Automated Simultaneous Multiple Feature Classification of MTI Data

Neal R. Harvey, James Theiler, Lee Balick, Paul Pope, John J. Szymanski,
Simon J. Perkins, Reid B. Porter, Steven P. Brumby, Jeffrey J. Bloch,
Nancy A. David, Mark Galassi

Space and Remote Sensing Sciences Group
Los Alamos National Laboratory
Los Alamos, New Mexico 87545

ABSTRACT

Los Alamos National Laboratory has developed and demonstrated a highly capable system, GENIE, for the two-class problem of detecting a single feature against a background of non-feature. In addition to the two-class case, however, a commonly encountered remote sensing task is the segmentation of multispectral image data into a larger number of distinct feature classes or land cover types. To this end we have extended our existing system to allow the simultaneous classification of multiple features/classes from multispectral data. The technique builds on previous work and its core continues to utilize a hybrid evolutionary-algorithm-based system capable of searching for image processing pipelines optimized for specific image feature extraction tasks.

We describe the improvements made to the GENIE software to allow multiple-feature classification and describe the application of this system to the automatic simultaneous classification of multiple features from MTI image data. We show the application of the multiple-feature classification technique to the problem of classifying lava flows on Mauna Loa volcano, Hawaii, using MTI image data and compare the classification results with standard supervised multiple-feature classification techniques.

Keywords: supervised classification, multiple features, evolutionary computation, lava flow classification

1. INTRODUCTION

Much interest has been shown in recent years in the development of feature extraction tools which can assist in the exploitation of the ever-increasing quantities of multi-spectral data that are becoming available. Creation and development of task-specific feature-detection algorithms is important, yet can be extremely expensive, often requiring a significant investment of time and effort by highly skilled personnel. At Los Alamos National Laboratory we have developed an automated system for the generation of feature extraction/classification tools, which we refer to as GENIE.

Our particular interest is the pixel-by-pixel classification of multi-spectral remotely-sensed images, both to locate and identify and also to delineate particular features of interest. The large number of features in which we are interested, together with the variety of instruments which are available, make the hand-coding of suitable feature-detection algorithms impractical. We therefore employ a supervised learning approach that can generate image processing pipelines capable of distinguishing features of interest. Until recently our approach has been to only consider the two-class problem (distinguishing a single class against a background of "other" classes), however, many applications require the segmentation of an image into a larger number of distinct features or land-cover types. To this end we have extended GENIE's capability to allow the simultaneous classification of multiple features/classes from multispectral data. The technique builds on previous work and its core continues to utilize a hybrid evolutionary-algorithm-based system capable of searching for image processing pipelines optimized for specific image feature extraction tasks.

To demonstrate and evaluate the system we gave it the task of classifying lava flows on Mauna Loa Volcano, Hawaii, from multispectral data obtained from the Multispectral Thermal Imager (MTI) Satellite. In order to have some bench-mark against which to compare GENIE's performance, we gave the same classification tasks to some standard, commonly-used supervised classification techniques.

Work supported by the U.S. Department of Energy.
Emails: {harve,jt,lbalick,papope,szymanski,s.perkins,rporter,brumby,jbloch,ndavid,mgalassi}@lanl.gov.

2. BRIEF OVERVIEW OF THE "GENIE" SYSTEM

The details of GENIE's algorithmic structure have been described previously in the literature,[1–7] so, in the interests of brevity, we provide only a brief overview of our system, and describe the modifications from the previous versions that have enabled the multiple-feature classification capability.

GENIE employs a classic evolutionary paradigm: a population is maintained of candidate solutions (*chromosomes*), each composed of interchangeable parts (*genes*), and each assessed and assigned a scalar fitness value, based on how well it performs the desired task. After fitness determination, the evolutionary operators of selection, crossover and mutation are applied to the population and the entire process of fitness evaluation, selection, crossover and mutation is iterated until some stopping condition is satisfied.

2.1. Environment

The environment for each individual in the population consists of *data* planes, each of these planes corresponding to a separate spectral channel in the original multi-spectral image, together with a *weight* plane and a *feature* plane. The weight plane identifies those pixels to be used in training – these are all the pixels for which the analyst has provided a class label. The actual delineation of separate feature/class pixels is given by the feature plane.

2.2. Chromosomes and Genes

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. Therefore the entire chromosome describes an algorithm consisting of a sequence of primitive image processing operations.

Each gene used in GENIE takes one or more distinct image planes as input, and produces one or more image planes as output. Input can be taken from any data planes in the training data image cube. Output is written to any of a small 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 earlier in the chromosome sequence.

Our "gene pool" is composed of a set of primitive image processing operators which we consider useful. These include spectral, spatial, logical and thresholding operators.

2.3. Backends

Final classification requires that the algorithm produce a single (discrete) scalar output plane, which identifies, for every pixel, the class to which it has been assigned. We have found it advantageous to adopt a hybrid approach which applies a conventional supervised classifier to a (sub)set of scratch and data planes to produce the final output plane.

To do this, we first select a subset of the scratch and data planes to be *answer planes*. The conventional supervised classifier "backend" uses the answer planes as input and produces a final output classification plane; in principle, we can use any supervised classification technique as the backend, but for the experiments reported here, we used either *Maximum Likelihood*, *Minimum Euclidean Distance* or *Minimum Spectral Angle*.[8–11] By implementing the conventional multiclass classifiers as backends in GENIE, we were able to extend GENIE to directly perform multiclass classification. These new backends also extend the options for using GENIE to do binary classification. We also included some experiments using two-class discrimination, and in this case we used the Fisher linear discriminant[12] as the backend.

2.4. Fitness Evaluation

The fitness of a candidate solution is given by the degree of agreement between the final classification output plane and the training data. It is based on a simple ratio of the total number of incorrectly classified training pixels over all classes to the total number of training pixels over all classes. If we denote the total error (total number of training pixels classified incorrectly over all classes) as E, and the total number of training pixels over all classes as N, then the fitness F of a candidate solution is given by

$$F = (1.0 - (E/N)) \times 1000 \tag{1}$$

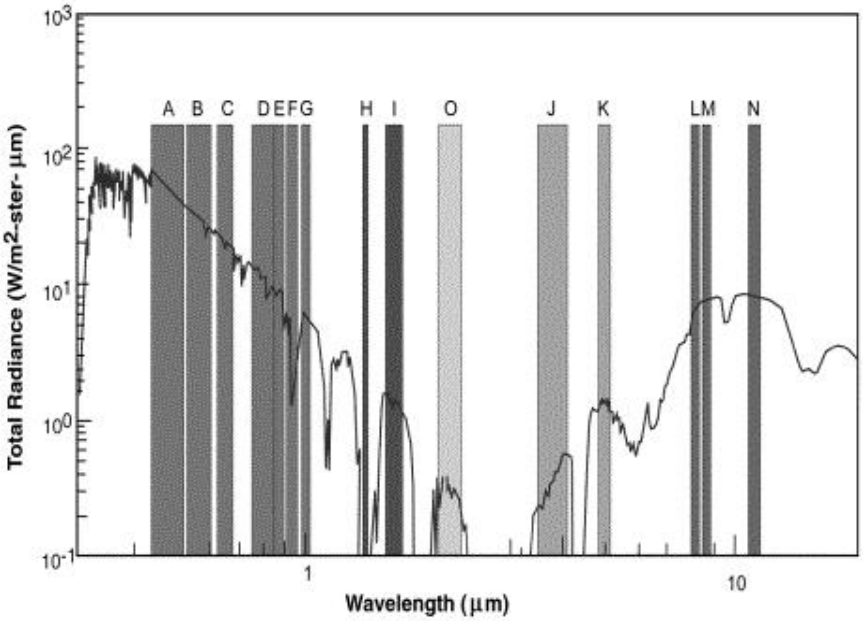

Figure 1. MTI Spectral Bands

Thus, a fitness of 1000 indicates a perfect classification result, i.e. no training pixels in any class have been classified incorrectly, and a fitness of 0 indicates an entirely incorrect classification, i.e. no training pixels in any class have been classified correctly.

3. THE MTI SATELLITE AND ITS DATA

The Multispectral Thermal Imager (MTI)[13] is a space-based research and development project sponsored by the U.S. Department of Energy (DOE). MTI's primary objective is to demonstrate technologies such as advanced multispectral and thermal imaging, image processing, and other associated technologies.

The MTI system consists of a single satellite that was launched into a polar, 360-mile-high orbit in February 2000 and carries an advanced multispectral and thermal imaging sensor. During its 3-year mission, the MTI satellite will periodically record images in 15 spectral bands, ranging from visible to long-wave infrared (three in the visible, five in the near infra-red, two in the short-wave infra-red, and five in the thermal infra-red). MTI's spectral bands are carefully selected to collect data needed to derive a broad range of information, including surface temperatures, materials, water quality, and vegetation health. To enhance accuracy, additional bands provide simultaneous information on atmospheric water vapor, aerosol content, and sub-visual cloud presence. Though the MTI system nominally images fifteen spectral bands (identified by the letters A – 0), the data is actually reported in terms of sixteen effective bands, because one spectral band, H, is duplicated to improve its signal to noise ratio. Fig. 1 provides information regarding the MTI spectral bands.

For further details about the MTI satellite, its mission and its data products, the interested reader is referred to Refs [13–15].

4. APPLICATION TO LAVA-FLOW CLASSIFICATION ON MAUNA LOA

4.1. Mauna Loa

Mauna Loa is a giant, active basaltic shield volcano which rises over 4 km above sea level, another 5 km above the north-central Pacific seafloor, and another 8 km above the isostatically depressed seafloor of the Pacific Plate, for a total volcanic height of 17 km. It is the most voluminous volcano on Earth, with a subareal surface of over 5,000 km^2 (half the Island of Hawaii). Mauna Loa is one of the Earth's most active volcanoes, having erupted more than 30 times since its first documented historical eruption in 1843. Mauna Loa has been selected as one

of 15 "Decade Volcanoes" by the International Association of Volcanology and Chemistry of the Earth's Interior (IAVCEI). This status provides opportunities for increased multidiscplinary and multinational efforts to understand volcanic processes.[16,17]

We chose to study the summit region of the Mauna Loa caldera for several reasons. Firstly, most of its surface is covered by broad lava flows with distinct edges and ranging over a thousand years of age. Here, mixed pixel effects during classification are minimized. However, there is spectral variation within flows due to several factors which can confound spectral classification including texture types (aa and pahoehoe), the development and character of glassy crusts, iron oxidation, silica veneers, and surface spalling. Secondly, a modern geological map has been published by the USGS[18,19] which can serve as validation data for the analysis. The map does not distinguish between many features apparent in the images but are not considered to be geologically important. The summit of Mauna Loa is about 4,000 m above sea level and there is a considerable range of elevations within the image (roughly 1500 m). Since the flows run downhill, many of them cover a large range of elevation and subsequent variation of atmospheric path length. Furthermore, basalt from the same flows appear at different slopes (mostly azimuthally) resulting in different illumination and heating of the same material.

MTI imagery, with spectral information ranging from visible to thermal IR wavelengths, and its high spatial resolution, is a natural choice for detailed mapping of Mauna Loa's lava flows and other volcanic features. Figure 2 (a) shows a "true color" representation of an MTI image taken of Mauna Loa. This image data was collected on 9th August 2001. For these experiments we used multispectral MTI data that covered the visible and infra-red regions. The data used had 16 bands, ranging from approximately 0.5 μm to 10 μm. The bands were coregistered and georeferenced[20] and had a nominal ground sample distance of 20 m.

The training data were provided by an analyst, who used ground truth available in the form of a USGS geological map of the island of Hawaii, as well as his own judgement, and the expert knowledge of others. Eleven (11) training classes were provided, based largely on the information provided in the USGS ground truth, but including input from other knowledgeable sources. Table 1 provides a list of the training classes, together with the appropriate color designation used in the training and results images shown in Fig. 2 (b) - (e) and Fig. 3 (b) - (e). It should be noted that for classes k3 and k5, as depicted on the USGS geological map, there were significant textural differences and subtle spectral differences within the class as defined in the USGS classification. We therefore chose to sub-divide these classes into two sub-classes each (i.e. k3(a), k3(b), k5(a) and k5(b)). Table 2 provides some additional information regarding the lava flow classes used in these experiments.[19] Note, with reference to the age information in Table 2, 1 ka = 1000 radiocarbon years before "present" (A.D. 1950), e.g. 1.5 ka is approximately A.D. 450. Pahoehoe and aa are physically distinct lava textural types derived from solidification under different rheological and flowage conditions. Individual Hawaiian lava flows generally include areas of both pahoehoe and aa texture, as well as textures gradational between the two types. Pahoehoe and aa have identical chemical and mineralogical compositions within individual lava flows and differ in their genesis only by their physical conditions of emplacement. Pahoehoe is formed by direct freezing in place from flowing liquid; aa reaches its final resting place as a complex aggregate of individual blocks which have formed by the fracturing of solidifying lava during flowage. Lava may transform from pahoehoe to aa during flow; the opposite has not been observed. Many Hawaiian basaltic lavas are initially erupted as pahoehoe, but commonly transform to aa downstream from eruptive events.[21] Even though pahoehoe and aa have essentially identical chemical and mineralogical compositions, they have different spectra due to multiple scattering and absorption/emission: aa is darker in the reflectance bands and flatter in the thermal.

5. COMPARISON WITH OTHER CLASSIFICATION TECHNIQUES

For our comparisons we essentially conducted three separate classification experiments, using the same data, but three different supervised classification techniques: (1) GENIE with its multiple-class classification capability for three different backends (maximum likelihood, minimum Euclidean distance and minimum spectral angle), (2) Conventional supervised multiple-class classification, using ENVI's maximum likelihood, minimum distance and spectral angle mapper and (3) Multiple GENIE two-class classifications, combined to produce a multiple-class classification. We trained the classifiers on one MTI image and tested the resulting classifiers on two different MTI images. We then made objective comparisons of the classification results on in-sample (training) and out-of-sample (test) data. The criterion for making the objective comparisons was the fitness achieved by each classification technique, as described above in Eq. 1.

Table 1. Class description and color mapping for training data and results images shown in Fig. 2 (b) - (e) and Fig. 3 (b) - (e).

Class Description	Color
Clouds	Red
Cloud Shadows	Green
Lava Flow, Type k1y	Blue
Lava Flow, Type k2	Yellow
Lava Flow, Type k3 (a)	Cyan
Lava Flow, Type k3 (b)	Magenta
Lava Flow, Type k4	Maroon
Lava Flow, Type k5 (a)	Sea Green
Lava Flow, Type k5 (b)	Purple
Ejecta	Orange
Sulphur/Sulphur Compounds	Sienna

Table 2. Age and weathering characteristics of Holocene lava flows of the Kau Basalt, southwest rift zone of Mauna Loa Vocano, Hawaii

Unit	Age (ka)	Dominant color of lava surface		Weathering character of pahoehoe surface
		Pahoehoe	Aa	
k5	0.0–0.2	Black	Dark Brown	Unweathered shiny glass surface
k4	0.0–0.2	Black to dark gray	Light brown	Original glassy surfaces well preserved; only slight weathering
k3	0.75–1.5	Gray to tan	Yellow-tan	Some original surfaces; thin weathering rind
k2	1.5–3.0	Light brown	Tan-orange	Original surfaces mostly destroyed
k1	3.0–10.0	Brown, orange brown and red brown	Red-orange	Original surfaces destroyed; deep weathering rind

The training image was MTI image number 0105721 and, as mentioned previously, was obtained on 9th August 2001. The test (out-of-training-sample) images were MTI image numbers 0031823 and 0107309 and were obtained on 9th October 2000 and 22nd October 2001, respectively.

5.1. Conventional Supervised Classification

Many implementations of standard supervised classifiers exist. One of the most widely used remote-sensing software packages is the ENvironment for Visualizing Imagery (ENVI),[22] which is built on IDL and is distributed by Research Systems, Inc.[23] Supervised classification techniques provided as part of the ENVI package were used in the comparison experiments with GENIE.

We used the GENIE and ENVI classifiers to clasify every pixel in the input data into one of the training classes, with no "unclassified" pixels being allowed. For applying the ENVI-supplied classifiers to out-of-training-sample data, the training data (reference spectra) used in the training was provided.

The following provides a brief description of the ENVI-supplied supervised classification techniques[24] used in the comparison experiments.

5.1.1. (MIN) Minimum Distance

The minimum distance supervised classification technique[8,24] computes the mean pixel vector of the "feature" class, and then assigns new pixels to the "feature" class based on the Euclidean distance from that pixel to the mean. For the multi-class case, the pixel is assigned to the feature whose mean value is the minimum distance from the pixel.

5.1.2. (MAX) Maximum Likelihood

Maximum likelihood classification is the most common supervised classification method used with remote sensing data,[8] and among the classifiers considered here, the one with the most free parameters. Here each class is modelled with separate multivariate gaussian distributions. New pixels are assigned to the class that had the highest probability of generating that pixel.

5.1.3. (SAM) Spectral Angle Mapper

The spectral angle mapper (SAM) technique[11] is motivated by the observation that changes in illumination caused by shadows, slope variation, sun position, light cloud, etc., approximately only alter the magnitude of a pixel's vector, rather than the direction. Therefore we can eliminate these effects by normalizing all pixel vectors to unit magnitude and then looking at the angle between a given pixel and the mean vector for the "feature" class. Pixels are assigned to the "feature" class if this angle is less than a user-defined threshold. For our experiments this threshold was set large enough that every pixel had to be assigned to a class, with no pixels left unclassified.

5.2. GENIE Classification

5.2.1. Direct Multiple-Class Classification

GENIE required modification from its two-class incarnation to enable multiclass classification. The fundamental modification was to add a conventional classifier as a backend. As mentioned earlier, GENIE searches for suitable algorithms for particular classification tasks. Therefore, GENIE's output is not a classification result, but a classification algorithm. Therefore, with respect to performing classification of data, once GENIE has been trained on a problem, it is then necessary to apply the algorithm found to the data (including the backend parameters found during training).

5.2.2. Combining Multiple Two-Class Classifications

As well as the direct approach for achieving multiclass classification from GENIE, we also investigated a scheme in which multiple binary classifications were combined to produce a single multiclass classification. For a K-class problem, we perform K binary classification tasks, where the k'th task treats the k'th feature as "true" and the other $K-1$ features as "false". Each of these K classification tasks produces a grayscale result, with a threshold that provides the best classfication for the two-class problem. We normalize this grayscale by first subtracting the threshold from all the grayscale values (so that the effective threshold is zero) and then separately histogram-equalizing the positive and negative values. Finally, we combine these K grayscale images into a single classification by taking, for each pixel, the value of k that has the largest grayscale value.

6. RESULTS

Table 3 shows the scores for all the classifiers on both the (in-sample) training and (out-of-sample) test data. This table also shows the average score for each classifier for both (out-of-sample) test data sets.

Fig. 2 (a) shows a true color (Red = band C: 0.62 – 0.68 μm, Green = band B: 0.52 – 0.60 μm, Blue = band A: 0.45 – 0.52 μm) representation of the training data, MTI Image number 0105721. Fig. 2 (b) shows the training classes provided for the training data overlaid on top of the true color image shown in Fig. 2 (a). Figure 2 (c) shows the result of applying the ENVI maximum likelihood classifier to the training data. Figure 2 (d) shows the result of applying the GENIE multiple-class algorithm using the maximum likelihood backend to the training data. Figure 2 (e) shows the result of applying the combined GENIE 2-class classifiers to the training data.

Fig. 3 shows the classification results for the best ENVI and GENIE classifiers on some out-of-sample (test) data. Fig. 3 (a) shows a true color representation of the test data, MTI Image number 0031821, (b) shows the training classes provided to enable a fitness score for the classifiers to be determined, the overlaid on top of the true color image, (c) shows the result of applying the ENVI spectral angle mapper classifier to the test data, (d) shows the result of applying the GENIE multiple-class algorithm using the minimum distance backend to the test data, and (e) shows the result of applying the combined GENIE 2-class classifiers to the test data.

Tables 4 – 7 show confusion matrices for four different classifiers. The columns correspond to true classes and the rows to the classes that were estimated by the given classifier. A good classifier has the largest elements on the diagonal. Large off-diagonal elements correspond to important misclassifications. For instance, Table 7 shows that the ENVI (SAM) classifier failed to identify any of the actual k4 material, mostly labelling it as k5(b).

Table 3. Scores obtained by classification techniques on the (in-sample) training and (out-of-sample) test data

	GENIE: Multiple-Class			ENVI Multiple-Class			GENIE
	Max.	Min.	S.A.M.	Max.	Min.	S.A.M.	Combined 2-Class
(In-sample) training data							
MTI Image # 0105721	979	905	930	942	723	784	946
(Out-of-sample) test data							
MTI Image # 0031823	152	659	378	125	263	389	390
MTI Image # 0107309	45	443	381	31	74	348	465
Average (test data)	99	551	380	78	168	368	428

Table 4. Confusion matrix for result of multiple-class GENIE with maximum likelihood backend applied to training data

GENIE (Max) Classes	Training Class											Total
	Clouds	Shadows	k1y	k2	k3(a)	k3(b)	k4	k5(a)	k5(b)	Ejecta	Sulphur	
Clouds	3411	15	3	16	272	6	22	0	1	0	1	3747
Shadows	12	668	0	2	44	1	11	53	56	1	0	848
k1y	0	0	494	0	0	0	0	0	0	0	0	494
k2	0	0	0	1908	49	7	13	60	0	0	0	2037
k3(a)	7	0	0	11	25148	68	0	0	0	0	0	25324
k3(b)	1	0	0	19	154	6562	81	2	0	0	0	6819
k4	1	1	0	19	64	121	3572	75	90	0	0	3943
k5(a)	0	0	0	13	0	0	12	4912	44	0	0	4981
k5(b)	0	0	0	0	0	0	13	0	20869	0	0	20882
Ejecta	0	0	0	0	0	0	0	0	0	585	0	585
Sulphur	4	0	0	0	11	0	0	0	0	0	116	131
Total	3436	684	497	1988	25742	6765	3724	5102	21060	586	117	69701

7. DISCUSSION

The first thing one can say regarding the results of these comparison experiments is that the best GENIE algorithm was able to outperform the best standard ENVI algorithm on the tasks given, for both in sample (training) and out-of-sample (test) data. The huge drop in performance from training data to out-of-training-sample data for all the ENVI classifiers seems quite remarkable. However, it is not really surprising given that there are 11 classes and

Table 5. Confusion matrix for result of ENVI maximum likelihood classification of training data

ENVI (Max) Classes	Training Class											Total
	Clouds	Shadows	k1y	k2	k3(a)	k3(b)	k4	k5(a)	k5(b)	Ejecta	Sulphur	
Clouds	3413	11	4	16	194	17	92	14	13	0	0	3774
Shadows	6	631	1	6	31	4	54	79	115	0	0	927
k1y	0	0	492	0	19	0	0	0	0	0	0	511
k2	0	1	0	1900	1322	153	60	132	3	0	0	3571
k3(a)	3	1	0	8	23884	133	51	2	3	0	0	24085
k3(b)	0	2	0	49	237	6370	580	0	0	0	0	7238
k4	0	31	0	3	34	88	2704	64	64	0	0	2988
k5(a)	0	1	0	5	0	0	183	4798	72	0	0	5059
k5(b)	0	0	0	0	0	0	0	11	20768	0	0	20779
Ejecta	0	0	0	0	0	0	0	0	0	586	0	586
Sulphur	14	6	0	1	21	0	0	2	22	0	117	183
Total	3436	684	497	1988	25742	6765	3724	5102	21060	586	117	69701

Table 6. Confusion matrix for result of multiple-class GENIE with minimum distance backend applied to out-of-training-sample data (MTI image number 0031823)

GENIE (Min) Classes	Training Class											Total
	Clouds	Shadows	k1y	k2	k3(a)	k3(b)	k4	k5(a)	k5(b)	Ejecta	Sulphur	
Clouds	4164	44	0	0	0	0	0	0	0	0	0	4208
Shadows	4	816	0	6	0	0	15	0	0	0	0	841
k1y	0	3	13	0	0	0	0	0	0	0	0	16
k2	0	46	0	146	5	0	1	0	3	0	0	201
k3(a)	0	0	157	0	9283	0	1	0	8	7	0	9456
k3(b)	0	1	86	5	3638	108	5	0	0	0	0	3843
k4	0	0	7	342	4279	2709	146	21	12	3	0	7519
k5(a)	0	31	0	363	17	65	425	952	133	1	0	1987
k5(b)	24	193	0	50	17	0	96	303	10129	3	0	10815
Ejecta	0	0	0	0	4	0	0	0	96	414	0	514
Sulphur	305	0	0	0	0	0	0	0	0	0	7	312
Total	4497	1134	263	912	17243	2882	689	1276	10381	428	7	39712

Table 7. Confusion matrix for ENVI spectral angle mapper applied to out-of-training-sample data (MTI image number 0031823)

ENVI (SAM) Classes	Training Class											Total
	Clouds	Shadows	k1y	k2	k3(a)	k3(b)	k4	k5(a)	k5(b)	Ejecta	Sulphur	
Clouds	4163	83	0	0	0	0	15	0	0	0	2	4263
Shadows	37	44	0	246	7219	1994	40	0	4	87	0	9671
k1y	0	0	61	0	0	0	0	0	0	0	2	63
k2	0	0	28	9	5215	8	0	0	0	0	0	5260
k3(a)	0	1	38	0	442	1	0	0	0	0	0	482
k3(b)	0	0	134	0	107	0	0	0	0	0	0	241
k4	14	2	1	14	2375	37	0	0	0	0	0	2443
k5(a)	2	1	1	64	72	48	17	0	0	0	0	205
k5(b)	8	987	0	511	45	756	608	1275	10376	3	0	14569
Ejecta	0	16	0	68	1768	38	9	1	1	338	0	2239
Sulphur	273	0	0	0	0	0	0	0	0	0	3	276
Total	4497	1134	263	912	17243	2882	689	1276	10381	428	7	39712

hence a lot of adjustable paramters, and therefore ample opportunity for overfitting. With the maximum likelihood classifiers, having a larger number of free parameters than the other classifiers, this problem with overfitting becomes even more apparent and is clearly evident in the results. The classifiers employing the maximum likelihood method (even for the GENIE classifiers) were the best performing on the training data but were by far the worst on the out-of-training-sample data. Of all the ENVI classifiers, the spectral angle mapper had the least drop in performance, but that is not surprising: it should be more robust to overall illumination effects and also has the least number of free paramters.

8. CONCLUSIONS

We have described how we have further enabled GENIE with the capability to generate classification algorithms able to simultaneously classify an image into multiple classes. We have also compared the performance of this new capability with that of standard, commonly-used techniques for doing this multiple-class classification. We have demonstrated that GENIE thus enabled still outperforms these standard classifiers for both training and out-of-training sample data. We have also described and demonstrated an alternative method for using GENIE to obtain multiple-class classifications. This method combines multiple two-class GENIE classifications. This was also able to outperform the standard classifiers both for in-sample and out-of-sample data, but not to the same extent as the more direct multiple-class GENIE technique.

REFERENCES

1. S.P. Brumby, J. Theiler, S.J. Perkins, N.R. Harvey, J.J. Szymanski, J.J. Bloch and M. Mitchell, "Investigation of Image Feature Extraction by a Genetic Algorithm", in *Proc. SPIE 3812* pp. 24–31 (1999).
2. J. Theiler, N.R. Harvey, S.P. Brumby, J.J. Szymanski, S. Alferink, S. Perkins, R. Porter and J.J. Bloch, "Evolving Retrieval Algorithms with a Genetic Programming Scheme", in *Proc. SPIE 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", in *Proc. SPIE 4132*, pp.72-82, 2000.
4. 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", in *Proc. SPIE 4099*, pp. 480–490, (2000).
5. N.R. Harvey, S. Perkins, S.P. Brumby, J. Theiler, R.B. Porter, A.C. Young, A.K. Varghese, J.J. Szymanski and J. Bloch, "Finding golf courses: The ultra high tech approach", in *Evolutionary Image Analysis, Signal Processing and Telecommunications*, Poli, et al. Springer-Verlag (2000).
6. A.B. Davis, S.P. Brumby, N.R. Harvey, K. Lewis Hirsch, and C.A. Rohde, "Genetic refinement of cloud-masking algorithms for the multi-spectral thermal imager (MTI)" in *Proc. IGARSS 2001*, Sydney, Australia, 9-13 July 2001.
7. 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, "Comparison of GENIE and Conventional Supervised Classifiers for Multispectral Image Feature Extraction", to be published in *IEEE Trans. Geoscience and Remote Sensing*, 40:2, (2002).
8. J.A. Richards and X. Jia, "Remote Sensing Digital Image Analysis", New York: Springer-Verlag, 1999.
9. R.A. Schowengerdt, "Remote Sensing: Models and Methods for Image Processing", Academic Press, San Diego, CA, 1997.
10. R.O. Duda, R.E. Hart, and D.G. Stork, "Pattern Classification", John Wiley and Sons, New York, 2001.
11. F.A. Kruse, A.B. Lefkoff, J.B. Boardman, K.B. Heidebrecht, A.T. Shapiro, P.J. Barloon, and A.F.H. Goetz, "The Spectral Image Processing System (SIPS) - Interactive Visualization and Analysis of Imaging Spectrometer Data," *Remote Sens. Environ.*, vol. 44, pp. 145–163, (1993).
12. C.M. Bishop, *Neural Networks for Pattern Recognition*, pp. 105–112, Oxford University Press (1995).
13. `http://nis-www.lanl.gov:80/nis-projects/mti/`
14. P.G. Weber, B.C. Brock, A.J. Garrett, B.W. Smith, C.C. Borel, W.B. Clodius, S.C. Bender, R.R. Kay, and M.L. Decker, "Multispectral thermal imager mission overview", *Proc. SPIE 3750*, pp. 340–346 (1999).
15. J.J. Szymanski, W. Atkins, L. Balick, C.C. Borel, W.B. Clodius, W. Christensen, A.B. Davis, J.C. Echohawk, A. Galbraith, K. Hirsch, J.B. Krone, C. Little, P. McClachlan, A. Morrison, K. Pollock, P. Pope, C. Novack, K. Ramsey, E. Riddle, C. Rohde, D. Roussel-Dupr, B.W. Smith, K. Smith, K. Starkovich, J. Theiler, and P.G. Weber. "MTI Science, Data Products and Ground Data Processing Overview." *Proc SPIE 4381*, pp. 195–203, (2001).
16. J.P Lockwood, and J.M. Rhodes, "Mauna Loa: A Decade Volcano", *Periodica di Mineralogia (Rome)*, Vol. **44**, pp. 45–47, (1995).
17. `http://www.soest.hawaii.edu/mauna_loa/`
18. `http://www.usgs.gov/`
19. E.W. Wolfe, and J. Morris, "Geological Map of the Island of Hawaii", *U.S. Dept. of the Interior, U.S. Geological Survey*, (1996).
20. J. Theiler, A. Galbraith, P. Pope, K. Ramsey, and J. Szymanksi, "Automated Coregistration of MTI Spectral Bands", *Proc. SPIE 4725*, (2002).
21. A.B. Kahle, A.R. Gillespie, E.A. Abbott, M.J. Abrams, R.E. Walker, G. Hoover, J.P Lockwood, "Relative dating of Hawaiian lava flows using multispectral thermal infrared images: a new tool for geologic mapping of young volcanic terranes", *Journal of Geophysical Research*, vol. 93, no. B12, pp. 15239-51–15361-3 (1988).
22. `http://www.rsinc.com/envi/index.cfm`
23. `http://www.rsinc.com`
24. `http://www.rsinc.com/Envi/tut2.cfm`

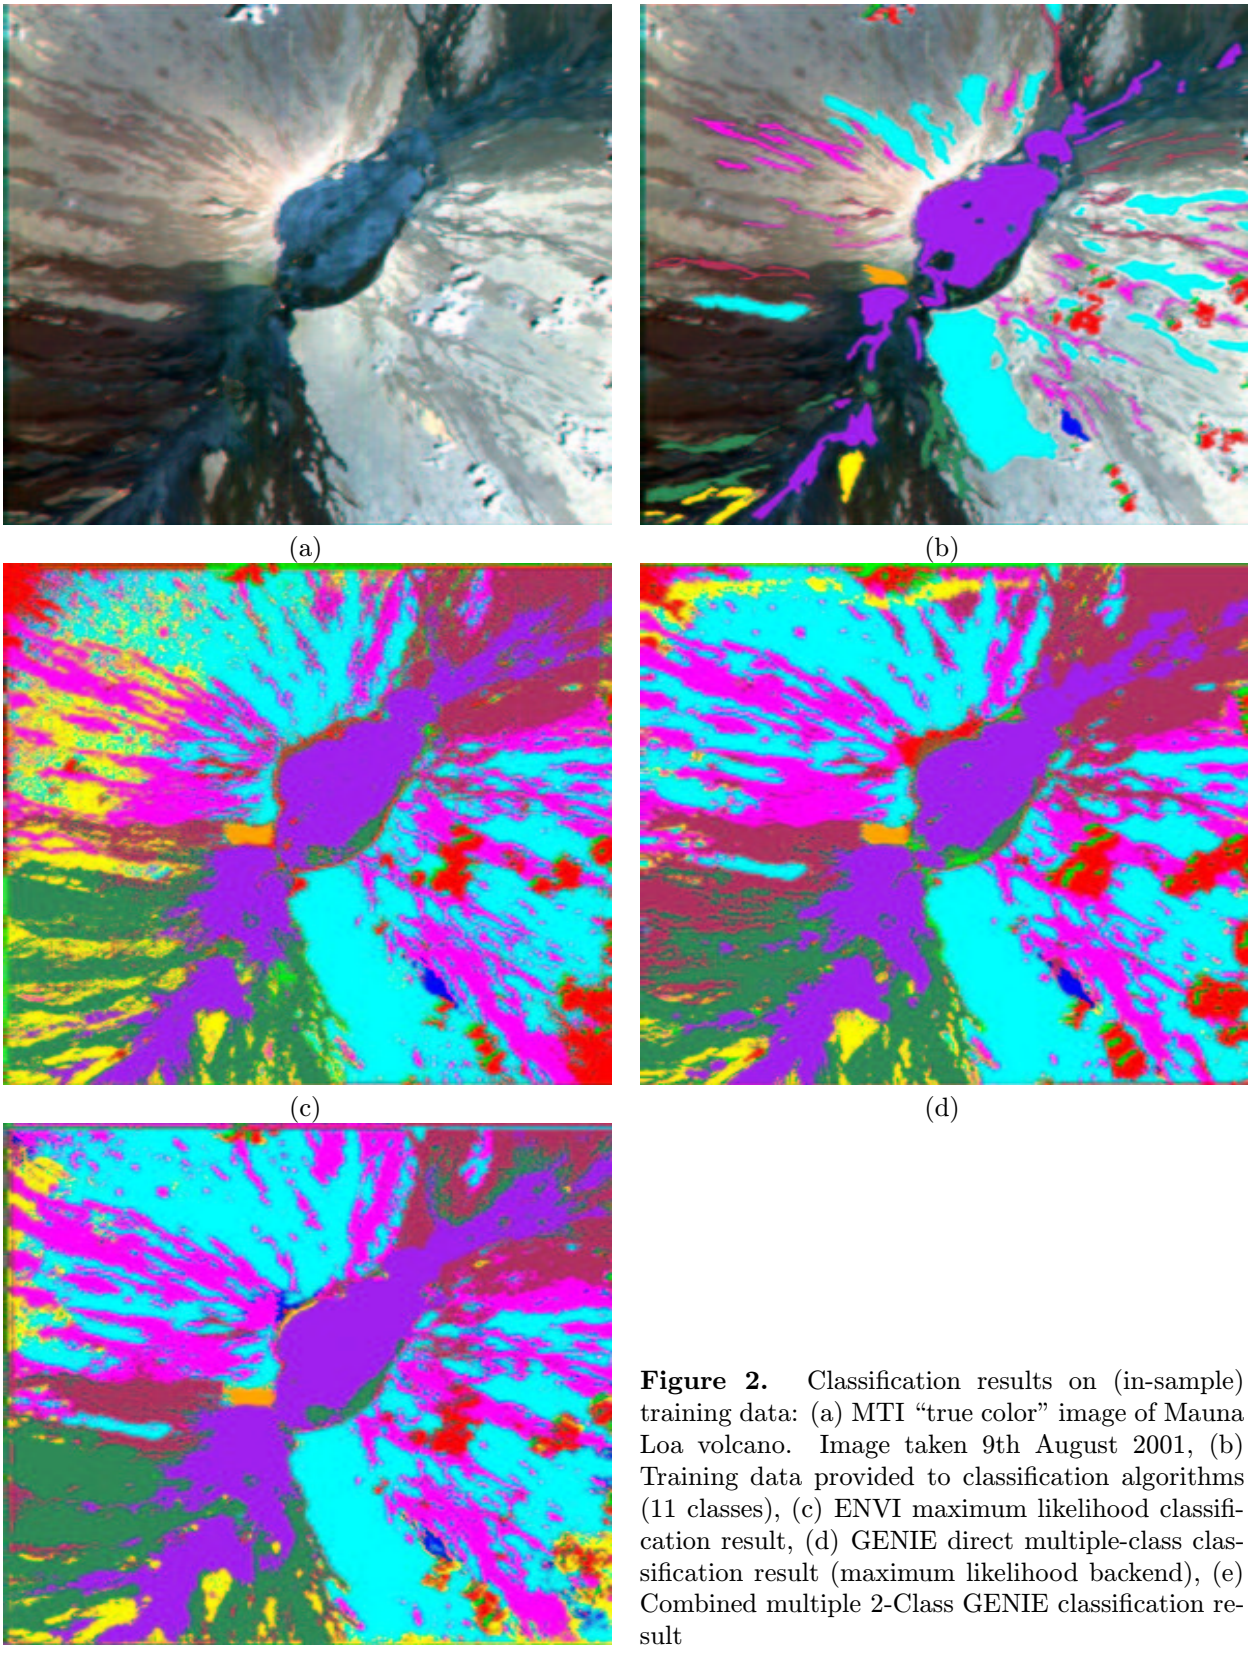

(a) (b) (c) (d) (e)

Figure 2. Classification results on (in-sample) training data: (a) MTI "true color" image of Mauna Loa volcano. Image taken 9th August 2001, (b) Training data provided to classification algorithms (11 classes), (c) ENVI maximum likelihood classification result, (d) GENIE direct multiple-class classification result (maximum likelihood backend), (e) Combined multiple 2-Class GENIE classification result

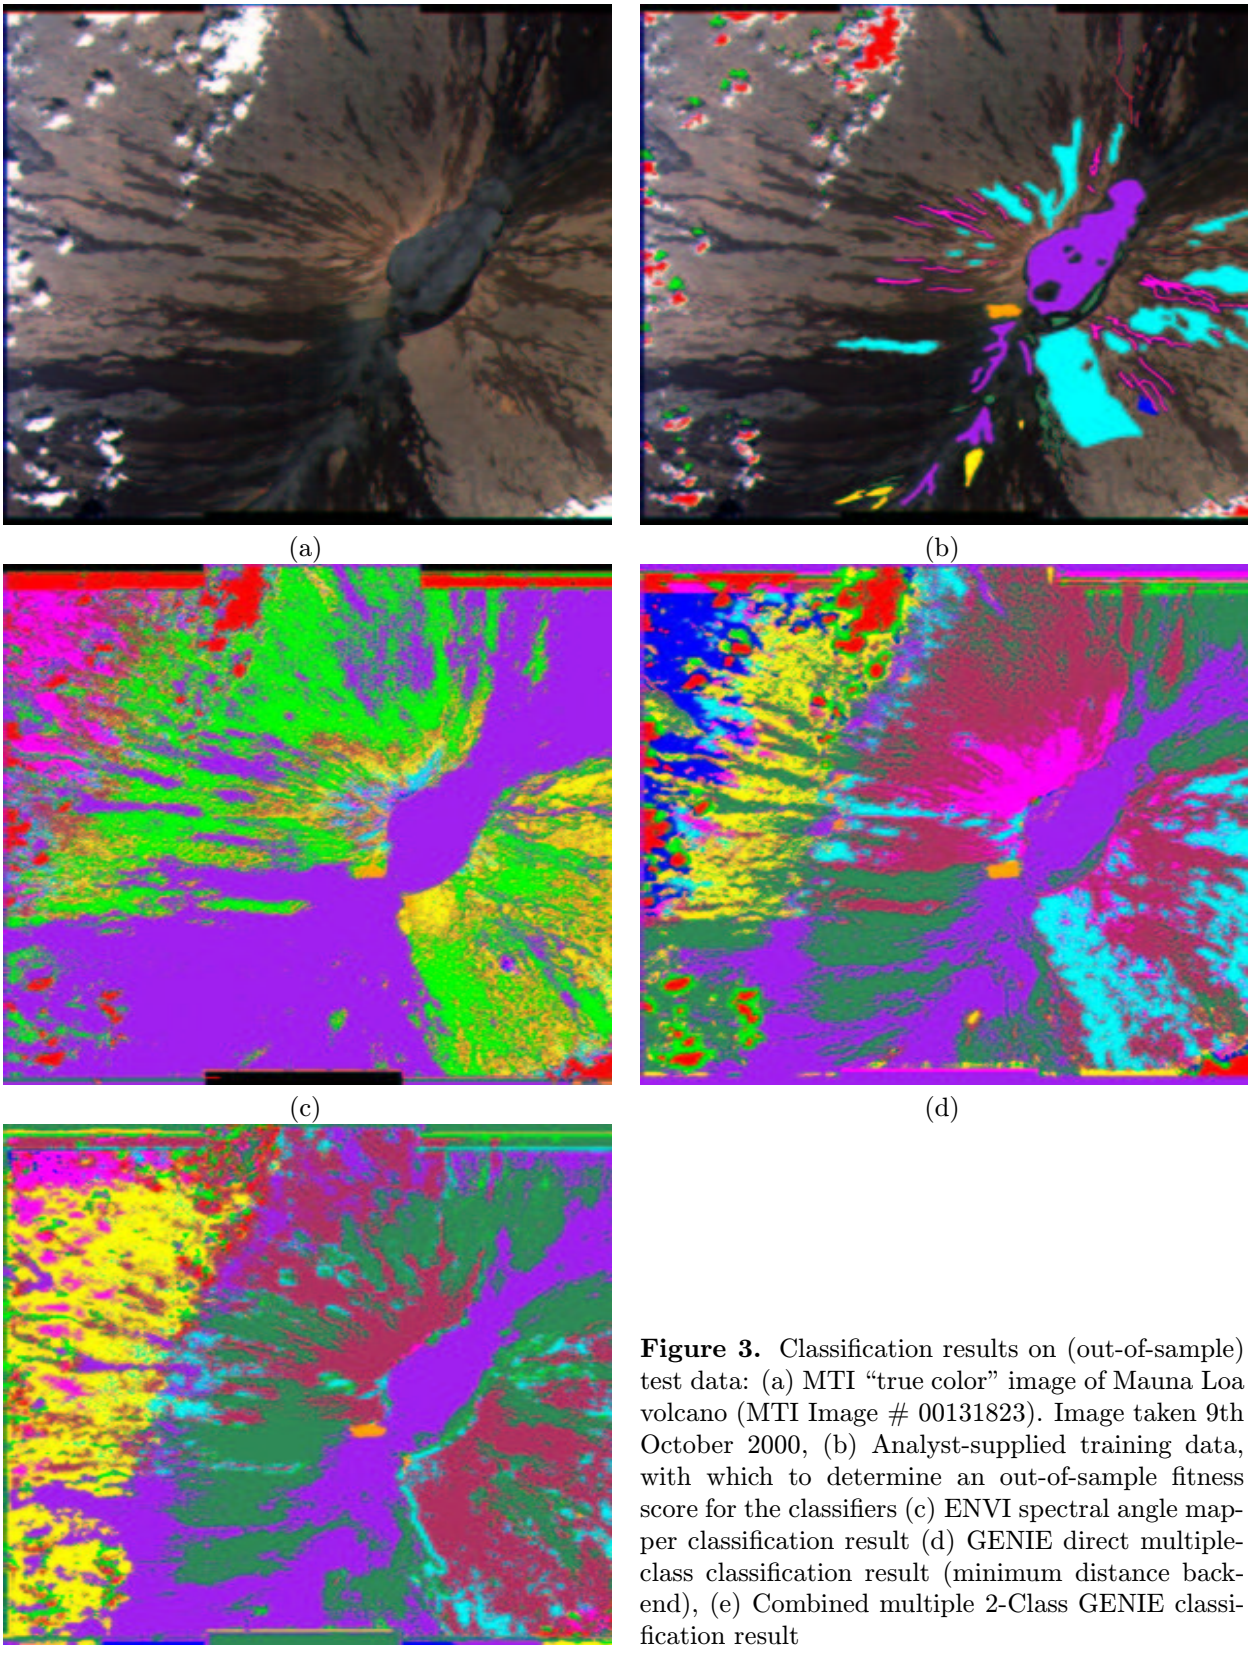

Figure 3. Classification results on (out-of-sample) test data: (a) MTI "true color" image of Mauna Loa volcano (MTI Image # 00131823). Image taken 9th October 2000, (b) Analyst-supplied training data, with which to determine an out-of-sample fitness score for the classifiers (c) ENVI spectral angle mapper classification result (d) GENIE direct multiple-class classification result (minimum distance backend), (e) Combined multiple 2-Class GENIE classification result