

Use of Shape Representation and Similarity in Classification of UXO in Magnetometry Data

Jim R. McDonald
AETC Incorporated
120 Quade Dr.
Cary, NC 27513

David Opitz
Rocky Mountain Adaptive Software.
2822 Carla Jo Lane
Missoula, MT 59803

Final Report

SERDP Project UX-1354



Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 08 AUG 2004		2. REPORT TYPE Final		3. DATES COVERED -	
4. TITLE AND SUBTITLE Use of Shape Representation and Similarity in Classification of UXO in Magnetometry Data				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) 1. Jim R. McDonald 2. David Opitz				5d. PROJECT NUMBER UX-1354	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) 1. AETC Incorporated 120 Quade Dr. Cary, NC 27513 2. Rocky Mountain Adaptive Software. 2822 Carla Jo Lane Missoula, MT 59803				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Strategic Environmental Research & Development Program 901 N Stuart Street, Suite 303 Arlington, VA 22203				10. SPONSOR/MONITOR'S ACRONYM(S) SERDP	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 22	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

1. Introduction

Over the past few years sophisticated passive and active metal detectors have been coupled with state-of-the-art (SOA) GPS systems to provide increasingly accurate real-time localization capability when conducting Unexploded Ordnance (UXO) surveys. However, clearing UXO ranges with these automated techniques invariably still requires digging 5-100 items for every recovered intact ordnance item. In an attempt to improve discrimination and reduce false alarms, target analysis techniques have been developed to process UXO survey data sets, emphasizing statistical analysis approaches applied to the output parameters of physics-based target fitting algorithms. The objective of this project is to go beyond the use of physics-based parameters when making decisions about ordnance classification. We have incorporated shape representation and similarity based on anomalies in the mapped data files to extract and exploit image features related to the target signatures. This provides an additional input for a discrimination/classification decision.

Shape information is an important component of the semantic content of the UXO target image and is a primary component of the visual decision making process used by the human analyst in the current interactive data analysis approach. Because this information is so important to the human in-the-loop, if it can be quantified and incorporated into the machine analysis of the data, it will provide an important classification tool. Indeed, if the target analysis process is ever to be fully and effectively automated, implementation of this step is imperative.

This project employed existing data sets taken by the Vehicular Multi-sensor Towed Array Detection System (MTADS) at the Badlands Bombing Range (BBR) during 2001.¹ In preparation for that survey a 10-acre area near the bull's eye (which had previously been cleared of UXO in 1999)² was seeded by ERDC with 25 projectiles (ten 105-mm projectiles, ten 155-mm projectiles, and five 8-in projectiles). These inert projectiles were buried in a variety of orientations at depths up to the maximum expected self-burial depths for these ground-fired projectiles. One additional pre-existing HE-filled buried projectile was discovered in the 10-acre area. This 10-acre area and an additional 100 acres were surveyed by the Vehicular MTADS. The entire data set was interactively analyzed using standard techniques by an analyst using the MTADS data analysis system (DAS); a target list was prepared and all targets on the list were dug by UXO technicians.

The pattern recognition team on this project was provided the complete ground truth (locations and identities of all dug targets, including ordnance and non-ordnance (metallic scrap) items from the 10-acre seeded area. The data for the additional 100 acres was provided to the pattern recognition team as a mapped data file. The dig results and target location ground truth for this area were blind to the group until final analysis results had been submitted to the P.I.

To begin the study, the pattern recognition team applied statistical analysis approaches based on the physics-derived fitting parameters resulting from the previously completed interactive analysis of the seed target area. These techniques are similar to those that have been applied by others using datasets from other ranges.³⁻⁶ They had not been previously applied to data from the BBR. Section 2 provides more background on these techniques.

The culminating step in this research project involved (a) developing parameter sets that describe shape information from the magnetic anomaly image map and (b) developing techniques to automatically filter these parameters and learn from them by submitting them to an inductive learning algorithm. The goal of the project was to develop an automated analysis method that is as accurate, if not more accurate, than the interactive human interpretation.

In the body of the report we describe:

1. The comparative results of applying Genetic Ensemble Feature Selection (GEFS) versus Artificial Neural Networks (ANNs),
2. Our analysis of the characteristics of the outliers (targets both UXO and scrap that defied correct characterization), and
3. A reanalysis of the data using Principal Components Analysis (PCA) and clustering techniques.

Based on these results, we concluded that additional information needed to be generated to give to the learning algorithm. We found that improving the inductive learning methods using only the physics-based parameters does not improve accuracy. Instead, one needs to investigate what further types of information are available from the raw MTADS dataset to give to an inductive learning algorithm.

The project then proceeded by:

1. Using a spatial array of the raw pixel values associated with apparent magnetic anomalies from the MTADS mapped data,
2. Adding these spatial arrays of raw pixel values to the physics parameters derived from the MTADS DAS dipole fitting, and
3. Resubmitting these parameters to the GEFS and ANN algorithms.
4. Automating the entire process by using a two-step approach with an existing commercial software product, Feature Analyst[®] to (a) find candidate target sites (i.e. we implemented an automated target picker) and (b) characterize these target sites using shape function information and shape filters resident in the Feature Analyst software.

The report concludes with an analysis of these results and a prediction of how they might be used to develop a completely-automated data analysis approach that would require only preprocessing the raw data to produce mapped data files.

2. Background

The Current MTADS Process

Modern geophysical UXO surveys are typically conducted under GPS control using systems such as the MTADS. This vehicular-towed array of magnetometer sensors employs a 0.25-m horizontal sensor spacing and a 50 Hz data rate. In a UXO geophysical survey this system typically records over 2 million sensor data points/acre. Currently the target analysis process begins by mapping these data to create a magnetic anomaly image map. Figure 1 shows an MTADS magnetic anomaly map prepared from the mapped data file for the 10-acre seed target area at the BBR. Figure 2 shows a 0.25-acre portion of the 10-acre area, similar to what would be viewed by the analyst during target picking. In the analysis process, the analyst visually inspects the data and eliminates anomalies that clearly fall outside the parameters that would be characteristic of UXO types known or suspected to be present on the range. This process obviously relies strongly on the experience of the analyst. He has access to a library of known UXO magnetic anomaly signatures at various orientations and depths for reference. However, more important in the prescreening process, is his knowledge of and ability to identify signatures that are atypical of UXO signatures. These include anomalies that are composed of contributions from numerous smaller objects, objects too large or misshapen to be UXO, or anomalies that are too small to meet the minimal threshold criteria for the smallest expected UXO.

After visual prescreening to eliminate non-candidate anomalies, the evaluation continues with the analyst using the computer mouse to select areas containing anomaly signatures of potential buried UXO items, as shown by the white outlines around anomalies in Figure 2. Data in each selected area are submitted to a physics-based modeling routine, which uses a magnetic dipole signal model. The model parameters (location, depth, size, orientation, and magnetic moment) are systematically varied until the difference between the measured data and the model is minimized. In the analysis window image, shown as a screen clip in Figure 3, the analyst has the option to further edit the data to remove multiple target signature contributions to the image, and to refit the signature. Ultimately, the analyst rejects the fit or accepts the fit and provides comment about the individual target, its identity, its likelihood of being UXO, or unique attributes of the fitting process. The quality of the fit approximating a magnetic dipole signature is a final numerical output parameter.

These physics-based parameters form the input for statistical analysis approaches (inductive learners) being developed in this project to improve and automate the target analysis process. The goal of incorporating inductive learning into the MTADS process is to (a) improve accuracy, (b) provide consistency in predictions, and (c) ultimately remove the man-in-the-loop required in the current process.

Inductive Learning: ANNs and GEFS

A computer program that learns from a set of labeled examples is called an inductive learner. A teacher provides the output for each example. The set of labeled examples given to an inductive learner is called a training set. The task of inductive learning is to generate from the training set a




Figure 1. Magnetic anomaly image of the 10-acre Seed Target area at the BBR prepared from the mapped data file.




Figure 2. Magnetic anomaly image of a 0.25 acre part of the Seed Target area. Signatures of 3 projectiles (a 105-mm, a 155-mm, and an 8-in) are shown. They are buried 3.5-4.5 feet deep.

concept description that correctly predicts the output of all future examples, not just those from the training set. Several previously studied⁷⁻⁹ inductive-learning algorithms have proven successful for interpreting imagery.¹⁰⁻¹² These algorithms differ both in their concept-representation language and in their method (or bias) for constructing a concept within this language. These differences are important because they determine which concepts a classifier will learn. For instance, some are more robust to noisy input data; some are well suited for real-valued inputs, and so on. The choice of a learning algorithm should be made by the user to match the task at hand. In this project we investigate various learning approaches to determine the ones best suited for distinguishing legitimate UXO from metallic scrap.




Figure 3. Screen shot of the MTADS analysis image window showing the fit of target 156, an 8-in projectile. Sensor data are shown on the left, the dipole model fit to its right. The physical fit parameters are given below the images.

Artificial Neural Networks⁹ (ANNs) have been successfully applied to a wide variety of real-world domains and are thus a logical choice for consideration for this project. Briefly, ANNs consist of nodes interconnected by weighted links. These networks propagate a set of input signals, representing an image's feature values, forward into a set of output signals that serve as the network's prediction. Nodes that are neither input nor output nodes are called hidden nodes and their sole function is to help map input values to output values. The nodes in the network contain an activation function that typically allows the network to make non-linear predictions. Training of a network consists of modifying the interconnection weights via a learning algorithm such as back propagation,⁹ which we use in this project.

In conjunction with the basic ANNs, we also include in our investigation a powerful inductive-learning method, called *ensembles*, which combines the output of many separate predictors. Figure 4 represents the basic ensemble approach. In Figure 4, neural networks are the classification method used. Each network in the ensemble (networks 1 through N) is trained using the training instances for that network. Then, the predicted output of each of these networks is combined to produce the output of the ensemble ($\hat{O}$ in Figure 4). Previous research has shown that ensembles generally produce more accurate predictions than the individual predictors within the ensemble.^{13, 14}




Figure 4. Predictor Ensemble.

Both theoretical research^{14,15} and empirical work^{16,17} have shown that a good ensemble is one in which (1) the individual networks are accurate, and (2) any errors that these individual networks make occur in different parts of the ensemble's input space. Many learners, when run too long, tend to overfit the training data. Overfitting occurs when the learning algorithm captures too much information on the training examples and not enough on the general characteristics of the domain as a whole. This can lead to poor generalizations. In Figure 5, the dashed outer lines represent predictions from learners. These lines do a good job in their predictions of the training data (the X's), but are not to be trusted outside these predictions. Note that the average of these two predictions (the solid center line) is relatively smooth and would probably make reasonable predictions outside the training points. This is the power of ensembles: Maximizing the difference between individual predictions (the dashed lines) results in smooth averages.

Feature selection in inductive learning is the task of identifying the set of features (i.e., inputs) that are given to a learning algorithm. Automated feature selection is a central problem in inductive learning because there are numerous possible low-level image features, for example pixel values, objects size, object outline characteristics, etc., that can be used by a learning algorithm. Rather than follow the traditional approach of finding one subset of features to give a learning algorithm, Dr. Opitz of RMAS developed the concept of *ensemble feature selection*¹⁸ and described a refined machine-learning algorithm called GEFS (for *Genetic Ensemble Feature Selection*) based on his earlier learning algorithm, ADDEMUP.^{15,17,19} In this approach, GEFS searches for a set of feature subsets that is not only germane to the learning task and learning algorithm, but actually promotes *disagreement* among the ensemble's classifiers. The rationale is that by maximizing the differences between individual well-fit predictors then taking an average of their results, the analyst will arrive at the most accurate prediction (see Figure 5). Since we only know the truth at the "X" points, the goal is to search for a set of predictors that go through most of the X's, but make their errors on different ones (i.e., we want a set of predictors that are accurate, yet diverse in their predictions).




Figure 5. Averaging predictions helps to prevent overfitting.

GEFS uses genetic algorithms (GAs) to generate a set of classifiers (neural networks in this project) that were developed to be accurate and diverse in their predictions. A key GEFS feature is its modified fitness function, which takes into account both accuracy and diversity when estimating the utility of a neural network to be included in the ensemble. Opitz¹⁸ compared the ensemble feature selection approach developed for GEFS's with neural networks as well as the two other state-of-the-art ensemble approaches, Bagging¹³ and Ada-Boosting.^{13,16} Results using 22 data sets showed that the simple and straightforward GEFS algorithm for creating the initial population, on average, produces better learned models than the three other approaches. Results also show that further running of the algorithm with the genetic operators improves performance. Thus, as the learning component of the proposed machine-learning system, GEFS can provide a good answer quickly and then if the user desires, can further improve on this performance. Finally, GEFS's performance is most impressive on domains that contain numerous possible input features,^{20,21} such as visual learning domains. In summary, GEFS is an ideal fit for UXO detection.

3. Studying the Utility of the Physics-Based Parameters

In the initial setup of the project we evaluated the performance of ANNs and GEFS using the physical descriptors provided from the human-in-the-loop interactive target analysis. We constructed a two-class (i.e. UXO and not-UXO) Seeded Site dataset from the vehicle magnetic anomaly data table provided in Appendix G from the NRL report.¹ The table includes the vehicular MTADS predictions of depth, size, moment, inclination, azimuth, and the fit quality for 170 targets that were dug to confirm the existence of ordnance as opposed to scrap. Of the 170 targets, 26 targets were confirmed to be the previously-seeded inert ordnance (ten – 105-mm, ten – 155-mm, and five – 8-in) or live UXO (one – 155-mm). For the two-class datasets in this report, the types of ordnance were undifferentiated.

We conducted the UXO classification tests using leave-one-out cross-validation experiments. With leave-one-out, the learning algorithm creates a model using all of the targets except one, and then predicts the classification of the left-out target using the generated model. This process is repeated until all targets have been left out once and only once. Thus, for the 170 targets, the process is repeated 170 times. The cumulative accuracy of the 170 held-out targets determines the estimate for the predictive accuracy. Leave-one-out has become the standard technique for inductive learning algorithms.

We conducted numerous experiments to tune the various learning parameters for both ANNs and GEFS. We found the error margins to be consistent across the plausible ranges of these parameters. In particular, we found a consistent minimum error of approximately 6.8%, with both ANNs and GEFS. The consistency of the value of the best error rate strongly suggests there are certain targets that cannot be distinguished (classified) using only the physics-based parameters predicted using the MTADS data sets. We confirmed this suspicion in the next step.

Data Characterization of Physical Parameters

In this section we characterize the data in an attempt to identify if the misclassified data are outliers with the current parameters and are thus impenetrable for inductive learning, in general. We started by performing a Principle Components Analysis (PCA) of the dataset. PCA is a common data characterization technique for identifying outliers. The proportion of variance accounted for by each component is:

0.43 0.23 0.17 0.12 0.05 0.01,

and the cumulative proportions are:

0.43 0.66 0.83 0.94 0.99 1.00

Because the proportion of variance smoothly decreases with the principal component, the data set would not be well represented by a reduction in dimension (as would be expected if the first few components captured nearly all the variation). This suggests that there is little correlation among the attributes of the dataset. Figures 6 and 7 show plots of the first 3 components (which capture 83% of the variation).

The plots show a well-behaved distribution with few outliers. This can make distinguishing between classes hard for an inductive learning algorithm.




Figure 6. Plot of the first three principal components – PC1, PC2, and PC3.

To investigate the interesting question of whether there were patterns resistant to classification and whether they were predominantly of a single class (i.e., UXO or clutter), distributions of pattern classification were tabulated. It was found that many of the incorrect patterns were never classified correctly in any of the 50 computed trials. In addition, the GEFS algorithm produced nearly the same number of misclassifications on the same data points as the ANN. The pattern numbers resistant to classification in both the single ANN and GEFS are given in Table 1. We examined these classification-resistant patterns and found when plotted in the PC1, PC2, and PC3 space, the resistant patterns lie well within the center of well-behaved points. Points at the center of the PCA clusters are the difficult patterns to classify, whereas outliers are typically readily classified.

We note that the classifiers seem to have the greatest problem misidentifying UXO as non-UXO. This is the less desirable of the two types of error since it is better to dig a few extra pieces of scrap than leave dangerous ordnance on site. It is apparent we need to supplement the physics-based parameters with additional information to improve learning accuracy.

4. Incorporating Size and Shape Information from the Mapped Data Images




Figure 7. Plot of PC2 vs PC1 (top), PC3 vs PC1 (center), and PC3 vs PC2 (bottom).

For inductive learning in general, representing the problem meaningfully to the learning algorithm is of the utmost importance. While ANNs and GEFS can theoretically learn complex concepts, in practice only well-represented inputs provide useful results.

Images contain spatial data. The traditional inductive learning approach for utilizing spatial information is to map a local 2-dimensional array of color information. However, images contain much more information. In fact, imagery contains six attributes upon which a human analyst can focus when classifying objects: size, shape, texture, pattern, association, and color.²² The research of Dr. Opitz has investigated using all 6 attributes and has resulted in the state-of-the-art commercial-off-the-shelf (COTS) system for recognizing objects in imagery called Feature Analyst (see www.featureanalyst.com), the cornerstone of this research project.

Feature Analyst works by taking a small and simple set of training examples digitized by the user, learns from the examples, and classifies the remainder of the image. A key component of Feature Analyst is its unique ability to create a number of proprietary descriptors that describe all six attributes (size, shape,

Table 1. Patterns Resistant to Classification (Misclassifications out of 50 Trials).

Pattern Number (Anomaly)	Class	Number of ANN Misclassifications	Number of GEFS Misclassifications
12	UXO	12	35
46	non-UXO	50	49
64	UXO	50	50
120	UXO	50	50
126	UXO	50	50
132	UXO	37	37
134	UXO	45	50
141	UXO	50	50
149	non-UXO	39	43
151	UXO	34	39
156	non-UXO	25	21
166	UXO	50	49

pattern, texture, associate, color) when classifying objects. In addition, Feature Analyst has the ability to improve the accuracy of its results with multiple passes. Figure 8 illustrates the approach that is taken to iteratively reduce clutter and/or retrieve false negatives. The classification improves with iterative passes; each new pass is designed to remove one form of error from the results of the previous pass.

Clutter is the most common form of error when recognizing specific targets in images; it was the focus for this research project. The objective of clutter mitigation is to remove false positives. Thus, the learning task is to distinguish between false positives and correctly identified positives. The user generates a training set by labeling the positive features from the previous classification as either positive or negative. The trained learner then reclassifies the positive instances from the previous pass. The negative instances from the previous pass are considered correct in clutter mitigation and are thus masked out. It is possible to remove too much clutter (i.e. miss legitimate UXO). One would want to provide a path to retrieve false negatives if UXO were missed. There are various approaches that can be tried. In this report, all classifier decisions were made on a yes/no basis. The prediction process is based upon a scaling factor, which can be adjusted to be more or less restrictive. Alternatively, the human analyst could be flagged to QC predictions that are uncertain, or were based upon poor or incomplete signatures.




Figure 8. The hierarchical feature extraction process.

We proceeded in our experiments using spatial information for UXO classification as follows:

- We established a baseline using the physics-based parameters alone on targets falling within the seeded site region having MTADS data.
- We established a second baseline using the raw pixel values (color) without the physics-based parameters. This is the simple and traditional approach and serves as a good baseline to compare the more comprehensive approach provided by Feature Analyst.

- We tested the integration of the physics based parameters with the spatial array of pixel values.
- Finally, we used the descriptors of Feature Analyst that go beyond the baseline pixel-based approach to classify the image.

For the experiment, we used the magnetometer data generated by the vehicular MTADS from the Seeded Site in the BBR to generate a GeoTiff raster having a pixel width of 0.25m. Each pixel in the raster is simply the average magnetic flux for all data points lying within the pixel boundaries. The raster extent was geo-referenced from the Northern-most and Western-most data coordinates.

Baseline 1: Physics Parameters Alone

The data set in the previous section consisted of the 170 dug targets (26 UXO and 144 scrap) associated with the 10-acre seed target area. For the remainder of this report, we reduced this dataset to 139 targets (25 UXO and 114 scrap) because 31 of the original targets were outside the seeded site region having MTADS data. Table 2 presents the results for the 139 targets using only the physics-based parameters. For the ANNs, the error rate of UXO classified as non-UXO is 24% while the error rate of non-UXO classified as UXO is 3.5%. The GEFS algorithm missed one more UXO than did the ANNs, but had an improved false positive rate. These results are similar to the predictions described in Section 3 for the 170 targets, and serve as a baseline for the remainder of the tests in this report.

Baseline 2: Pixel-Based Spatial Representations

Feature Analyst offers numerous pixel-based representation patterns. For this study, we tested several types of input filters, including Bull's-eye filters, Square-pattern filters, and Manhattan-pattern filters. Figure 9 shows sample input representations available within Feature Analyst. With target recognition, mapping input pixels to values for learning is a difficult problem because it is necessary to include spatial information without overwhelming the learner. The baseline approach maps pixel values from a square or circular window of the pixels in the image to the inputs of the learner. To reduce the amount of information, the Bull's-eye patterns take a cue from nature and the use high densities of optic nerves at the center of the field of view and lower densities at the periphery. UXO comprises varying shapes and sizes. Spatial input filters are generally sensitive to the target size in the image; however, one can append multiple filter types in Feature Analyst by using its Retrieve Missed Objects functionality described above.

Table 2. Results using the physics-based parameters alone on the 139 targets included in the MTADS database.

Artificial Neural Networks			
	Predicted UXO	Predicted Non-UXO	Error (%)
Actual UXO	19	6	24.0
Actual non-UXO	4	110	3.5
Error (%)	17.4	5.2	7.2
GEFS			
	Predicted UXO	Predicted Non-UXO	Error (%)
Actual UXO	18	7	28.0
Actual non-UXO	2	112	1.8
Error (%)	10.0	5.9	6.5




Figure 9. Sample Input Representations: Manhattan 5 x 5 (left), Square 7x7 (center), and Bull's Eye with a width of 13 (right). Each square represents a pixel. A filled-in pixel represents a pixel value given to the inductive learning algorithm where the overall pattern is centered on the target UXO.

Feature Analyst automatically rotates patterns during training, thus allowing UXO to be recognized in various orientations. Using the training examples generated by Feature Analyst®, we conducted multiple trial experiments using both ANNs and GEFS. From multiple trial experiments with various training options, we determined mapping a square 7 by 7 pixel region centered at each candidate site produced the best results; however, the variability of results for different input representations was minor. On sites with more diverse or smaller ordnance, either the pixel size or the mapping pattern may have to be adjusted to accommodate the site.

Table 3 contains the results of the experiment. For the analysis employing the pixel-base target signatures alone, the ANN and GEFS both miss 32% of the actual UXO, while the ANN gives a false positive rate of 3.5% and GEFS gives a false positive rate of 2.6%. The ability of the image analysis to correctly identify UXO as UXO is marginally inferior to the physics-based analysis based upon the 170 target dataset. From these values, it seems that the physics-based parameters alone provide more information than the base pixel-based target signatures to achieve a slightly better classification performance using either the ANN or the GEFS method.

Table 3. Results using a Square 7 by 7 window of pixel values.

Artificial Neural Networks			
	Predicted UXO	Predicted Non-UXO	Error (%)
Actual UXO	17	8	32.0
Actual non-UXO	4	110	3.5
Error (%)	19.0	6.8	8.6
GEFS			
	Predicted UXO	Predicted Non-UXO	Error (%)
Actual UXO	17	8	32.0
Actual non-UXO	3	111	2.6
Error (%)	13.0	6.7	7.9

Combined Pixel-Based Target Signatures and Physics-Based Parameters

We next appended the physics-based parameters with the localized array of pixel values. Table 4 gives the results using all physics descriptors added to the Square 7-by-7 input representation. Note the combined information does not significantly improve accuracy. For the ANNs, the combined dataset gives identical error rates obtained using the target signatures alone and somewhat less accurate results than the physics-based parameters alone. The GEFS results for this trial are identical to those obtained with the physics-based rates alone.

Table 4. Results using the physics-based descriptors appended with the square 7 by 7 window of pixel values.

Artificial Neural Networks			
	Predicted UXO	Predicted Non-UXO	Error (%)
Actual UXO	17	8	32.0
Actual non-UXO	4	110	3.5
Error (%)	19.0	6.8	8.6
GEFS			
	Predicted UXO	Predicted Non-UXO	Error (%)
Actual UXO	18	7	28.0
Actual non-UXO	2	112	1.8
Error (%)	10.0	5.9	6.5

5. Using Size and Shape Information from Feature Analyst

In traditional UXO target analysis, an analyst visualizes a localized area of the mapped data file generated from magnetometer data collected by a detection system. Based on the visual attributes of individual localized perturbations (target anomalies) in this image, the analyst determines which targets to select for fitting and computation of the physics-based parameters. Predictions based on these parameters, and the subjective visual attributes associated with each target anomaly, determine which targets are chosen to dig. The remainder of this report investigates the application of Feature Analyst's two-pass approach as an alternative to the traditional processes described above. The first pass is used to select the potential targets of interest (it serves as an automated target picker). The second pass then uses the shape and size descriptors calculated in the Feature Analyst software to distinguish between the UXO and not-UXO targets that survived the first pass. This approach is fully automatic, although it provides the opportunity for an analyst to fine-tune the process.

First-Pass Filtering. In preparation for the first pass analysis, Feature Analyst is used to learn the UXO target signatures based on the size of the anomaly footprint and signal intensity of the positive magnetic flux regions. The positive flux regions of two known UXO were selected as training examples. Based on these two examples, Feature Analyst learned an automated model that filtered the image to 84 candidate target sites. It is interesting to note that this set of 84 sites is a direct subset of the 139 sites provided by the human analyst. Moreover, the 84 sites include *all the known 25 UXO sites*, but contain 56 fewer non-UXO targets (with 59 remaining) than the human analyst selection. While this approach identified all positive UXO targets in the Seeded Site, this selection of first pass parameters may need modification in other fields containing additional types of UXO. The adaptive nature of Feature Analyst allows it to alter its first pass (and second pass) model tuning them to a site's characteristics.

Second-Pass Shape Classification. In the second stage of the two pass analysis; Feature Analyst learns to distinguish between positive and negative examples using the shape of the UXO target signatures. After specifying the known classifications of the targets in the Seeded Site area, the Remove Shapes tool computes the shape descriptor patterns. Using these shape patterns, an analyst can create a model to filter and reclassify the targets selected in the first pass. The Feature Analyst shape descriptors used in this study are: number of vertices, number of holes, perimeter length, area of shape, area of bounding rectangle, compactness, second order moments, rotational invariant measure, scale (size) invariant measure, and the average pixel values in the shape.

Table 5 presents the results of the experiments using the shape descriptors described above on the 84 Seeded Site targets specified in the first pass. The learning algorithms, and in particular GEFS, showed an increase in accuracy. GEFS missed only 3 actual UXO while producing only 3 false positives.

Table 5. Results using the Feature Analyst for pass 1 (auto-picker) and pass 2, the shape classification of UXO.

Artificial Neural Networks			
	Predicted UXO	Predicted Non-UXO	Error (%)
Actual UXO	20	5	20.0
Actual non-UXO	2	112	1.8
Error (%)	9.1	4.3	5.0
GEFS			
	Predicted UXO	Predicted Non-UXO	Error (%)
Actual UXO	22	3	12.0
Actual non-UXO	3	111	2.6
Error (%)	12.0	2.6	4.3

6. Target Classification on the North, South, and West Sites Using the Two-Pass Model

As described above, the two-pass analysis approach using the Feature Analyst was developed and applied to the 10-acre seed target area, and shown to be significantly more effective than the approaches described in earlier Sections. The final exercise in this project was to test and evaluate this most effective approach using survey data from the 100-acre vehicular MTADS survey. For data handling purposes, in the original MTADS analysis, the survey was broken into a South site, 2 West sites (I and II), and 2 North sites (I and II). In preparation for this effort, using Feature Analyst we constructed a model from the complete set of 84 shape patterns generated from the seeded site target signatures.

Because the Pattern Recognition Team had the BBR Demonstration Report,¹ they knew the total number of live UXO recovered from the 100-acre Vehicular Survey Area. That information extracted from Table 9, page 47 of the Demonstration Report¹ is reproduced below. The team did not, however, have the ground truth differentiating individual targets between UXO and clutter until all classification analyses was completed.

Based on the results in Sections 4 and 5, we developed two approaches to apply on the 100-acre survey:

- Using the targets previously selected by the human analyst (Column 2 in Table 6) as the first pass, and the learning model generated from GEFS and Feature Analyst as the second pass.

- (b) Using Feature Analyst for both the first pass (automated target picker) and second pass, as described in detail in Section 5.

Case (a) is equivalent to human analyst filtering, followed by Feature Analyst shape-based classification (i.e., the second pass). The second approach, Case (b) is the fully-automated approach using Feature Analyst for both the first and second pass. The first pass of the fully automated system generated the following numbers of candidate UXO targets: SouthII – 61, WestI – 192, WestII – 10, NorthI – 30, NorthII – 41, which were reclassified during the second pass. The second pass for each approach above is the same – a shape based classification using GEFS and Feature Analyst.

Table 6. Live ordnance recovered on the 100-acre survey area (excluding the Seed Target Area).

Survey Area	Targets Analyzed & Excavated	Live UXO Recovered
South Survey	70	4
West Survey	179	3
North Survey	52	2

Using the 301 targets selected by the human analyst as the first pass, the second pass learner classified the targets as shown in Table 7. When these 301 targets were submitted to the shape filter routine, 24 of them were classified as UXO, the remaining 277 as not-UXO. Using the ground truth data, Table 8 shows the disposition of the 24 targets that were classified as UXO. Six of the 24 targets (highlighted in yellow) were correctly classified as UXO. Eighteen targets (highlighted in white) were false positives. Three targets were misclassified as not-UXO (false negatives). These are highlighted in orange in Table 8.

When the targets were dug on the 100-acre area following the MTADS vehicular survey and analysis, the MTADS dig list contained 301 targets. The MTADS analysis broke the classification into 6 probability categories. The first 3 categories classified 142 of the targets as more likely UXO than not-UXO. This classification captured all of the UXO, with no false negatives. In fairness however, compared to the human analyst performance on the 10-acre seeded site all six MTADS probability categories had to be dug to capture all the UXO targets. The combined human analyst and shape filter cooperative analysis in this project was significantly better at distinguishing not-UXO than was the human analyst alone. This clutter rejection came at the price of missing 3 of the 9 UXO targets, however.

Table 7. Target classification predictions for the 100-acre survey using the human analysis for Pass 1.

Survey Area	Targets Specified by Human Analyst (Pass 1)	Targets Classified As UXO (Pass 2)	Targets Classified As Not-UXO (Pass 2)
South Survey	70	8	62
West (I+II) Surveys	179	13	166
North (I+II) Surveys	52	3	49

Table 8. Targets selected by the human analyst and subsequently classified as UXO using the shape filter.

MTADS Target Analysis								Remediation Results
Target ID	UTM X(m)	UTM Y(m)	Depth (m) Predicted	Size (m)	Fit Quality	MTADS Analyst Comments	UXO Category	Field Comments
N-35	723194.07	4838767.71	0.25	0.126	0.764	possible 155 E/W	2	2 Pieces of Frag 4-in long, 3 in apart
S-178	723123.94	4838030.15	-0.250	0.125	0.838	not likely UXO	5	1.5 lb of Frag
S-204	723351.89	4838091.00	-0.250	0.137	0.966	possible 155	2	Wire & Frag, 1 lb
S-224	723125.84	4838121.22	-0.250	0.125	0.891	unlikely 105/155	3	2 m-50 Fuzes, 10 Pieces of Frag, 4lb
S-240	723372.93	4838188.87	-0.250	0.152	0.907	possible 155	1	Sheet Metal, 5 lb
S-241	723293.08	4838216.84	-0.250	0.152	0.854	155 with remnant?	3	Large Frag, 8 lb
W-7	722644.49	4837597.68	0.19	0.068	0.840	trash	6	Auto body sheet metal, 4 lb
W-64	722812.81	4837639.08	0.18	0.189	0.935	likely 8-in projectile	1	Auto Tie Rod, 6 lbs
W-67	722823.13	4837632.81	0.83	0.108	0.789	likely clutter	5	Frag, 4 lbs
W-89	722583.52	4837685.93	0.62	0.090	0.937	possible 105	3	Wire & Metal Scrap, 2lb
W-98	722601.65	4837703.25	0.42	0.083	0.894	likely not UXO	4	Frag, 10 lb
W-121	722585.09	4837728.51	1.06	0.126	0.873	likely clutter	5	Frag, 2lb
W-123	722536.63	4837760.42	0.26	0.131	0.827	low end of a 155	3	A piece of steel
W-141	722533.65	4837785.82	0.35	0.084	0.953	possible 105	3	frag, 1lb & 3ft of steel banding
W-161	722600.67	4837864.38	0.45	0.093	0.927	possible 105	2	Frag, 12 pieces, 2 lb
W-162	722569.38	4837850.76	0.24	0.135	0.882	possible 155	2	Chrome auto bumper, 2.5ft long
W-50	722370.02	4837631.41	1.63	0.350	0.915	likely not UXO, but too large to leave	3	Hot Dirt
W-84	722720.96	4837678.06	0.94	0.106	0.897	clutter	6	Collection of Frag 1-10 in
N-30	723171.94	4838753.04	0.39	0.177	0.979	dig up a 155	1	155mm HE-Filled, Fuze Sheared Off
N-49	723093.11	4838930.06	0.66	0.238	0.980	Perfect 8-in @ 1 m	1	8-in Projo, HE-Filled, Mech Time Fuze
S-199	723104.72	4838072.68	-0.250	0.156	0.936	possible 155 with clutter	2	8in Projo, HE-Filled, Mech Time Fuze
S-218	723196.12	4838120.70	-0.250	0.142	0.934	likely 155	1	155mm Projo, HE-Filled, Nose 45deg down
S-230	723292.50	4838135.03	-0.250	0.219	0.967	8 in proj deep	1	8in Projo, HE-Filled, Mech Time Fuze
W-80	722806.33	4837668.81	0.95	0.188	0.966	likely 8in proj	1	155mm Projo, HE-Filled
W-91	722432.36	4837674.54	0.63	0.122	0.942	105/155mm	2	155mm Projo, HE-Filled
W-153	722870.63	4837833.02	1.58	0.179	0.909	unlikely, but possible very deep 8-in	3	8in Projo, Lying Flat HE-Filled, Mech Time Fuze,
S-104	722958.73	4838109.38	0.60	0.164	0.947	155 mm	1	155mm Projo, HE-Filled, fuze broken off

Using the fully automated approach, the two-pass cooperative method selected and classified the targets as shown in Table 9.

The automated target picker, which thresholded on the target footprint, selected 334 targets. When these targets were submitted to the shape filter routine, 27 were classified as UXO and 307 as not-UXO. The automated target picker rejected several targets that had been selected by the human analyst. It rejected one target, S-104, a legitimate UXO. Interestingly, this same target, when picked by the human analyst (and classified as a high probability UXO) was classified as not-UXO by the shape filter. Table 10 shows the disposition of the 27 targets classified as UXO in the fully automated two pass approach. This list

Table 9. Target predictions for the 100-acre survey using the automated target picker for Pass 1 and shape filters for Pass 2.

Survey Area	Targets Specified by Automated Picker (Pass 1)	Targets Classified As UXO (Pass 2)	Targets Classified As Not-UXO (Pass 2)
South Survey	61	9	52
West (I+II) Surveys	202	11	191
North (I+II) Surveys	71	7	64

contains the 6 new targets not chosen by the human analyst. Again, there were ultimately the same 6 true positives (highlighted in yellow) and 3 false negatives as appeared in Table 8. The false negatives are highlighted in orange, the false positives in white.

As in the prior case, the cooperative analysis was significantly better at correctly distinguishing not-UXO than was the human analyst, operating alone. The fully automated target picker performed almost as well as the human analyst. This may be unique to this site, however. The relatively unsophisticated automated target picker will likely not perform as well at a site with more complex geology, a wider range of UXO threats, or a much more dense target environment.

Figures 10-13 provide visual representations of the data shown in Tables 8 and 10. These targets are typical of the data used to establish the criteria for the automated target picker and to train the shape classifier. The images are screen captures taken during the target analysis process using the human-interactive MTADS DAS. The false-color presentations are in pixel format, showing data from individual sensor tracks. The cross-track sensor separation is 0.25 m. The down-the-track data points are about 6 cm apart and usually appear continuous on these presentation scales. The color image on the left shows the data; on the right the best dipole fit model to the data is shown. The black and white inset in each figure is a plot of the down-the-track readings from individual sensors. These data are called by clicking the detrend tool, which is available to the analyst. The human analyst uses the fit information shown below the images, the 2-dimensional image presentations, and the various other data (extracted by using the expert analysis tools, e.g. the detrend tool) to make a classification decision, which is then entered into the comment box and becomes part of the target analysis report. The pattern recognition approach developed in this project can either replace the human analyst, providing consistent and accurate results, or minimally be used as an additional and valuable source of information for the analyst.

Figures 10 and 11 show the two extremes of magnetic anomaly types. Figure 10 shows a classical dipole signature from an 8-inch projectile whose long axis is lined roughly parallel to the Earth's magnetic field. The data have high signal-to-noise. Effectively each sensor track shows signal dominated only by components of the dipole signature. There is very little clutter interfering with the fit and the coherence of the model fit is very high.

Figure 11 shows a magnetic anomaly image from a collection of shrapnel pieces that resulted from the detonation of a projectile after it had impacted the earth. The presence of the individual shrapnel pieces is shown in the down-the-track sensor readings in the black and white inset. The color image on the right shows the result of modeling this collection of returns as a single dipole target. The decision that the human analyst must make is "can this collection of returns be due to an intact projectile (or bomb) which also has a lot of clutter (or tail fins) scattered about?" The automatic shape analysis tool must use these 2-dimensional presentations from ground truth to train on and extend this knowledge to the analysis of unknown anomaly signatures. Our technique does this by using a spatial array of pixels and shape descriptors described in Section 6.

Table 10. Targets chosen by the automated target picker and classified as UXO by the shape filter. Targets highlighted in orange are UXO that were missed by the picker or misclassified by the shape filter.

MTADS Target Analysis								Remediation Results
Target ID	UTM X(m)	UTM Y(m)	Depth (m) Predicted	Size (m)	Fit Quality	MTADS Analyst Comments	UXO Category	Field Comments
N-2	723091.30	4838616.02	0.74	0.111	0.897	possible 105 mm	3	
N-?	723085.92	4839029.63				Not an MTADS Pick		
N-?	723220.29	4838968.51				Not an MTADS Pick		
N-35	723194.07	4838767.71	0.25	0.126	0.764	possible 155 E/W	2	2 Pieces of Frag 4-in long, 3 in apart
N-45	723279.99	4838851.48	0.78	0.115	0.882	low probability 105	3	10 Pieces of Frag, 6-12 in.
S-178	723123.94	4838030.15	-0.250	0.125	0.838	not likely UXO	5	1.5 lb of Frag
S-204	723351.89	4838091.00	-0.250	0.137	0.966	possible 155	2	Wire & Frag, 1 lb
S-224	723125.84	4838121.22	-0.250	0.125	0.891	unlikely 105/155	3	2 m-50 Fuzes, 10 Pieces of Frag, 4lb
S-240	723372.93	4838188.87	-0.250	0.152	0.907	possible 155	1	Sheet Metal, 5 lb
S-?	723357.48	4838201.78				Not an MTADS Pick		
S-?	723361.13	4838188.86				Not an MTADS Pick		
W-44	722514.91	4837602.77	0.65	0.105	0.953	possible 105	2	Frag, 3 lbs
W-121	722585.09	4837728.51	1.06	0.126	0.873	likely clutter	5	Frag, 2lb
W-123	722536.63	4837760.42	0.26	0.131	0.827	low end of a 155	3	A piece of steel
W-126	722611.71	4837749.13	0.74	0.121	0.977	possible 155	1	Frag Pile, 2.5 lb
W-127	722642.24	4837769.75	0.38	0.144	0.899	155mm E/W	1	Frag, 4lb & barb wire
W-162	722569.38	4837850.76	0.24	0.135	0.882	possible 155	2	Chrome auto bumper, 2.5ft long
W-50	722370.02	4837631.41	1.63	0.350	0.915	likely not UXO, but too large to leave	3	Hot Dirt
W-84	722720.96	4837678.06	0.94	0.106	0.897	clutter	6	Collection of Frag. 1-10 in
W-?	722464.35	4837905.72				Not an MTADS Pick		
W-?	722568.98	4837697.21				Not an MTADS Pick		
N-30	723171.94	4838753.04	0.39	0.177	0.979	dig up a 155	1	155mm HE-Filled, Fuze Sheared Off
N-49	723093.11	4838930.06	0.66	0.238	0.980	Perfect 8-in @ 1 m	1	8-in Projo, HE-Filled, Mech Time Fuze
S-199	723104.72	4838072.68	-0.250	0.156	0.936	possible 155 with clutter	2	8in Projo, HE-Filled, Mech Time Fuze
S-218	723196.12	4838120.70	-0.250	0.142	0.934	likely 155	1	155mm Projo, HE-Filled, Nose 45deg down
S-230	723292.50	4838135.03	-0.250	0.219	0.967	8 in proj deep	1	8in Projo, HE-Filled, Mech Time Fuze
W-80	722806.33	4837668.81	0.95	0.188	0.966	likely 8in proj	1	155mm Projo, HE-Filled
W-91	722432.36	4837674.54	0.63	0.122	0.942	105/155mm	2	155mm Projo, HE-Filled
W-153	722870.63	4837833.02	1.58	0.179	0.909	unlikely, but possible very deep 8-in	3	8in Projo, Lying Flat HE-Filled, Mech Time Fuze,
S-104	722958.73	4838109.38	0.60	0.164	0.974	155 mm	1	155mm Projo, HE-Filled, fuze broken off




Figure 10 Screen capture of the MTADS analysis of target N-49.




Figure 11. Screen capture of the MTADS analysis of target W-67.




Figure 12. Screen capture image of the MTADS analysis of target W-153.




Figure 13. Screen capture image of the MTADS analysis of target S-104.

Figures 12 and 13 show two of the three false negative classified targets in the same presentation format shown in Figures 10 and 11. Target W-153 lies very close to the bull's eye. Therefore, it has a lot of clutter interference from shrapnel and auto body parts. The bull's eye target was a stack of automobiles that had been reduced to scrap metal by projectile detonations. This 8-inch dud projectile was buried relatively deeply; therefore the signal intensity was much smaller than that of target N-49. The human analyst, knowing he was working very near the bull's eye, made allowances for the generalized clutter and designated the target to be dug as UXO. The shape function utility compared the image to the typical clutter training image and logically classified the anomaly as not-UXO.

The situation depicted in Figure 13 is less clear cut. Target S-104 is a 155-mm projectile, buried about 2 feet deep. It was not selected by the automated target picker and therefore, was not submitted for classification in Table 9. In Table 8, the human analyst chose the target for analysis and it was therefore, submitted to the classifier. The shape classifier, however, classified the target as not-UXO. The projectile is characterized by a relatively high signal-to-noise ratio. It probably is oriented northeast/southwest (or vice versa) giving almost the maximum skew to the dipole orientation away from magnetic north. The anomaly is characterized by some, but not very much, adjacent clutter. This anomaly points out the necessity for creating a good set of examples upon which to train, and exposes the need for follow-on research.

7. Conclusions

We draw the following conclusions from this study:

- a. A human analyst, working with a physics-based modeling analysis system and high resolution mapped data files, can achieve excellent detection results at relative uncomplicated sites like the BBR. However, the analyst's ability to distinguish between intact ordnance and clutter still requires digging 5-20 targets for each recovered UXO.
- b. We have shown, by exhaustive studies in this project, that the physics parameters generated by a human-in-the-loop analysis approach, such as at the BBR survey, cannot be significantly improved by the use of inductive learning algorithms that focus on the physics-based fitting parameters. In other situations, other researchers working with different datasets from more complex sites have drawn different conclusions. The generality of our observations may depend upon both the complexity of the survey site and the quality of the datasets and the human-assisted analysis.
- c. In the early parts of this study, the use of raw spatial pixel values from the MTADS data when appended in a cooperative analysis with the physics-based predictions, did not improve classification accuracy.
- d. In our final study, the first-pass technique that we developed provides a quick and effective automated target picker. With these datasets this approach can effectively replace the man in the loop for target selection, or at a minimum, can provide a valuable aid to the analyst. This automated technique can be adapted to different fields and situations.
- e. Our second-pass technique provides the ability to use shape and intensity information from the MTADS data to improve the automated discrimination of UXO and other metallic scrap. As a result, our new techniques hold great promise in reducing the numbers of non-UXO that must be dug.

In general, the results of the cooperative analysis of the BBR data using shape filters as a classification tool suggest some tantalizing possible improvements in classification. Using these pattern recognition tools in conjunction with either a primitive automatic target picker, or the physics-based analysis as a pre-

screeners, reduced the dig list by more than a factor of two. This significant advance in correctly identifying clutter targets came at the expense of some missed UXO targets.

Because of the limited resources available for this study, the scope of our efforts had to be restricted. There are several areas where significant improvements and evaluations could be made. There was neither time nor resources available to implement a probabilistic evaluator in the shape function analysis. To critically evaluate the performance of our new tools it will be necessary to develop this capability. It will allow the generation of ROC curves and the evaluation of this performance approach against others. Equally importantly, it would allow classification thresholds to be adjusted to emphasize either UXO detection or clutter rejection.

Our discussions in the previous section have illustrated the importance of extensive ground-truthed data sets to achieve adequate training. It is important that the training sets have both extensive UXO and clutter targets. Moreover, it is important that these training data be associated with the actual site to be analyzed. Indigenous clutter and geological interferences are critical contributors to noise. Noise sources ultimately establish the detection and classification floor because of signal-to-noise limitations.

The BBR site was chosen for this study because of its simplicity. From an ordnance perspective, it is almost a single use site. Background clutter (except near the bull's eye) and geological interferences were low, and generally, UXO targets were well separated and isolated from each other. Finally, this is the most extensive area we are aware of that has ever been studied with high quality surveys and 100% remediated.

These shape function filters, combined with a cooperative analysis approach need to be further developed and evaluated at other ranges against a more complex mix of UXO and background challenges.

8. References

1. "Airborne MTADS Demonstration on the Impact Area of the Badlands Bombing Range, September 2001," J.R. McDonald, David Wright, Nagi Khadr and H.H. Nelson, NRL/PU/6110--02-453, May 2002
2. "MTADS Unexploded Ordnance Operations at the Badlands Bombing Range, Air Force Retained Area, Pine Ridge Reservation, SD, September, 1999," J.R. McDonald, H.H. Nelson, R. Robertson, and R.A. Jeffries, NRL/PU/6110--00-424
3. "Signal Processing and Modeling for UXO Detection and Discrimination in Highly Contaminated Sites," Leslie Collins, SERDP Project 1281, <http://www.serdp.org/research/UX/UX-1281.pdf>
4. "Feature Based UXO Detection and Discrimination," Dean Keiswiter, SERDP Project UX-0210, <http://www.serdp.org/research/UX/UX-0210.pdf>
5. "Bayesian Approach to UXO Site Characterization with Incorporation of Geophysical Information," Sean McKenna, SERDP Project UX-1200, <http://www.serdp.org/research/UX/UX-1200.pdf>
6. "Spatial Statistical Models and Optimal Survey Design for Rapid Geophysical Characterization of UXO Sites," William E. Doll, SERDP Project UX-1201, <http://www.serdp.org/research/UX/UX-1201.pdf>

7. **An Introduction to Genetic Algorithms**, M. Mitchell, MIT Press, Cambridge, MA, 1996
8. **Programs for Machine Learning**, C4.5, J. Quinlan, Morgan Kaufmann, San Mateo, CA, 1993
9. "Learning Internal Representations by Error Propagation," Rumelhart et al., In: **Distributed Processing: Explorations in the Microstructure of Cognition**, Vol. 1, 1986, pp. 318-363
10. "Learning to Recognize Volcanoes," Burl, et al., *J. Machine Learning*, **30** 165-194 (1998)
11. "Artificial Neural Network Classification Using a Minimal Training Set: Comparison to Conventional Supervised Classification," Hepner, et al., *Photogrammetric Eng. and Remote Sensing* **56**, 469-473 (1990)
12. "Top Ten Lessons Learned in Automated Cartography," D. McKeown, *Technical Report CMU-CS-96-110*, Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, 1996
13. "Bagging Predictors," L. Breiman, *J. Machine Learning*, **24**, 123-140 (1996)
14. "Popular Ensemble Methods: An Empirical Study," D. Opitz and R. Maclin, *J. Artificial Intelligence Research*, **11**, 169-198 (1999)
15. "Actively Searching for an Effective Neural-Network Ensemble," D. Opitz and J. Shavlik, **Perspectives in Neural Computing**, pp. 79-97, Springer Verlag, 1999
16. Shapire, et al., *Proc. of the 14th Int. Conference on Machine Learning* pp. 322-330, 1997
17. "Actively Searching for an Effective Neural-Network Ensemble," D. Opitz and J. Shavlik, *J. Connection Science*, **8**, 337-353 (1996)
18. "Feature Selection for Ensembles," D. Opitz, *Proc. of the 16th National Conference on Artificial Intelligence*, pp. 379-384, 1999
19. "Connectionist Theory Refinement: Genetically Searching the Space of Network Topologies," D. Opitz and J. Shavlik, *Journal of Artificial Intelligence Research*, **6**, 177-209 (1997)
20. "Hazard Assessment Modeling: An Evolutionary Ensemble Approach," D. Opitz, et al., *Genetic and Evolutionary Computation* **2**, 1643-1651 (1999)
21. "Neural Network Ensemble for Control," D. Opitz, et al., *Proceedings of the International Conference on Control and Applications*, pp. 464-467, 1999
22. "How to use Aerial Photographs for Natural Resource Applications," J. Caylor, USDA Forest Service, Nationwide Forestry Applications Program. p. 328, 1988