Optimizing Supervised Learning for Pixel Labeling and Material Classification

Mark Rahmes, Morris Akbari, Mike McGonagle
Harris Corporation, Government Communications Systems
Melbourne, Florida
mrahmes@harris.com, makbari@harris.com, mmcgonag@harris.com

ABSTRACT

The visualization and simulation industry has a demonstrated interest in classification products for sensor simulation. The challenge lies in providing highly accurate material classification of remotely sensed imagery while significantly reducing the time and cost to create products. Visualization and simulation products for material classification are created by merging and mosaicking multi-source satellite and aerial imagery of different resolutions on an elevation surface to provide realistic, geo-specific terrain features. This requires that all image data is orthorectified, seamlessly co-registered, tonally balanced and feather blended into mosaics from source data of different resolution. To achieve highest accuracy at faster speed and lower cost, we apply an innovative, optimal pixel-labeling process to the mosaic imagery. This process is based on artificial intelligence (AI) algorithms using Nash Equilibrium and game theoretic analyses to help solve the problem of feature extraction through supervised classification. This can be viewed as a constant sum game, whereby the players are pixel data points that take part in the game to decide their class memberships. A player's land cover classification strategies are based on four different supervised learning algorithms: k-Nearest Neighbors (KNN), Decision trees using a classification and regression tree (CART), Normal/Naïve Bayes probabilistic graphical model, and support vector machine (SVM). Within this formulation, we used a weighted reward matrix for consistent labeling of feature pixels and classification factors, resulting in higher accuracy and precision when compared to the individual machine learning algorithms alone.

ABOUT THE AUTHORS

Dr. Mark Rahmes has 22 years of experience at Harris Corporation as an Electrical/Computer Engineer and Senior Research Scientist. He earned his BSEE from The Citadel, MSEE from Duke University and PhD in Operations Research from Florida Tech. Dr. Rahmes is a retired U.S. Navy Reserve Captain and served 22 years as a Commanding Officer, Engineering Duty Officer and Surface Warfare Officer. He currently has 40 patents granted, and has published 35 IEEE, ASPRS, ILMF, and SPIE conference papers. At Harris Corporation, Dr. Rahmes serves in the capacity of a Principal Investigator and Chief Software Engineer on various programs while advancing image processing research and development. He is a member of Tau Beta Pi and Phi Kappa Phi Honor Societies.

Morris Akbari has 22 years of technical and managerial experience in Remote Sensing, Image Processing and GIS. He earned his BS and MS degrees in Engineering from the University of Tennessee. Mr. Akbari currently has 8 patents in image processing, and has published 7 ASPRS, ILMF, and I/ITSEC conference papers. He serves as Senior Geospatial Engineering Manager for the ImageLinks program within Harris Corporation. Mr. Akbari is involved with development of advanced image processing and remote sensing algorithms as well as new product development. He is a member of ASPRS.

Mike McGonagle has 14 years of experience at Harris Corporation as a Software Engineer. He earned his BS in Mathematics from Rose-Hulman Institute of Technology and MS in Mathematics from Purdue University. Mr. McGonagle's work experience and interests include computational mathematics, scientific computing, image processing, computer vision and machine learning.

Optimizing Supervised Learning for Pixel Labeling and Material Classification

Mark Rahmes, Morris Akbari, Mike McGonagle
Harris Corporation, Government Communications Systems
Melbourne, Florida
mrahmes@harris.com, makbari@harris.com, mmcgonag@harris.com

INTRODUCTION

Remote sensing requires that image analysts have the capability to identify regions in imagery that correspond to a particular object or material. The automatic extraction of image areas that represent a feature of interest requires two specific steps: The first is to accurately classify the pixels that represent the region while minimizing misclassified pixels; the second is a vectorization step that extracts a contiguous boundary along each classified region which, when paired with its geo-location, can be inserted in a feature database independent of the image [3].

The amount of available high-resolution satellite imagery, and the increasing rate at which it is acquired, simultaneously present interesting opportunities and difficult challenges for the simulation and visualization industry. Updating material classification product databases frequently using high-resolution panchromatic and multispectral imagery is only feasible if the time and labor costs for extracting features, such as pixel labeling, and producing products from the imagery, are significantly reduced. Our solution is designed to provide flexible and extensible automated workflows for land use land cover (LULC) pixel labeling and material classification. The products of the workflows undergo an accelerated review and quality control process for feature extraction accuracy by geospatial analysts [11].

The ultimate goal is to greatly reduce the quantity of data an analyst must review, while maintaining the high quality of the resulting products. The data reduction is achieved through batch processing the area of interest (AOI) to identify those feature classes in which analysts are interested. Our solution uses game theory to extract pixel labels, provide tools for analyst review and post processing, and produce inputs to our material classification process. Batch processing is initiated by the process workflow manager specifying the input AOI imagery, processing parameters and the output products desired. The process is comprised of components as shown below in Figure 1.

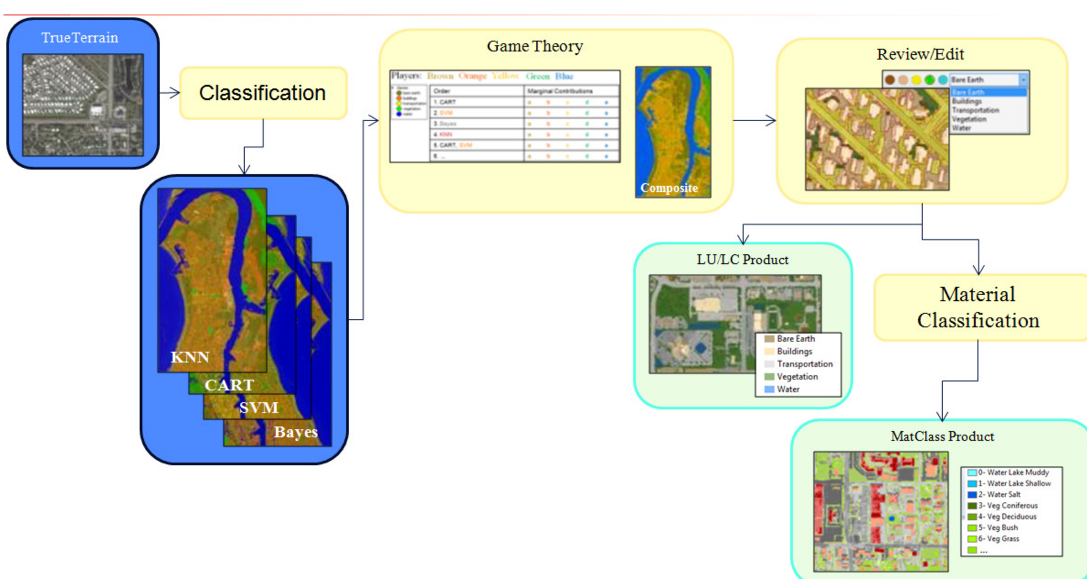

Figure 1. Pixel Labeling System Overview

Statistical pattern recognition requires a statistical relationship between features and class membership of a pattern. It consists of three steps: feature selection and feature extraction; selection of a distribution/density function and estimation of its parameters; and computation and test of a decision boundary. The distribution/density function may be selected based on understanding how features vary given the imaging process. Parameter estimation is based on samples of a training data set. The decision boundary consists of those locations in feature space where the class changes according to the computed maximum probability. The location of the decision boundary may include the cost of error. Testing the quality of the classifier may be dependent or independent of the selected distribution/density function [15]. In our solution, feature selection and extraction is accomplished through supervised classification; estimation of parameters is accomplished through training; and our decision boundary accuracies are measured via Receiver Operating Characteristic (ROC) Curves.

A large body of research in supervised learning deals with the analysis of single label data, where training examples are associated with a single label from a set of disjoint labels. However, training examples in several application domains are often associated with a set of labels. Such data are called multi-label. The categorization of textual data, such as documents and web pages, is perhaps the dominant multi-label application. Recently, the idea of learning from multi-label data has attracted significant attention from many researchers, motivated by an increasing number of new applications, such as semantic annotation of images and video, functional genomics, music categorization into emotions and directed marketing [16].

This paper discusses results of a multi-label application where our algorithms make a cognitive decision as to the best land use classification per pixel. It addresses Material Classification, the current methods for achieving it, and the more affordable and efficient solution we propose.

MATERIAL CLASSIFICATION

Material classification is the semantic assignment, or labeling, of a color or multi-spectral image pixel to an index representing a material or group of materials making up a material mixture. The purpose of the assignment is to provide additional information – beyond the spectral characteristics of the pixel – to aid in the development of correlated sensor simulations and geo-specific content generation. Traditional material classification within the supervised learning process can pose some challenges:

Limited Training Samples – Remote sensing imagery is rich with information on spectral and spatial distributions of distinct surface materials. Owing to its numerous and continuous spectral bands, hyperspectral data enables even more accurate and reliable material classification than panchromatic or multispectral imagery. However, high-dimensional spectral features, and the limited number of available training samples for supervised learning, have caused difficulties in material classification, such as overfitting in learning, noise sensitiveness, overloaded computation, and lack of meaningful physical interpretability [20].

Thousands of Variables – The task is made even more challenging by the fact that the number of spectral channels available for the detailed analysis of the materials is very large. The dimensionality of hyperspectral data may range from dozens to thousands of variables (spectral bands), which can prevent the successful application of standard pattern recognition techniques – especially in small sample size situations. This is known as the curse of dimensionality. To avoid these adverse effects on many learning systems, it is common to apply as a preprocessing step, one of the many existing feature extraction (FE) or dimensionality reduction (DR) techniques [17].

Finding the Discriminator – Feature extraction methods are also employed to establish more concentrated features for separating different materials, as not every spectral band contributes to material identification. Among them, discriminative feature extraction methods learn a suitable subspace where one can expect the separability between the different classes to be enhanced. Typical methods widely used for hyperspectral imagery include linear discriminant analysis and nonparametric weighted feature extraction, which design proper scatter matrices to effectively measure the class separability [13].

Object material identification in spectral imaging combines the use of invariant spectral absorption features and statistical machine learning techniques. The relevance of spectral absorption features for material identification casts the problem into a pattern recognition setting by making use of an invariant representation of the most discriminant band-segments in the spectra. The identification problem is a classification task, which is effected based upon those invariant absorption segments in the spectra that are most discriminative between the materials. To robustly recover those bands that are most relevant to the identification process, discriminant learning is used [5].

Integration of geometrical features, such as the characteristic scales of structures, with spectral features is used for the classification of hyperspectral images. The spectral features, which only describe the material of structures, cannot distinguish objects made by the same material but with different semantic meanings (such as the roofs of some buildings and the roads). The use of geometrical features is therefore necessary. Moreover, since the dimension of a hyperspectral image is usually very high, a linear unmixing algorithm is used to extract the end members and their abundance maps in order to represent compactly the spectral information [8].

Enhancement of commercial satellite imagery is accomplished by merging and mosaicking multi-source satellite and aerial imagery of different resolutions on an elevation surface to provide realistic geo-specific terrain features. This requires that all data is orthorectified, seamlessly co-registered, tonally balanced, pan-sharpened and feather blended mosaics created from different resolution source data.

We use this pan-sharpened product as opposed to original multispectral imagery to perform classification, as the pan-sharpened product has higher fidelity. We determine the two dominant materials, as well as the relative abundance of each material, for each pixel in the data set. Available at the same pixel resolutions and precisely correlated to the true color product, the material classification data set is ideal for creating various sensor views to accompany out-the-window views within the Simulation Image Generator. Material classification products can be used to create night vision, IR, and radar visual databases or for mapping high detail, geotypical textures with real-world accuracy. Output is typically made available in Geotiff format. While current technology offers robust and proven classification capabilities, there are some applications that will benefit from our new approach to imagery feature extraction.

SUPERVISED CLASSIFICATION

Supervised classification techniques play a key role in the analysis of hyperspectral images, and a wide variety of applications can be handled by successful classifiers, including: land-use and land-cover (LULC) mapping, crop monitoring, forest applications, urban development, mapping, tracking and risk management. Conventional classifiers treat hyperspectral images as a list of spectral measurements. Classifiers use both spectral and spatial information. In addition, to reduce the redundancy of features and address the so-called curse of dimensionality, different supervised feature extraction (FE) techniques are considered. A possible way to improve the extraction of spatial information is to use different types of segmentation methods. Image segmentation is a procedure that can be used to modify the accuracy of classification maps [6].

Supervised Classification as a Learning Technique – To get the semantic assignment or labeling of a color or multi-spectral image pixel for material classification, we use supervised learning methods. Supervised classification requires a training set representative of the real world information to "learn" about the information to properly classify and predict the selected input objects or feature representations. In addition to the limited number of training samples available, mentioned earlier, the size and characteristics of the training set can have a noticeable effect on the results in both accuracy and precision. Other factors to consider for the training set are heterogeneity of data, redundancy in the data, and presence of interactions and non-linearity.

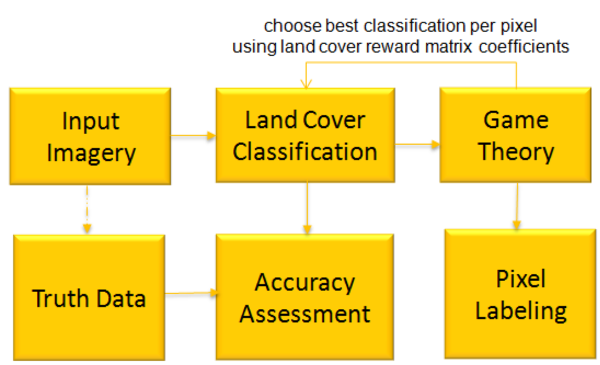

Figure 2. Supervised Land Cover Classification System

In our application, supervised classification is a supervised learning technique generated from examples selected from multispectral imagery, saved and communicated by means of a training set in the form of a shapefile. This training consists of a portion of truth from image data. Our system diagram is shown in Figure 2.

In order to assign a classification of features over an image, we apply supervised learning to the imagery. Supervised learning creates a classifier model that can infer the classification of a test sample using knowledge acquired from labeled training examples. In our case, the trained classifier predicts if a small area of an image is a particular feature or not, and we do this over the whole test image. Each small image area is turned into a feature vector, and it is this vector that is passed to the classifier. To train the classifier, the image areas are manually labeled with a feature type and turned into feature vectors. The feature vector and label pairs are inputs to a machine-learning algorithm that produces a classifier model.

To achieve the best results, we are looking for models with a good bias/variance trade-off leading to a higher generalization, which allows us to apply a trained model to a wider variety of imagery. For example, if a model were trained over desert imagery yet also works well when testing over forested areas, then the model generalized well. There are many choices for machine-learning algorithms for training of data sets. In our implementation, we used the available algorithms from the machine learning library (MLL) provided by the Open Computer Vision C++ library (OpenCV). The four algorithms considered are: k-Nearest Neighbors (KNN), Decision trees using a classification and regression tree (CART), Normal/Naïve Bayes probabilistic graphical model (PGM), and support vector machine (SVM) [18].

- KNN is the most simplistic algorithm, and simply looks at the k points (k is a chosen odd integer) in the training set that is closest in feature space distance to the test sample. KNN selects the feature class for the test based on the class label of the majority of the k closest training points.
- CART uses the training data to create a tree, where each leaf node has a class label determined by the class label of the majority of training examples reaching that leaf. The internal nodes of the tree are questions based on the feature vectors; it branches based on the answers. When a test vector is applied to the tree, the vector obtains the label of the leaf it reaches.
- The Bayes PGM treats test feature vectors as probabilistic evidence and infers the hidden classification state. Bayes PGM algorithm is naïve because it makes the assumption that the evidence variables are independent, even though they frequently are not. The training data allows the PGM to learn the weights on the graph edges that maximize our expectation of correct inference.
- The most sophisticated algorithm, SVM (with linear kernel), locates a linear separator of the training data with maximum margin. Training points that lie on the margin are considered as the support vectors. A simple linear calculation of a test vector with the support vector solution will generate a positive or negative value that indicates a feature classification of the test.

How well each model works is dependent upon feature properties, the quality and quantity of training data, and the parameter settings for the individual algorithms. Extensive validation of the results will need to be considered in order to properly select the optimal model and model parameters for a given problem. If the labeled training data is distributed very non-linearly, then a linear learning method will be unlikely to fit the data well, resulting in a high bias, but may be generalized to some degree. If the training data is linearly separable and we use a highly non-linear based learning algorithm, then we will likely over-fit the data, suffer from high variance, and not be able to generalize well with the resulting output. If we have too little training data, or the data is not a representative sample of the feature space, then accuracy and precision will be negatively affected. However, if we are suffering from a high bias, additional training data can actually make a model fit worse. We have tested each model on a few different images and geographic areas to understand how well each might work in practice.

CREATING A TRAINING SET

We experimented with several methods to generate a viable training set, with an eye toward developing a repeatable procedure. We utilized high-resolution satellite multispectral imagery for our imagery. The imagery training set consisted of points split approximately evenly among feature classes. The points consisted of a centroid from every feature in the truth set, which are polygonal shape files, and a set of random points from the non-feature polygons. The truth set could be seeded from the existing database.

We independently applied four supervised learning algorithms discussed above to several study areas. We refer to this as a local training set, meaning that the training samples are drawn from, and applied to, the same image as shown in Figure 3. The concept of operations is to train on a small area and test on the larger surrounding area. Minimizing the number of false negatives (misses) is our first priority, and keeping the number of false positives down is the second priority. Extraction feature results were evaluated by comparing the extracted features with a truth set, which consisted of a feature shapefile created using automated methods and modified manually to meet established extraction requirements. A point grid shapefile was created within a regular 20 meter grid. This grid was then reduced to include only those points in the grid that lay with the union of the truth and the extracted features. The trimmed shapefile is modified to show relationships with the autonomously extracted features and truth set.

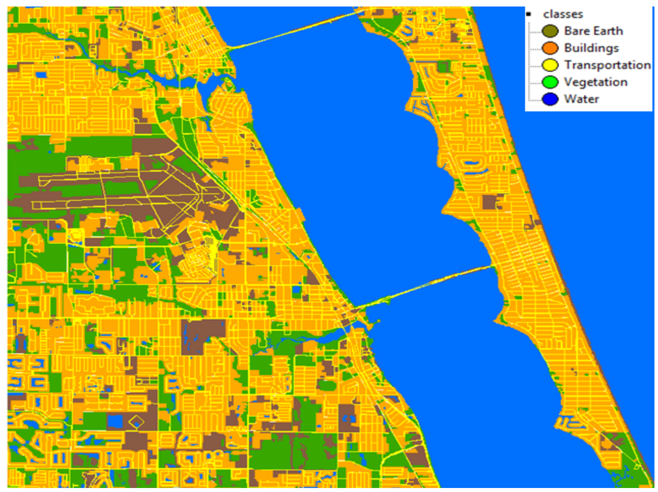

Figure 3. Truth Data for Land Use/Land Cover Features

Although collection of truth is somewhat subjective, these land cover features were collected via typical specifications. In areas where only larger buildings exist (for example, an industrial complex), building areas were not collected. In this case, individual buildings would have been collected due to their size and significance, and an industrial complex would have been added. It is important to remember that there is a lot of grey area when determining what should be collected and what should just stay individual buildings. Different analysts can and will collect the same area differently with both views considered as acceptable.

There is work being done that could help to better close the gap between the difference in feature extraction by analysts. To make informed decisions, an expert has to reason with multi-dimensional, heterogeneous data and analyze the results. Items in such datasets are typically represented by features. However, as argued in cognitive science, features do not yield an optimal space for human reasoning. In fact, humans tend to organize complex information in terms of prototypes or known cases rather than in absolute terms. When confronted with unknown data items, humans assess them in terms of similarity to these prototypical elements. Interestingly, an analogous, similarity-to-prototype approach, where prototypes are taken from the data, has been successfully applied in machine learning. Combining such a machine learning approach with human prototypical reasoning in a Visual Analytics context requires integration of similarity-based classification with interactive visualizations. To that end, the data prototypes should be visually represented to trigger direct associations to cases familiar to the domain experts. Highly interactive visualizations are used to explore data and classication results in terms of dissimilarities to visually represented prototypes. This approach not only supports human reasoning processes, but is also suitable to enhance understanding of heterogeneous data [9].

A NEW APPROACH: ENHANCED PIXEL LABELING THROUGH GAME THEORY

Automated artificial intelligence is especially desirable in commercial applications and applications of machine learning in fields beyond its traditional boundaries, such as medicine, marketing, etc. An effective and efficient general explanation method would also be a useful tool for comparing how a model predicts different instances. How different models predict the same instance feature value's contribution is defined as the difference between the model's initial prediction, and a model's average prediction across perturbations of the corresponding feature. In other words, one can look at how the prediction would change if we ignore the feature value. This myopic approach can lead to serious errors if the feature values are conditionally dependent, which is especially evident when a disjunctive concept (or any other form of redundancy) is present. Rigorous theoretical analysis and explanation methods have been performed and link with known concepts from game theory, thus formalizing some of its desirable properties [14].

We have developed a novel, modified game-theoretic perspective to solving the problem of supervised classification that takes the best pixel label derived from four supervised classifications. This is a game whereby the "players" are pixel data points that participate in the game to decide their class memberships. The land cover classification strategies are based on several different supervised learning algorithms. Within this formulation, we use a weighted reward matrix for consistent labeling of feature pixels and classification factors resulting in higher accuracy and precision.

The purpose of classification algorithms is to label pixels that have spectral signatures that fall within a distribution defining a region of interest. A pixel belongs to a classification set when the distance, in feature space, between the pixel's spectral signature and the signature of a representative set of pixels is small. Classification algorithms vary in how the feature vector (and therefore feature space) is defined, how the distance metric is defined, how a representative set of pixels or distribution is determined, and in the algorithm by which pixels matches are identified. Nevertheless, they all share the concept of goodness-of-fit, a per-pixel score measuring how well a pixel actually fits the target spectral distribution. The goal is to accurately identify the boundary of a spatially consistent set of pixels that belong to a region of interest, with the intent of extracting that region as a distinct feature [3].

The Players – The players of the game can be categorized into two distinct groups: labeled players – those that already have knowledge of their membership; and unlabeled players – those that do not have this knowledge as the game begins. [4]. The so-called labeled players of the game can be further distinguished based upon the strategies they follow without hesitation, coming from their membership information. In this regard, it can be argued that the labeled players do not play the game to maximize their payoffs since they have already chosen their strategies. In fact, the transduction game can easily be reduced to a game with only unlabeled players where the definite strategies of labeled players act as bias over choices of unlabeled players [4].

The Reward Matrix – We focus on creating a reward matrix of five rows and four columns per pixel. The five rows are the land cover classifications as shown in Table 1. The four columns are the supervised learning algorithm confidence values. Other simulations have accounted for games involving more than two players [12]. Some have used a relaxation algorithm based on the game theory for scene labeling problems. Labeling pixels that maximize the a posteriori probability are Nash equilibrium points of the game, and that all the Nash equilibrium points are local maxima that have been proven. Relaxation algorithms exist showing efficiency and rapid convergence [19].

Table 1. Reward Matrix per Pixel

	Bayes	**CART**	**KNN**	**SVM**
Bare Earth	a11	a12	a13	a14
Buildings	b21	b22	b23	b24
Transportation	c31	c32	c33	c34
Vegetation	d41	d42	d43	d44
Water	e51	e52	e53	e54

Pixels as Players – One method for feature extraction is derived from a special branch of the theory, known as non-cooperative n-person game theory. The basic idea is to consider the pixels as the players and the labels as the strategies. The maximization of the payoff function of the game corresponds to the maximization of the a posteriori probability of the labeling. Here, we are interested in using game theory to find the maximum payoff. The game should be a cooperative one: all the players have to take part in a coalition in order to find the global maximum. A play is a Nash equilibrium if none of the players can improve his or her expected payoff by unilaterally changing his strategy. It is known that Nash points always exist in n-person games with pure or mixed strategies [2]. It is important to note that any instance of the proposed graph transduction game will always have Nash equilibrium in mixed strategies [10].

Game Theory Applied to Feature Extraction – Feature Extraction is accomplished by evaluating the Nash equilibrium value. The equilibrium value serves as the metric to determine the likelihood of feature accuracy for manual evaluation to update a given database. A stronger Nash equilibrium value would prompt an analyst to manually inspect the data area for updating a geospatial database. Feature extraction, especially in diverse urban environments, often suffers from feature misclassification. Improving algorithms for feature extraction is an ongoing process. Three topics for future work include: study how to optimize training sets; apply remote rules to a similar landscape, such as a neighboring (rather than distant and different place); and modification of supervised classification rules to generate a value for ROC curve, to use for selecting the threshold. Game theory can cope with both asymmetric and/or negative similarities. The classification task is to estimate the pixel classification of the unlabeled pixels based exclusively on the information available for the possible five features [2].

Linear programming is useful for solving game theory problems and finding optimal strategies. We can define:
x_1= probability that the blue player chooses bare earth
x_2= probability that the blue player chooses buildings
x_3= probability that the blue player chooses transportation
x_4= probability that the blue player chooses vegetation
x_5= probability that the blue player chooses water
As an example using the reward matrix in Table 1, we show the linear programming solution for the constant sum game as follows:

$$
\begin{aligned}
&\max v\\
&\text{s.t.}\\
&v - a_{11}x1 - b_{21}\,x2 - c_{31}x3 - d_{41}x4 - e_{51}x5 \le 0\\
&v - a_{12}x1 - b_{22}\,x2 - c_{32}x3 - d_{42}x4 - e_{52}x5 \le 0\\
&v - a_{13}x1 - b_{23}\,x2 - c_{33}x3 - d_{43}x4 - e_{53}x5 \le 0\\
&v - a_{14}x1 - b_{24}\,x2 - c_{34}x3 - d_{44}x4 - e_{54}x5 \le 0\\
&x1 + x2 + x3 + x4 + x5 = 1\\
&x1, x2, x3, x4, x5 \ge 0
\end{aligned}
$$

The initial solution for optimal player's mixed strategy in terms of probabilities: **x** = (x1,x2, x3, x4, x5).

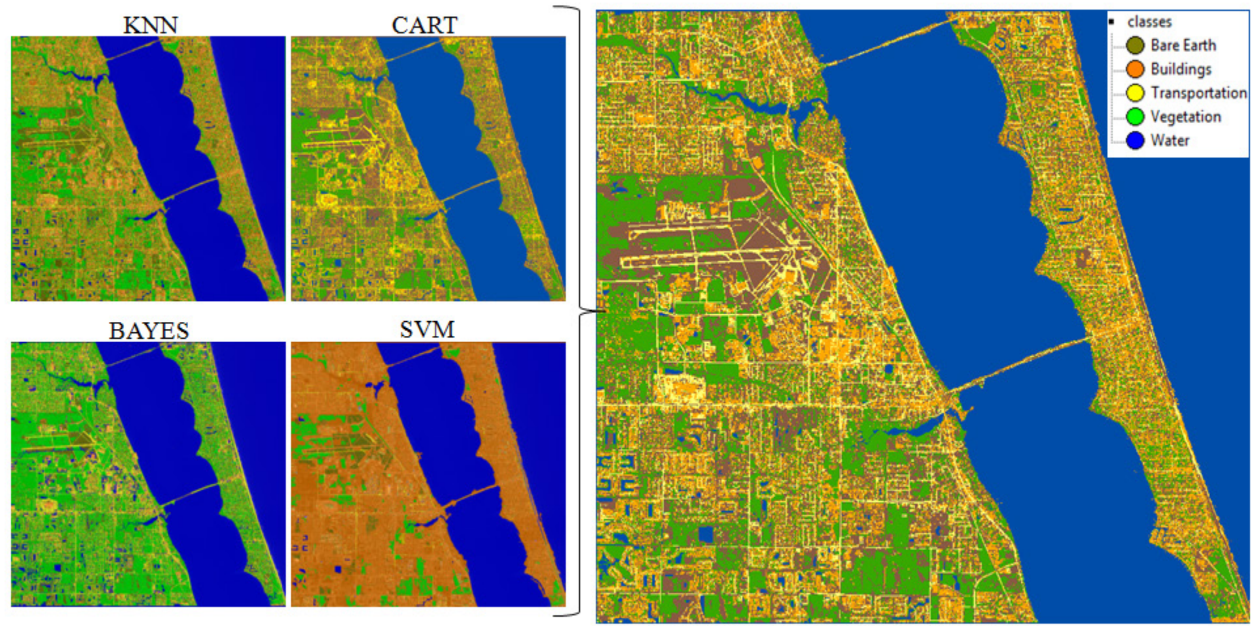

Figure 4. Enhanced Pixel Labeling

We perform a variety of manipulations and analyses before generating the shapes that have been identified as features. A measure of certainty, or confidence value, is returned as part of the statistics from our supervised classification process and is attributed to each of the shapefile features. The confidence values within a feature polygon are calculated and assigned to each of the grid points. Grid points for true negatives and false negatives were assigned a predictive value of zero, since they do not occur within our output. The features extracted autonomously in Figure 4 and the truth set in Figure 3 show that the described method produces exceptional results as shown in Table 2.

Creation of truth data is accomplished using accelerated graphical software that allows the user to edit extracted feature segments with one click and to generate a polygon shapefile that represents truth. This process is currently implemented as a menu of command buttons and tools, which are executed interactively. In a production environment with well-defined input and output, the command buttons could be fully automated. The output is the final truth shapefile. Note that while the one-click tools include editing tools, all evaluations were performed on automated extraction.

ASSESSING THE RESULTS

To evaluate the output of the production process, we use a uniform 20m point grid over the imagery. Attributes to support a confusion matrix are populated (hit, miss, correct rejection, and false alarm) using the extracted features and the truth set. At 20m spacing, the point grid is dense enough to show small variations, but because the sample size is so large (in the millions), they do not swing the results. Missing large features counts more than missing small ones; getting close to the truth counts, but an exact match is not required.

Table 2. Accuracy Improvement for Game Theory with Area Under ROC Curve

%	Bayes	CART	KNN	SVM	Nash
Bare Earth	75	77	75	76	82
Buildings	60	67	78	57	73
Transportation	71	61	56	51	59
Vegetation	85	79	81	58	84
Water	99	98	98	95	99
Nash Improvement	7	15	9	60	

Table 2 shows the improvement using our modified game theory algorithm. The results show a superior land cover classification product that saves on editing costs. The improvement achieved with the composite pixel labeling was better than any individual supervised classification method alone. The composite map is the best pixel label from KNN, CART, SVM, and Bayes supervised learning.

The ROC score is defined simply as the area under the ROC curve. For example, if the decision is purely random, then the ROC score should be very close to 0.5. On the other hand, for a perfect classifier, its ROC curve consists of (0, 0) to (0, 1) and then to (1, 1), and its ROC score is 1. Clearly, the higher the ROC score, the better the classifier performance. It is important to note that our work is not simply a binary classification problem, but a multiple classification problem. Instead of determining if one feature is present or not, we are determining if one of the five features is present. When you have more than a binary input of features, the curse of dimensionality will start having an effect [1]. Our solution helps with dimensionality by improving pixel decision making with multiple parameter or algorithm inputs.

SUMMARY

Feature extraction, especially in diverse urban environments, often suffers from feature misclassification. The accuracy and completeness of feature data collection of any kind depends on many factors, including the quality of the source imagery and the skill of the analyst. Improving algorithms for feature extraction is an ongoing process. Three topics for future work include: study how to optimize training sets; apply remote rules to a similar landscape, such as a neighboring (rather than distant and different place); and modification of supervised classification rules to generate a value for ROC curve, to use for selecting the threshold.

With the increasing number of new algorithms emerging, there is a need to assess their performance and sensitivities to various kinds of environments. This need is best addressed by developing standard image sets with adequate classification schemes and sufficiently representative training and testing samples. Only when more comprehensive test data sets exist to cover major environmental types of the world can we make more appropriate selection of algorithms for a particular application of remote sensing classification [7].

Until then, the results of supervised classification using local rules are promising and show that this method could be used to generate an LULC product that is fed into the material classification process to produce the final visualization and simulation output products.

The solution described in this paper contains a software and process workflow for efficient, targeted feature extraction. The automated batch process provides analysts with a select subset of features or change areas from an area of interest to review. Our solution combines the capabilities of traditional pixel-based methods and object-based image analysis. The results of the image analysis can be used directly, or with an existing vector database to identify features and change.

The goal of this solution is to reduce the labor cost to frequently update material classification product databases, given the large volume of high-quality satellite imagery available. Our system has resulted in a 20% labor cost reduction with richer classification especially in with vegetation features. The game-theoretic interpretation for imagery feature extraction is appealing, and it can cope with both asymmetric and/or negative similarities. It provides a robust method for optimizing pixel classification labeling and is intrinsically a multi-class approach.

REFERENCES

[1] Andrew P. Bradley. The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern Recognition*, 30(7):1145–1159, 1997.

[2] Berthod, Marc, et al. "Bayesian image classification using Markov random fields." *Image and Vision Computing* 14.4 (1996): 285-295.

[3] Bandukwala, F., "Extracting spatially and spectrally coherent regions from multispectral images," *Computer Vision and Pattern Recognition Workshops, 2011 IEEE Computer Society Conference*, pp.82,87, 2011.

[4] A. Erdem and M. Pelillo. Graph transduction as a non-cooperative game. *Neural Computation* (in press) *(preliminary version in GbR 2011)*.

[5] Fu, Zhouyu and Robles-Kelly, Antonio. "Discriminant Absorption-Feature Learning for Material Classification.." *IEEE T. Geoscience and Remote Sensing* 49 , no. 5 (2011): 1536-1556.

[6] Ghamisi, P.; Benediktsson, J.A.; Cavallaro, G.; Plaza, A., "Automatic Framework for Spectral–Spatial Classification Based on Supervised Feature Extraction and Morphological Attribute Profiles," *Selected Topics in Applied Earth Observations and Remote Sensing, IEEE Journal*, 2014.

[7] Li, C.; Wang, J.; Wang, L.; Hu, L.; Gong, P. Comparison of Classification Algorithms and Training Sample Sizes in Urban Land Classification with Landsat Thematic Mapper Imagery. *Remote Sens.* 2014, *6*, 964-983.

[8] Luo, Bin and Chanussot, Jocelyn. "Supervised Hyperspectral Image Classification Based on Spectral Unmixing and Geometrical Features" *Signal Processing Systems* 65, no. 3 (2011): 457-468.

[9] M. Migut, J. Gemert, M.Worring, Interactive Decision making using Dissimilarity to visually represented Prototypes, *IEEE Symp. on Visual Analytic Science and Technology*, 139 147, 2011.

[10] Nash, J.F., 1951. Non-cooperative games. Annals of Mathematics 54 (2), 286–295.

[11] Rahmes, M., Yates, J.H., Brown, P., Akbari, M., Tipping, A., "Cost Effective Evaluation of Change Detection", *ASPRS 2012 Annual Conference*, Tampa, Florida, October 30, 2012.

[12] L.S. Shapley. A value for n-person games. In H.W. Kuhn and A.W. Tucker, editors, *Contributions to the Theory of Games, Vol. II*, number 28 in Annals of Mathematics Studies, pages 307.317. Princeton Univ, 1953.

[13] Shi, Qian, Zhang, Liangpei and Du, Bo. "Semisupervised Discriminative Locally Enhanced Alignment for Hyperspectral Image Classification" *IEEE T. Geoscience and Remote Sensing* 51 , no. 9 (2013): 4800-4815.

[14] Strumbelj, Erik, and Igor Kononenko. "An efficient explanation of individual classifications using game theory." *The Journal of Machine Learning Research* 11 (2010): 1-18.

[15] Klaus D. Toennies, "Statistical Decision Making", 2008.

[16] Grigorios Tsoumakas, Ioannis Katakis, and Ioannis P. Vlahavas. "Mining Multi-label Data".Data Mining and Knowledge Discovery Handbook 2010: 667-685

[17] Volpi, M.; Matasci, G.; Kanevski, M. & Tuia, D.; Multi-view feature extraction for hyperspectral image classification, In *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning ESANN 2013, Bruges, Belgium*, 2013.

[18] Wang, Chong, David Blei, and Fei-Fei Li. "Simultaneous image classification and annotation." *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*. IEEE, 2009.

[19] Yu, S., Berthod, M., A Game Strategy Approach for Image Labeling, Computer Vision and Image Understanding 61(1):32-37, 1995.

[20] Yuntao Qian, Minchao Ye, and Jun Zhou. "Hyperspectral Image Classification Based on Structured Sparse Logistic Regression and 3D Wavelet Texture Features", *IEEE Transactions on Geoscience and Remote Sensing*, 2012, pp.1-16.