Genetic programming approach to extracting features from remotely sensed imagery

Steven P. Brumby, James Theiler, Simon Perkins,
Neal R. Harvey, and John J. Szymanski
Space and Remote Sensing Sciences,
Los Alamos National Laboratory, Mail Stop D436,
Los Alamos, New Mexico 87545, USA.
brumby@lanl.gov

Abstract - *Multi-instrument data sets present an interesting challenge to feature extraction algorithm developers. Beyond the immediate problems of spatial co-registration, the remote sensing scientist must explore a complex algorithm space in which both spatial and spectral signatures may be required to identify a feature of interest. We describe a genetic programming/supervised classifier software system, called Genie, which evolves and combines spatio-spectral image processing tools for remotely sensed imagery. We describe our representation of candidate image processing pipelines, and discuss our set of primitive image operators. Our primary application has been in the field of geospatial feature extraction, including wildfire scars and general land-cover classes, using publicly available multi-spectral imagery (MSI) and hyper-spectral imagery (HSI). Here, we demonstrate our system on Landsat 7 Enhanced Thematic Mapper (ETM+) MSI. We exhibit an evolved pipeline, and discuss its operation and performance.*

Keywords: Evolutionary Computation, Genetic Programming, Image Processing, Remote Sensing, Multi-spectral Imagery, Panchromatic imagery

1 Genetic programming with supervised classification

GENIE [1-4] is an evolutionary computation (EC) software system, using a genetic algorithm (GA) [5-7] 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, such as, e.g., golf courses [8]. GENIE can sequentially extract multiple features for the same scene to produce terrain classifications [9]. GENIE has been described at length elsewhere ([1-4]), 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 follows via the evolutionary operators of selection, crossover, and mutation, applied to the most fit members of the population. The process of fitness evaluation and reproduction with modification is iterated until some stopping condition is satisfied.

The algorithms assembled by GENIE will generally combine spatial and spectral processing, and the system was in fact designed to enable spatio-spectral image processing experimentation. Each individual chromosome in the population consists of a fixed-length string of genes. Each gene in Genie corresponds to a primitive image processing operation, and so the whole chromosome describes an algorithm consisting of a sequence of primitive image processing steps. 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.

1.1 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 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 manual marking 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 a 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.

1.2 Representation of Image-Processing Algorithms

Traditional genetic programming (GP) [10] uses a variable sized (within limits) tree representation for algorithms. Our representation differs in that it allows for reuse of values computed by sub-trees, i.e. the resulting algorithm is a graph rather than a tree. The image processing algorithm that a given chromosome represents can be thought of as a directed acyclic graph where the non-terminal nodes are primitive image processing operations, and the terminal nodes are individual image planes extracted from the multi-spectral image used as input. In our representation, the total number of nodes is fixed, although not all of these may actually be used in the final graph, and crossover is carried out directly on the linear representation.

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 having a flat structuring element. The sizes and shapes of the structuring elements used by these operators is 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 number of input arguments specifying from which planes input is read, output arguments specifying to which planes output is written, and any additional parameters that 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 or scratch 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 [1] 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.

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.

1.3 Supervised Classification and Fitness Evaluation

Each candidate image-processing algorithm generates a number of intermediate feature planes (or "scratch" 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 [11]), 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 misclassified pixels (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.

1.4 Example: Remote sensing of forest fires

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). Initial estimates of forest damage included 17,000 acres (6,900 ha) suffering 70-100% tree mortality (Fig. 1). 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. During and after the fire, remote-sensing data was acquired from a variety of aircraft-based and satellite-based sensors, including Landsat 7, to evaluate the impact of the fire.

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 (AVHRR) [12] on the NOAA Polar-orbiting Operational Environmental Satellite (POES) series, and the Thematic Mapper (TM) and Enhanced Thematic Mapper (ETM+) instruments on the Landsat series of Earth observation satellites [13], 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, using, e.g., thresholds on brightness temperature [14-18] in AVHRR band 3 (3.7µm), and mapping of post-fire burn scars. For the present work, we are interested in mapping and classifying the post-fire burn scar.

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 [19] apply a combination of spectral image segmentation and hierarchical clustering to the mapping and analysis of fires in Mediterranean forests. Kushla and Ripple [20] 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.

We now report on the application of a machine learning technique to the classification of forest fire burn severity using Landsat 7 ETM+ multispectral imagery. The details of this automatic classification are compared to a manually produced burn classification, which was derived from field observations and human photo-interpretation of high-resolution aerial color/near infrared photography.

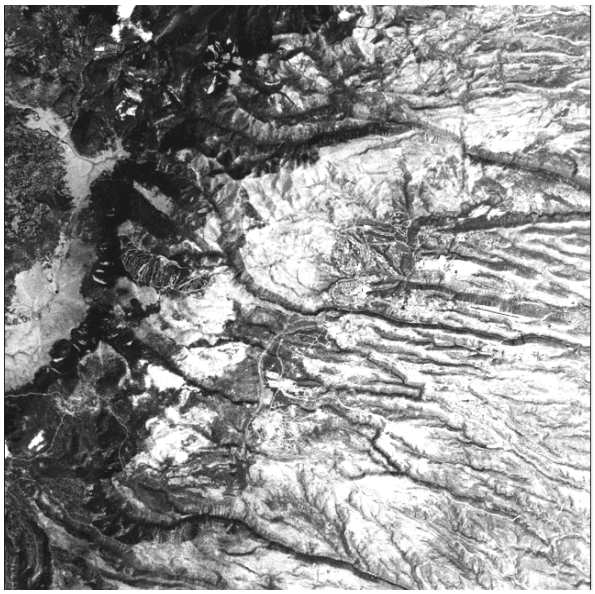

Figure 1. Landsat 7: Post-fire, July 19, 2000: Bright region in center of image is the burn scar (also see Fig.2). Los Alamos town lies against the underside of the burn scar. Topography changes from forested mountains (left) to bare mesas.

2 Training and results

2.1 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 [21] web site. We used a post-fire Landsat scene from July 17, 2000, Path 34 and Row 35. 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 larger than needed for this study. Hence, we spatially subset the image to a 1000 pixel x 1000 pixel region centered on the Los Alamos National Laboratory. We chose not to use the 60m thermal or 15m panchromatic data in the following analysis. In future work, we plan to extend our software to combine data of different spatial resolutions.

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 feet, falling off to the Rio Grande river valley at approximately 6,000 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., [22]). 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 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 almost certainly "burn", and several regions as almost certainly "non-burn", as shown in Fig. 3. The BAER Team assign "burn severity" on the basis of tree mortality – 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. Note that we have also evolved algorithms from training data based purely on photo-interpretation of the 30m Landsat scene, and have obtained similar results. This is most likely due to the fact that in the case of the Cerro Grande wildfire, the burn damage was sufficiently catastrophic that simple inspection of the 30m imagery allows accurate marking of "burn" and "non-burn" regions.

2.2 Evolved Image-Processing Algorithm

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 processing algorithms, which is very small compared to the search space of possible algorithms allowed by our representation. This evolution required approximately 7 hours of wall-clock time running on a 500MHz Linux/Intel Pentium-class workstation. This evolution time is a one-off cost: the final "best" image processing tool evolved by Genie only requires a few 10's of seconds to execute on the training scene, and is similarly fast executing on other Landsat 7 imagery.

Figure 2. BAER Team burn-severity map overlaid on a topographic map: Medium gray region marks high severity burn, pale gray region marks low severity/un-burned region: http://www.baerteam.org/cerrogrande

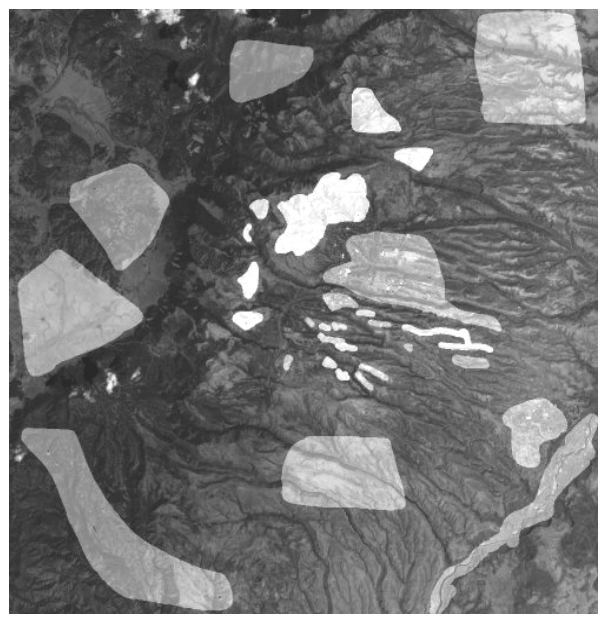

Figure 3. Training Data over Landsat imagery: White patches mark "burn" regions. Gray patches mark "non-burn" regions. Note: this image is presented at a larger spatial scale than Figure 2.

The best evolved image-processing algorithm had the chromosome (using our mnemonic convention to label primitive image processing operators),

[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 (Fig. 4):

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

The optimal threshold found by GENIE, given the training data, was 0.3437. Converting A to a Boolean mask at that threshold value produces Fig. 5. In relation to the BAER map (Fig. 2), we see that the system has extracted the high, medium, and low severity burn regions, but also presents a number of false positives. On inspection, these turn out to correspond to two physical categories of land cover: bare ground/rock, and cloud shadows. The histogram of A shows a bimodal distribution (Fig. 6), as expected if the burn/non-burn classes are separable. Adjusting the threshold on A to fall at the between-peak minimum of the histogram at 0.7930 (a different optimization criterion for the threshold than that used by default by GENIE) produces a new Boolean mask (Fig. 7), in which almost all the false positives have been removed, and the remaining pixels marked as "burn" correspond very closely to the high severity burn regions in the BAER map. For more details, see [23].

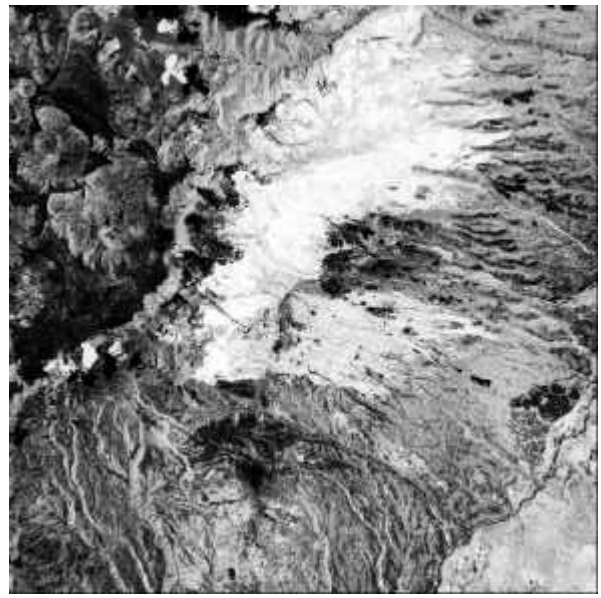

Figure 4. Real-valued Answer Plane: Regions classified as "burn" are bright. This image has been histogram-equalized to increase contrast.

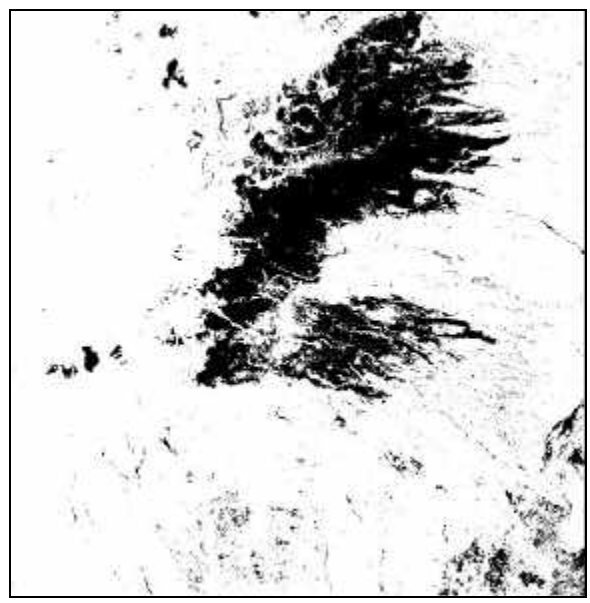

Figure 5. Burn mask: pixels classified as "burn" are shown black. GENIE determines an optimal threshold for converting the real-valued answer plane to a Boolean mask. Misidentified pixels are mostly cloud shadows (e.g., compact regions on left), or bare ground/rock (lower right and bottom).

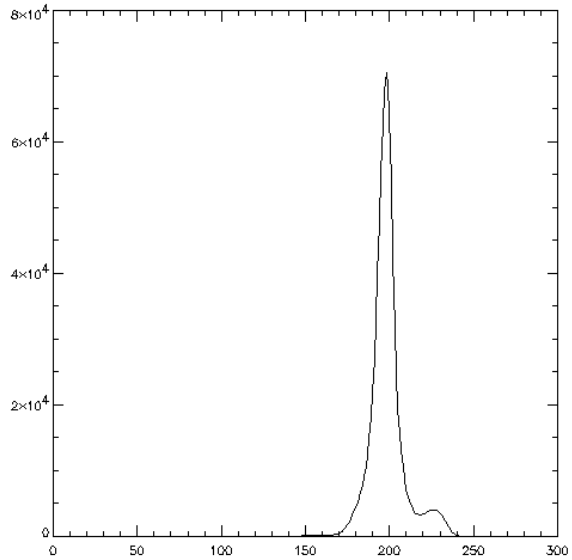

Figure 6. Histogram of the Answer Plane: The bimodal distribution indicates that "burn" and "non-burn" are indeed separable classes.

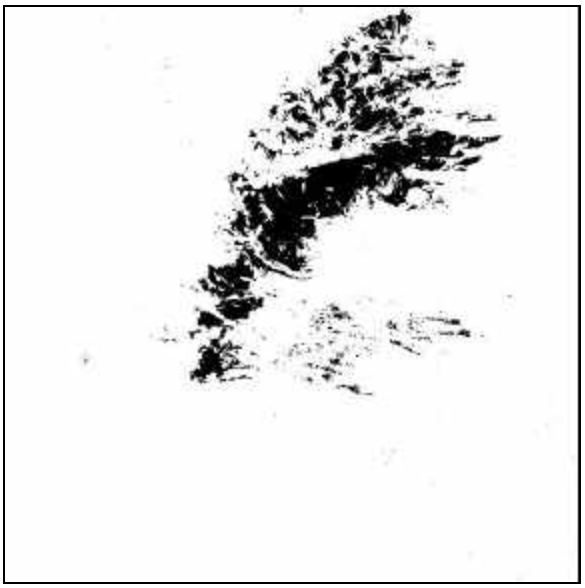

Figure 7. Final burn mask: Thresholding the answer plane at the between-peak minimum of the bimodal distribution produces this burn mask, which has almost eliminated false positives. There is substantial agreement with the details of the BAER map (Fig.2).

2.3 Application to Non-Training Data

The evolved algorithm can now be applied to any scene. 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. The result is shown in Fig. 8. We claim

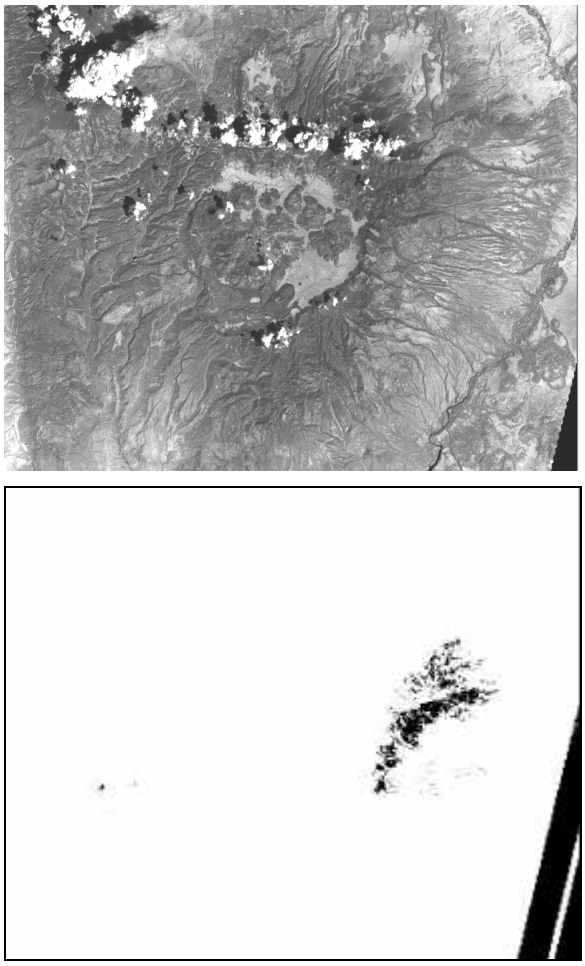

Figure 8 Testing the evolved image-processing algorithm: Extended region (top: image size is 80km by 60km) and burn mask (bottom). The evolved image-processing algorithm continues to work well, except at the edge of the Landsat swath (bottom image, lower right). The small black region on the left of the burn mask represents a true detection of a second recent wildfire, the 1999 Stable wildfire.

that this overall result is quite reasonable, and only fails where the Landsat swath ends (which can be easily masked out). Of particular interest is the persistent detection of a severe burn site on the Western side of the Jemez mountains, Fig. 9, which cannot obviously be excluded due to 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, including 200-400 acres of high-severity crown fire. As GENIE had no knowledge of this fire during its training, we find this detection, together with the reasonable behavior of the evolved image-processing algorithm over this large region, as quite encouraging for the future usefulness of this machine learning technique.

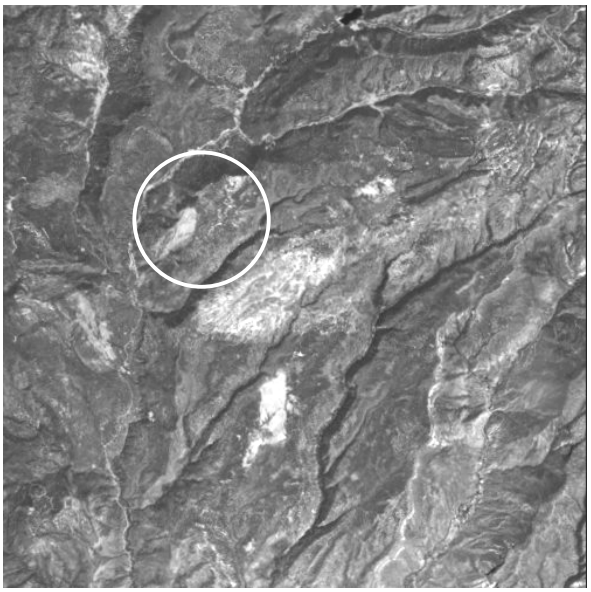

Figure 9. Detail of the second detected burn: Grayscale image, ETM+ band 7. Location, surrounded by the white circle, agrees with the known location of the 1999 Stable stream wildfire.

Figure 10. Burn severity classification, overlaid on a grayscale image of the region neighboring Los Alamos town (top center). A color version of this image is available at http://www.daps.lanl.gov/genie/gallery.html. High severity burn regions dominate the mountains north and west of town (top left, shown in dark gray) and also occur in isolated pockets south of the town. Medium severity (pale gray) and low severity (white) burn regions border and link the high severity regions. The pattern of burn severity agrees well with the BAER map and with field survey data.

2.4 Classification of burn severity

The algorithm described above successfully maps high-severity burn regions, but we are also interested in mapping medium and low severity burn regions. There are several possible approaches to this multiple-category classification problem. The GENIE software is designed to extract one feature at a time, so for this problem we chose to construct three hierarchical training sets, based on the BAER team burn severity map. These consisted of a high severity burn vs. (medium or low severity or no burn) classification problem, a (high or medium severity burn) vs. (low severity or no burn) classification problem, and a (high or medium or low severity burn) vs. no burn classification problem. In this way, we expect the difference regions between classifications to correspond to the individual high, medium and low burn severity regions.

For this work, we used airborne Daedalus 3600 MSI, which provides 6 bands of visible through thermal IR at approximately 25 meter spatial resolution, and which is quite similar to the 30 meter Landsat 7 ETM+ bands. Full details of the algorithms will be presented elsewhere, as they are quite similar to the algorithm described at length above. Similar results have been obtained with Landsat 7 ETM+.

We chose to combine these individual results by overlaying the “true” pixels of each classification , as shown in Fig. 10 for a region neighboring the town of Los Alamos (a color version of this figure is available at http://www.daps.lanl.gov/genie/gallery.html). Compared to the BAER map, we found good qualitative agreement for all three burn severity categories. More high and medium severity burn is indicated south of the town than was indicated on the BAER map. We have compared this result with field survey data collected post-fire, and again find good qualitative agreement between the evolved map and these field reports. We are currently working to quantitatively validate this result.

3 Conclusions

We have investigated evolution of an image-processing algorithm to extract wildfire burn scars in Landsat 7 ETM+ imagery, and have described the operation of the evolved algorithm in some 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). The algorithm can be confused by dark cloud shadows, and by bare ground/rock outcrops which are physically very similar to the charred remains of the severely burned forest, but adjustment of its final threshold can significantly improve this behavior. Applying the algorithm outside the training area showed that it continued to produce reasonable results over a large spatial region, and in fact was able to detect a second small wildfire on the west side of the Jemez mountains (September/October 1999 Stable stream wildfire). We find these results quite encouraging for the future application of this machine learning technique.

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: Jeffrey J. Bloch, Reid B. Porter, Mark Galassi, Kevin Lacker, and Melanie Mitchell. This work was supported by the U.S. Department of Energy and Department of Defense.

References

[1] 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.

[2] 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.

[3] N.R. Harvey, S.P. Brumby, S.J. Perkins, R.B. Porter, J. Theiler, A.C. Young, J.J. Szymanski, and J.J. Bloch, *Parallel evolution of image processing tools for multispectral imagery*, Proc. SPIE, Vol. 4132, pp. 72-82, 2000.

[4] S. Perkins, J. Theiler, S.P. Brumby, N.R. Harvey, R.B. Porter, J.J.Szymanski, and J. J. Bloch, *GENIE: A hybrid genetic algorithm for feature classification in multi-spectral images*, Proc. SPIE, Vol. 4120, pp 52-62, 2000.

[5] J. H. Holland, *Adaptation in Natural and Artificial Systems*, University of Michigan, Ann Arbor, 1975.

[6] I. Rechenberg, *Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution, Fromman-Holzboog*, Stuttgart, 1973.

[7] L. Fogel, A. Owens and M. Walsh, *Artificial Intelligence through Simulated Evolution*, Wiley, New York, 1966.

[8] 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.

[9] S.P. Brumby, N.R. Harvey, S. Perkins, R.B. Porter, J.J. Szymanski, J. Theiler, and J.J. Bloch, *A genetic algorithm for combining new and existing image processing tools for multispectral imagery*, Proc. SPIE, Vol. 4049, pp. 480-490, 2000.

[10] J. R. Koza, *Genetic Programming: On the Programming of Computers by Natural Selection*, MIT, Cambridge, 1992.

[11] For example, see C.M.Bishop, *Neural Networks for Pattern Recognition*, pp.105–112, Oxford University, 1995.

[12] 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

[13] Landsat TM and ETM+ are described on web site http://landsat7.usgs.gov

[14] 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.

[15] 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.

[16] 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.

[17] 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.

[18] 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.

[19] 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.

[20] 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.

[21] For details on USGS EarthExplorer, see http://edcsns17.cr.usgs.gov/EarthExplorer.

[22] R.A. Schowengerdt, *Remote Sensin*g, 2nd ed., Academic, San Diego, 1997.

[23] S.P. Brumby, N.R. Harvey, J.J. Bloch, J. Theiler, S. Perkins, A.C. Young, and J.J. Szymanski, *Evolving forest fire burn severity classification algorithms for multi-spectral imagery*, to appear in Proc. SPIE, 2001.