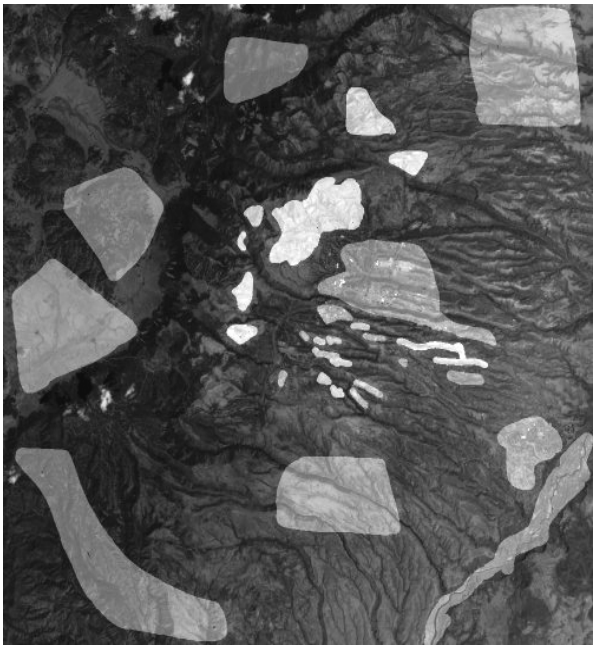

Figure 3. Training Data over raw imagery: White patches mark "burn" regions. Gray patches mark "non-burn" regions. Note: this image is presented at a larger spatial scale than Figure 2.

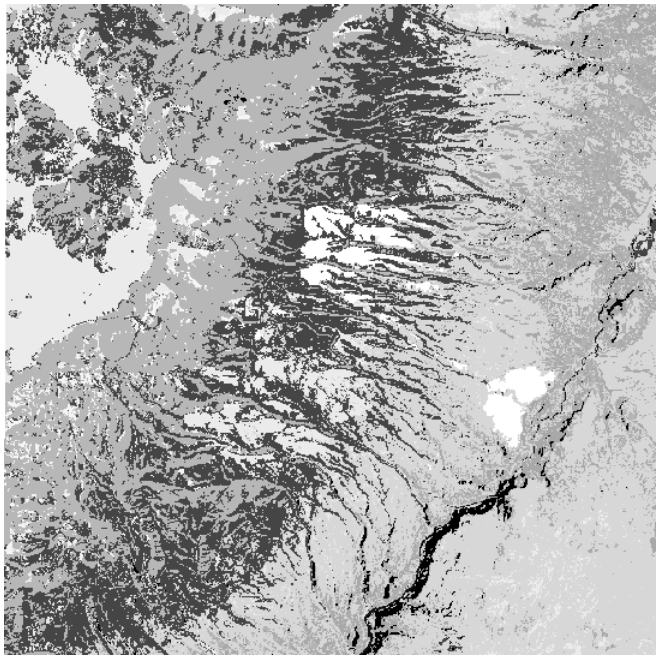

Figure 4. Land cover map for Los Alamos region: based on unsupervised Isodata clustering of a cloud-free Landsat 5 TM image, with clusters merged manually based on field data. (S. W. Koch, private communication.)

low burn severity corresponds to grass fire and low tree mortality, medium severity burn classification implies crown fire and majority tree mortality (more than half of the trees in the marked region are dead), and the high severity burn classification requires that 70 – 100% of the trees are dead. The Cerro Grande wildfire tended to produce either high severity or low severity burn, with only a relatively small fraction of the burn classified as medium burn severity in the BAER Team maps. This was mostly due to the over-grown nature of the Ponderosa pine/mixed conifer forest which suffered most of the damage.

Vegetative land-cover in the Jemez Mountains ecosystem can be usefully described by approximately a dozen land-cover categories. Forest in the region is a mixture of Ponderosa Pine, conifer trees (including Douglas fir and White fir), Aspen, and deciduous species including Gambel's Oak. Open and shrubland categories include alpine meadows (including the grasslands of the Jemez caldera), Piñon–Juniper shrubland, sparsely vegetated mesa tops, and bare outcrops of volcanic tuff and basalt. A pre-fire land-cover map of Los Alamos (Fig. 4) used a semi-automated technique[26] to produce a map with 12 land-cover classes. The Isodata algorithm (see, e.g., Ref. 11) was used to cluster a 30km x 30km subset of a Landsat 5 TM scene (30 meter resolution, 6 channel multi-spectral imagery, collected in the Fall of 1996). The complexity of the scene was reduced by cropping out the area immediately surrounding the town of Los Alamos and the Los Alamos National Laboratory, using additional ground survey data. Fifty centers were used for this unsupervised clustering, and the resulting clusters were reduced by hand to 12 final land-cover classes using field data from approximately 60 ground points. We registered this map to our new Landsat 7 scenes and used individual classes (e.g., Ponderosa pine) and the union of sets of related classes (e.g., all the forest classes: Ponderosa pine, mixed conifer, and aspen) to provide training data for GENIE. In fact, we were unable to directly use the individual land-cover classes of this map for the interesting reason that between the time of the Landsat 5 data collection and the present, more than one forest fire had significantly changed parts of the local environment in the region of interest (Dome wildfire of 1997). Instead, we used this existing land-cover map as a guide to the hand mark-up of training data.

3.2. Example Evolved Image-Processing Algorithms: Burn Scar and Forest Finder

The system was run with a population of 50 chromosomes, each having a fixed length of 20 genes, and 3 intermediate feature ("scratch") planes. The GA was allowed to evolve for 30 generations, in this case, evaluating 1282 distinct candidate image

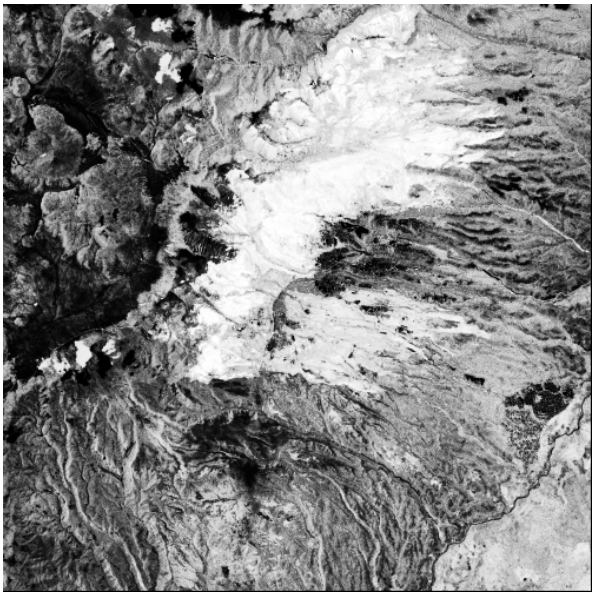

Figure 5. Real-valued Answer Plane: We use a Fisher Discriminant to find the optimal linear combination of evolved "signature" planes into a real-valued answer plane. Regions which will tend to be classified as "burn" are bright. This image has been histogram-equalized to increase contrast.

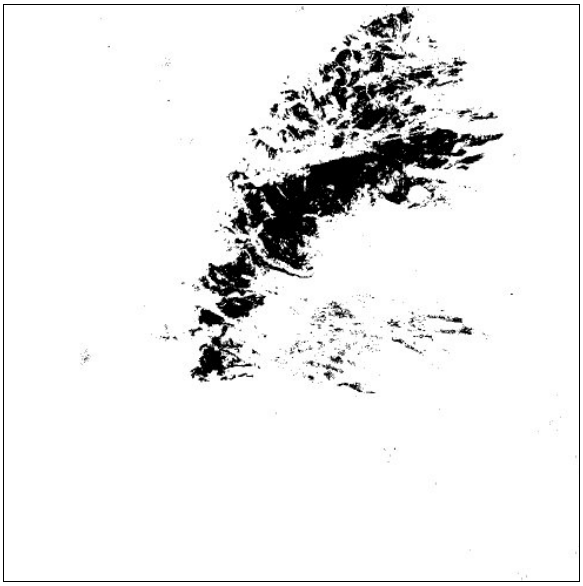

Figure 6. Final burn mask: Thresholding the answer plane produces this burn mask. There is substantial agreement with the details of the BAER map (Fig.2).

processing algorithms, which is very small compared to the search space of possible algorithms given our representation. This required approximately 7 hours of wall-clock time running on a 500MHz Linux/Intel Pentium 2 workstation.

The best evolved image-processing algorithm had the chromosome,

[OPEN rD1 wS1 1 1][ADDS rD4 wS3 0.34][NEG rS1 wS1][MULTP rD4 rS3 wS2]
[LINCOMB rS1 rD6 wS3 0.11][ADDP rS1 rS3 wS1][SUBP rS1 rD5 wS1]

In words, the image-processing algorithm works as follows. Note that GENIE converts the byte-valued raw data to real-valued data (64 bit doubles) and keeps that precision through all its calculations.

1. Data plane D1 (ETM+ band 1, visible blue 0.48μm) undergoes a grayscale morphological opening operation (node 1. OPEN) using a "circular" structuring element with diameter equal to 3 pixels (equivalent to a 3x3 square with corners removed) and the result is written to scratch plane S1,
2. The negative of this plane is taken (node 3. NEG), i.e., S1 $\rightarrow$ – S1,
3. The new S1 is linearly combined (node 5. LINCOMB) with data plane D6 (ETM+ band 7, medium wavelength infrared (MWIR) 2.22μm) with linear weights: 0.11*S1 + 0.89*D6 and the result written to scratch plane S3 (its final value),
4. Scratch planes S1 and S3 are summed (node 6. ADDP), and the difference (node 7. SUBP) of this sum and data plane D5 (ETM+ band 5, MWIR 1.65μm), S1 + S3 – D5, is written to S1 (its final value),
5. Data plane D4 (ETM+ band 4, near infrared 0.83μm) has a constant, 0.34 times a DATASCALE variable equal to the range of the input raw data values, added to each pixel (node 2. ADDS) and is multiplied by D4 again to form the linear combination D4*D4 + (0.34*DATASCALE)*D4, which is written to scratch plane S2 (its final value).

The final values of S1, S2, and S3 are then combined in the linear sum, where the coefficients and intercept have been chosen by the Fisher discriminant, as described in Section 2.3, above, to produce our real-valued answer plane A (Figure 5):

$$A = 0.0147*S1 - 0.0142*S2 + 0.0134*S3 + 1.554$$

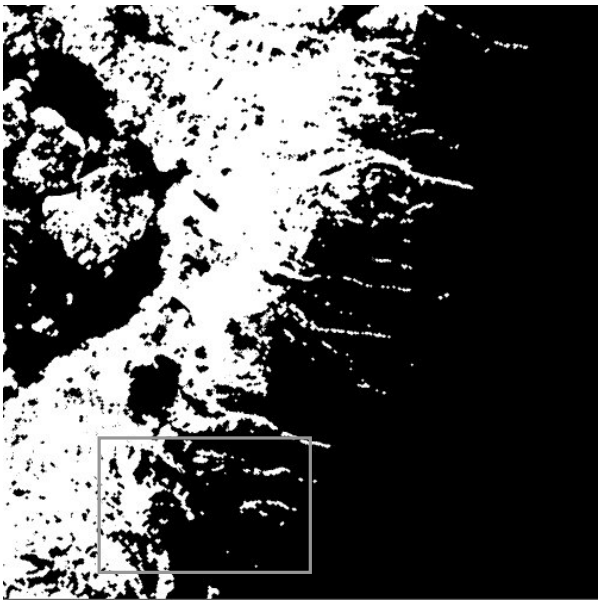

Figure 7. Forest result: White regions are forest, dark regions are non-forest. The region within the gray rectangle reveals loss of forest due to the Dome wildfire (1997), a separate wildfire incident.

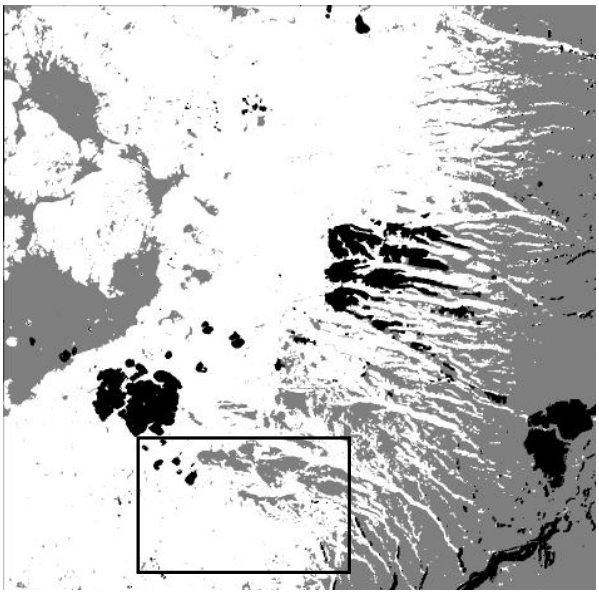

Figure 8. Forest test data: Combined mixed conifer, ponderosa pine, and aspen classes from the Los Alamos land-cover map (Fig. 4). White regions are forest, gray regions are non-forest, and black regions are unclassified. The black rectangle shows pre-Dome wildfire forest cover (compare to Fig. 7).

Converting A to a Boolean mask at a threshold value of 0.7930 produces Figure 6. In relation to the BAER map, Figure 2, we see that the system has extracted the high severity burn region, and the spatial details of this classification correspond very closely to the high severity burn regions in the BAER map. To check the reasonableness of our algorithm's performance, we ran the image-processing algorithm over a larger fraction of the Landsat scene, encompassing the entire Jemez mountain range. This resulted in the detection of a severe burn site on the Western side of the Jemez mountains, which could not be explained by cloud shadows or data drop-out. In fact, this turns out to be a true detection of a second wildfire, the Stable wildfire (affecting Stable Stream and School House Mesa in the Jemez Mountains of northern New Mexico), which destroyed approximately 800 acres of forest in September/October of 1999.

We next present results for evolving a forest finder algorithm. The system was run with manually marked-up training data based on inspection of the existing pre-fire land-cover map (described in Section 3.1 above). GENIE was run with a population of 50 chromosomes, each having a fixed length of 20 genes, and 3 intermediate feature ("scratch") planes. The GA evolved for 50 generations, and produced an algorithm that generated the result shown in Figure 7. We can compare this to the existing land-cover map's combined classes of Ponderosa Pine, Mixed Conifer, and Aspen shown in Fig. 8. We see that the classifications show good qualitative agreement, except in the region of the burn scar left by the Dome wildfire of 1997. On further inspection, the grayscale answer plane produced by GENIE, Fig. 9 (which is automatically thresholded by GENIE to produce Fig. 7), contains much structure, and suggests a possible "forest abundance" interpretation for the grayscale values. The quantitative validity of this possible interpretation will be the study of future work. We stress that there is no reason to expect this behavior, as the only training data provided to GENIE was Boolean.

In this way, we were able to use GENIE to rapidly evolve a set of a dozen feature extraction algorithms for the pre-fire scene. Each classifier appeared to generalize well when presented with progressively larger regions of the Landsat 7 Path/Row tile. Combination of these individual results allows construction of a traditional land-cover map, and in the next section we describe preliminary results of our investigation of a method to do this.

3.3 Post-processing of GENIE results: Unsupervised clustering of feature planes

GENIE is designed to extract one feature at a time, and post-processing is required to combine multiple binary classifications for different features into a single map of terrain classification. On inspection, the set of feature classifiers overlap spatially (i.e., at least two classifiers disagreed) on approximately 20% of the pixels, so a scheme for resolving conflicts between

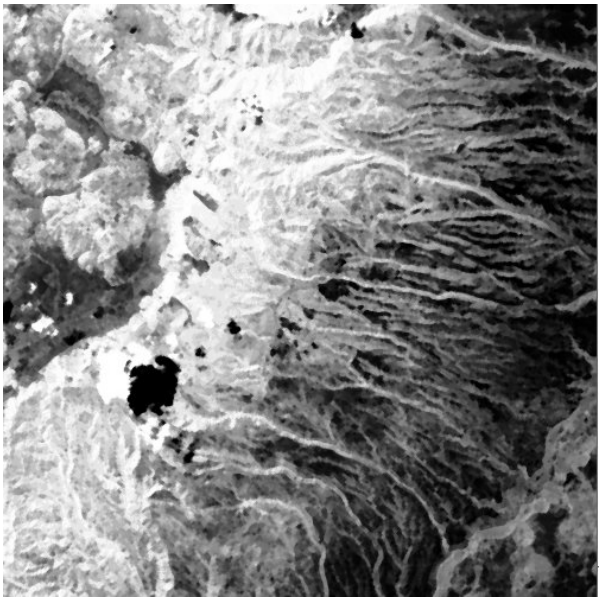

Figure 9. Grayscale Forest Result: Regions are bright as some function of their "forest" content. This grayscale result is automatically thresholded by GENIE to produce the Boolean classification shown in Fig. 7. A histogram equalization has been applied to enhance spatial details.

classifiers is required. Previous work with GENIE using multiple features[21] considered a "three-color trick" for visualizing the GENIE results: three Boolean feature classifications can be combined by mapping the features to the RGB channels of a color display, which gives 3 pure and 5 mixed color classes. For the case of a dozen features for which there is no restriction on more than one classifier claiming any given pixel, many ways of combining results exist.

One approach would be to examine the grayscale answer plane produced by each feature classifier, and set the label of each pixel according to which classifier had the strongest response (an "abundance" interpretation for each classifier's grayscale output). Since the original Landsat TM-based land-cover map used Isodata unsupervised clustering and supervised re-merging of the raw TM data, we decided to explore k-means clustering[27] of several GENIE feature planes. We now report preliminary, qualitative results of this approach, and will present quantitative results elsewhere[28].

As an initial test, we used three evolved feature planes: forest (described in detail in the previous section), grasslands/shrublands, and bare ground. The concept is that these feature planes represent useful axes in land-cover feature space. As "physical" features of the surface, k-means clustering on these planes will explore "natural" segmentation of the scene features. We ran k-means with 12 clusters for 200 iterations, which was sufficient for the algorithm to completely converge (i.e., class assignments are stable). We also evolved a cloud finder and cloud shadow finder, and overlaid these on the classified image. The final result is shown in Fig. 10. As always with k-means, cluster labels given by the algorithm are arbitrary, and we are free to complete the identification of centers with land-cover classes. We see that the result is strongly qualitatively consistent with the existing land-cover map in the neighborhood of Los Alamos, and that the features classes appear to coincide with physical land-cover classes throughout the Jemez Mountains region.

4. CONCLUSIONS

We have investigated evolution of image-processing algorithms to extract wildfire burn scars and land-cover classes in Landsat 7 ETM+ imagery from two time periods, pre- and post-fire, and have described the operation of some evolved algorithms in detail. The evolved algorithm shows a good qualitative fit to the published BAER Team burn-severity map of the May 2000 Cerro Grande/Los Alamos wildfire, specifically in comparison to their high-severity burn class (70-100% tree mortality regions). K-means clustering of GENIE feature planes shows promise for production of multi-feature land-cover maps. We find these results quite encouraging for the future application of this machine learning technique.

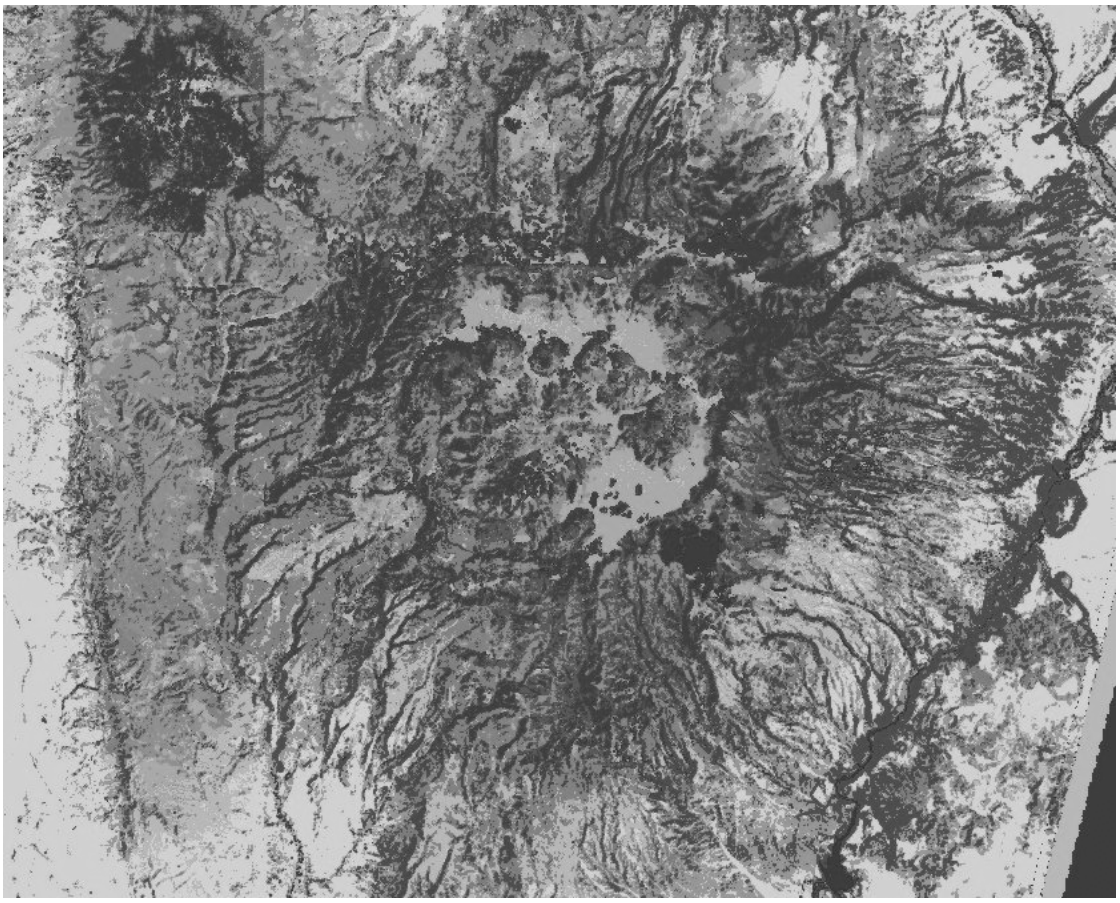

Figure 10. GENIE Land-cover classification: K-means clustering of the GENIE forest, grassland/ shrubland, and bare ground feature planes produces a land cover map that shows excellent agreement with the pre-fore land-cover map of Los Alamos (Fig. 4) but now extends to a much larger region (80km x 60km). 10 Categories are present: Mixed conifer forest, Ponderosa Pine forest, Aspen/Deciduous, Grassland, three densities of Piñon-Juniper shrubland, Unvegetated, Cloud/shadow, and unclassified. The cloud/shadow class is the result of GENIE feature extraction algorithm, and is overlaid on the clustered image.

ACKNOWLEDGEMENTS

The authors would like to thank Leslie Hanson, Steven Koch, and Randy Balice of the Ecology Group for useful discussions and access to data used in this work. One of the authors, S.P.B. wishes to thank C. S. Plesko for useful discussions. The GENIE system is the result of the combined efforts of a team of people at LANL, including, in addition to the authors of this paper: Mark Galassi, Kevin Lacker, Melanie Mitchell, Catherine Plesko, and Reid Porter.

REFERENCES

1. The National Oceanographic and Atmospheric Administration (NOAA) POES satellites and the AVHRR instrument are described on the NOAA web site http://www.ncdc.noaa.gov
2. Landsat TM and ETM+ are described on the U.S. Geological Survey (USGS) web site http://landsat7.usgs.gov
3. Y.J. Kaufmann, C.J. Tucker, and I. Fung, “Remote sensing of biomass burning in the tropics”, J. Geophysical Research, Vol 95, No. D7, pp. 9927-9939, 1990, and references therein.
4. Y. Rauste, E. Herland, H. Frelander, K. Soini, T. Kuoremäki, and A. Ruokari, “Satellite-based forest fire detection for fire control in boreal forests”, Int. J. Remote Sensing, Vol. 18, No. 12, pp. 2641-2656, 1997.

5. N.P. Minko, N.A. Abushenko, V.V. Koshelev, “Forest fire detection in East Siberia forests using AVHRR/NOAA data”, Proc. SPIE, Vol. 3502, pp. 192-200, 1998.
6. R. Lasaponara, V. Cuomo, V. Tramutoli, N. Pergola, C. Pietrapertosa, and T. Simoniello, “Forest fire danger estimation based on the integration of satellite AVHRR data and topographic factors”, Proc. SPIE, Vol. 3868, pp. 241-252, 1999.
7. S.H. Boles and D.L. Verbyla, “Effect of scan angle on AVHRR fire detection accuracy in interior Alaska”, Int. J. Remote Sensing, Vol. 20, No. 17, 3437-3443, 1999.
8. A. Lobo, N. Pineda, R. Navarro-Cedillo, P. Fernandez-Rebollo, F.J. Salas, J.-L. Fernández-Turiel, and A. Fernández-Palacios, “Mapping forest fire impact from Landsat TM imagery”, Proc. SPIE, Vol. 3499, 340-347, 1998.
9. J.D. Kushla and W.J. Ripple, “Assessing wildfire effects with Landsat thematic mapper data”, Int. J. Remote Sensing, Vol. 19, No. 13, 2493-2507.
10. R.A. Schowengerdt, *Remote Sensin*g, 2nd ed., Academic, San Diego (1997).
11. J. A. Richards and X. Jia, *Remote Sensing Digital Image Analysis*, 3rd ed., Springer, Berlin, 1999.
12. R. S. Lunetta and C. D. Elvidge (editors), *Remote sensing change detection*, Ann Arbor, Chelsea (1998).
13. S. P. Brumby, et al., “Evolving forest fire burn severity classification algorithms for multi-spectral imagery”, to appear in Proc. SPIE, Vol. 4381. Presented at Aerosense 2001: SPIE's 15th Annual International Symposium on Aerospace/ Defense Sensing and Controls, Orlando, Florida, April 16-20, 2001.
14. S.P. Brumby, J. Theiler, S.J. Perkins, N.R. Harvey, J.J. Szymanski, J.J. Bloch, and M. Mitchell, “Investigation of feature extraction by a genetic algorithm ”, Proc. SPIE, Vol. 3812, pp. 24-31,1999.
15. J. Theiler, N.R. Harvey, S.P. Brumby, J.J. Szymanski, S. Alferink, S.J. Perkins, R. Porter, and J.J. Bloch, “Evolving retrieval algorithms with a genetic programming scheme ”, Proc. SPIE, Vol. 3753, pp.416-425,1999.
16. N. R. Harvey, J. Theiler, S. P. Brumby, S. Perkins, J. J. Szymanski, J. J. Bloch, R. B. Porter, M. Galassi, and A. C. Young,, "Image Feature Extraction: GENIE vs Conventional Supervised Classification Techniques", submitted to IEEE Transactions on Geoscience and Remote Sensing, 2001.
17. J. H. Holland, *Adaptation in Natural and Artificial Systems*, University of Michigan, Ann Arbor (1975).
18. I. Rechenberg, *Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution, Fromman-Holzboog*, Stuttgart (1973).
19. L. Fogel, A. Owens and M. Walsh, *Artificial Intelligence through Simulated Evolution*, Wiley, New York (1966).
20. N.R. Harvey, S. Perkins, S.P. Brumby, J. Theiler, R.B. Porter, A.C. Young, A.K. Varghese, J.J. Szymanski, and J.J. Bloch, “Finding golf courses: The ultra high tech approach”, Proc. Second European Workshop on Evolutionary Computation in Image Analysis and Signal Processing (EvoIASP2000), Edinburgh, UK, pp 54-64, 2000.
21. S.P. Brumby, et al., “A genetic algorithm for combining new and existing image processing tools for multispectral imagery”, Proc. SPIE, Vol. 4049, pp. 480-490, 2000.
22. J. R. Koza, *Genetic Programming: On the Programming of Computers by Natural Selection*, MIT, Cambridge (1992).
23. For example, C.M.Bishop, *Neural Networks for Pattern Recognition*, pp.105 –112, Oxford University (1995).
24. For details on USGS EarthExplorer, see `http://edcsns17.cr.usgs.gov/EarthExplorer`.
25. See `http://www.rsinc.com/envi`.
26. S. W. Koch (Ecology Group, Los Alamos National Laboratory), private communication.
27. For example, R.O. Duda, P.E. Hart, and D.G. Stork, Pattern Classification (2nd ed.), Chapter 10, Wiley-Interscience, New York, 2001.
28. S. P. Brumby, et al., in preparation, 2001.