

A New Tool for Information Extraction and Mining from Satellite Imagery Available from Google Maps Engine

Sergio Bernabé and Antonio Plaza

Hyperspectral Computing Laboratory

Department of Technology of Computers & Communications

University of Extremadura, Avda. de la Universidad s/n, 10071 Cáceres, Spain

Emails: sbernabe@alumnos.unex.es, aplaza@unex.es

ABSTRACT- The wealth of satellite imagery available in web mapping service applications such as Google Maps, which now provides high-resolution satellite images from many locations around the Earth, has opened the appealing perspective of performing information extraction and mining tasks via the Google Maps application programming interface (API). In fact, the introduction of Google's mapping engine (the Google Maps service is free for non-commercial use) has prompted a worldwide interest in satellite imagery exploitation. In this paper, we present a new tool for information extraction and mining from satellite images available Google Maps. Our main focus in this work will be on describing the interface and how the user can interact with the tool. The functionalities included in the tool comprise the possibility to perform unsupervised classification of image portions selected by the user (at the maximum zoom level) using ISODATA and K-Means classifiers, followed by spatial post-processing based on majority voting. Selected regions in the classified portion can then be used to train a more sophisticated classifier able to map larger image areas in a manner transparent to the user. Several examples of use, focused on the analysis of different locations around the Iberian Peninsula and comprising different types of information extraction case studies (vegetation, semi-arid environments, urban areas, etc.) are illustrated and thoroughly described.

1. INTRODUCTION

The wealth of satellite imagery (Landgrebe, 2003) available in web mapping service applications such as Google Maps, which now provides high-resolution satellite images from many locations around the Earth, has opened the appealing perspective of performing classification and retrieval tasks via the Google Maps application programming interface (API) and other external libraries such as SwingX-WS. In fact, the introduction of Google's mapping engine prompted a worldwide interest in satellite imagery exploitation. The combination of an easily pannable and searchable mapping and satellite imagery tool such as Google Maps with advanced image classification and retrieval features has the potential to significantly expand the functionalities of the tool and also to allow end-users to extract relevant information from a massive and widely available database of satellite images (the Google Maps service is free for non-commercial use).

By using the Google Maps API or external libraries such as SwingX-WS, it is possible to embed the full Google Maps site into an external website application. Other similar services currently available comprise Yahoo Maps and OpenStreetMap. The characteristics of Yahoo Maps are similar to those available in Google Maps (although the spatial resolution of the satellite imagery available in Yahoo Maps is generally lower than the resolution of the

image data available from Google Maps). On the other hand, the OpenStreetMap follows a different approach. It is a collaborative project aimed at creating a free editable map of the world. Its design was inspired by sites such as Wikipedia. As shown by Fig. 1, the Google Maps service offers important competitive advantages, such as the availability of high resolution satellite imagery, the smoothness in the navigation and interaction with the system, the availability of a hybrid satellite view which can be integrated with other views (maps, etc.) and adequate adaptivity for general-purpose desktop applications. Despite its competitive advantages, the possibility to perform unsupervised or supervised classification of satellite images (Richards, 2005) is not available in Google Maps, despite image classification is widely recognized as one of the most powerful approaches in order to extract information from satellite imagery (Soille, 2003; Benediktsson et al., 2005; Bruzzone et al., 2006).

The remainder of the paper is organized as follows. Section 2 describes the programming libraries used to develop the proposed system and the integration of the different modules. Section 3 presents the unsupervised and supervised classification techniques considered in its implementation. Section 4 describes the integration of the different software modules. Section 5 presents an experimental evaluation of the proposed system which has been conducted by comparing the obtained classification

results with those provided by commercial software, such as the popular Research Systems ENVI package. Finally, section 6 concludes with some remarks and hints at plausible future research.

	Google maps	YAHOO! LOCAL	Openstreetmap
No Restrictions of use	No, up to certain limit	Yes	Yes
Hybrid satellite view	Yes	Yes, low visibility	No
High resolution imagery	Yes	No	No
Zooming levels	Very high quality	High quality	Medium quality
Error correction	Low	Low	Very high
Smoothness in navigation	Very high	High	High
Adaptivity for desktop applications	High	High	High

Figure 1. Comparison between the main functionalities of Google Maps, Yahoo Maps and OpenStreetMap.

2. LIBRARIES

In this section we describe the different steps that we followed for the integration of different software components in libraries towards the design of our proposed classification system for Google Maps satellite imagery. Our first approach to the design of such system was directed towards the exploitation of the Google Maps API in order to access the satellite imagery. However, we realized that the Google Maps API is mainly intended for exploitation in web-based applications and our goal was to develop an executable application in the PC. Another approach was to use Google Web Toolkit (GWT), an open source set of tools that allows web developers to create and maintain complex JavaScript front-end applications in Java. However, the fact that GWT is also intended for web-based applications led us to further discard this toolkit as the baseline package for the development of our tool. Finally, our approach to address the aforementioned issues was to resort to the SwingX-WS library. SwingX-WS attempts to simplify the use of web services (in the broad sense) by providing APIs that sit on top of existing libraries. It contains a set of JavaBeans for interacting with web services, but it allows embedding of those services into a standard application (this feature fit very well our desired functionality). The initial beans included in SwingX-WS comprise support for several Google web services such as searching news, video, images, and financial data, as well as a generic tile based mapping component. The SwingX-WS beans have been designed with graphical configuration in mind and work very well inside of a JavaBeans aware editor

such as NetBeans. The bean that we particularly exploited in the development of our application is JXMapView, a generic viewer for tile based map servers. For illustrative purposes, Fig. 2 shows a flowchart of our interaction with different libraries for accessing the satellite data available in Google Maps and our final choice, given by the use of SwingX-WS library with the JXMapView component.

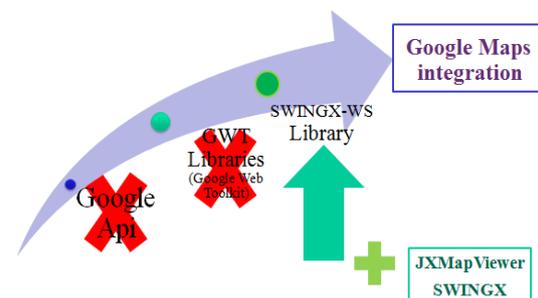


Figure 2. Flowchart describing our interaction with different libraries for accessing the satellite data available in Google Maps and our final choice (SwingX-WS library with the JXMapView component).

In addition to JXMapView, other additional SwingX-WS and SwingX modules were used in the development of our application. In the following, we provide a brief overview of these additional modules:

- **JXMapKit.** The JXMapKit is a pair of JXMapViews preconfigured to be easy to use with common features built in. This includes zoom buttons, a zoom slider, and a mini-map in the lower right corner showing an overview of the map.
- **TileFactoryInfo.** This component encapsulates all information specific to a map server and is needed for the considered development.
- **GeoPosition.** Provides the geographical latitude and longitude and allows centering of the satellite image in a certain geographic location, thus allowing us to access such locations from our application.
- **Painter.** This API allows developers to be able to customize the background painting of a JPanel. Since many components within SwingX extend JPanel, the developer can implement custom painting on many parts of SwingX.
- **CompoundPainter.** Painters can be combined together by using the CompoundPainter, which uses an array to store several Painters, and the order in which these items should be painted in the application.

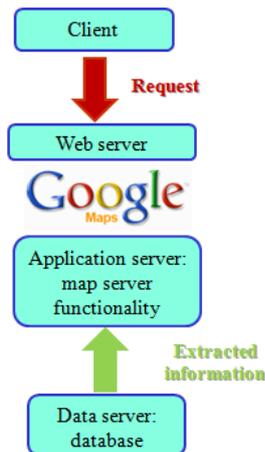


Figure 3. Flowchart describing the interaction procedure with Google Maps to obtain the satellite data.

Fig. 3 describes the overall procedure for interacting with Google Maps via the SwingX-WS library to obtain the satellite data. As shown by Fig. 3, a client (user) sends a request to a web server via SwingX-WS. The web server interacts with an application server that provides the map server functionality. Finally, the application server interacts directly with a database from which the extracted information is provided.

3. CLASSIFIERS

The proposed system incorporates functionalities of unsupervised clustering and supervised classification, all followed by spatial post-processing based on majority voting (Gamba et al., 2004).

Unsupervised clustering aims at grouping pixels if feature space, so that pixels belonging to the same cluster are spectrally similar (Richards, 2005). In our implementation, we have used the well-known ISODATA (Ball and Hall, 1965) and k-means (Hartigan and Wong, 1979) algorithms for this purpose. For instance, the ISODATA is a squared-error clustering method. The algorithm starts with a random initial partition of the available pixel vectors in the original image into c candidate clusters. It then iteratively optimizes this initial partition so that, at each iteration, a partition of the original set into c clusters is computed. A relevant issue for the ISODATA algorithm is how to set the number of clusters in advance.

In our work, the aforementioned choice is left to the end-user, who can adjust the quality of the clustering by interactively setting this parameter. Once a preliminary segmentation of the original image has been achieved via unsupervised clustering, a supervised procedure can be applied to classify other

different areas based on the training sites selected in a different spatial location. In our tool, this can be accomplished using the minimum distance and maximum likelihood classifiers (Richards, 2005). Conventional maximum likelihood classification is based on the assumption that the probability distribution for each spectral (color) class is of the form of a multivariate normal model with dimensions which equal the number of color bands. A spatial post-processing module can be applied to refine the outcome of the segmentation by simply sliding a square neighborhood window (with sizes ranging from 3×3 to 7×7 pixels) centered in each classified pixel and applying a majority voting procedure in which the central pixel is assigned to the most predominant class in the neighborhood window.

4. INTEGRATION

The final implementation of our system consists of the integration of the different software modules developed (unsupervised and supervised classifiers) with the functionalities provided by the SwingX-WS libraries (provided by the SwingX libraries), in the form of a general-purpose desktop application. For this purpose, we have resorted to the Java programming language (see Fig. 4), which is a multi-platform environment that simplifies porting of our tool to different environments. Specifically, we have resorted to the NetBeans platform, which allows applications to be developed from a set of modular software components called modules. In this framework, applications can install modules dynamically and any application can include the update center module to allow users of the application to download digitally-signed upgrades and new features directly into the running application. Reinstalling an upgrade or a new release does not force users to download the entire application again. The platform offers reusable services common to desktop applications, allowing developers to focus on the logic specific to their application.

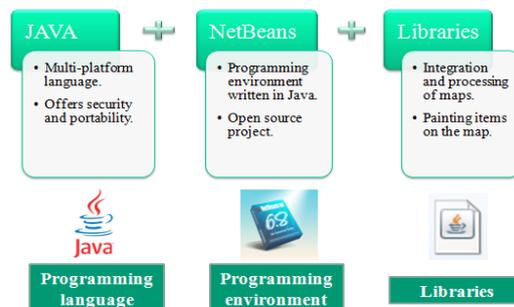


Figure 4. Different software modules used for the development of our application.

5. EXPERIMENTAL VALIDATION

In this section, we perform an experimental validation of the unsupervised and supervised classification features of our tool tested using satellite images obtained from Google Maps across different locations. The experimental validation of unsupervised (ISODATA, k-Means) and supervised (minimum distance, maximum likelihood) classification algorithms has been conducted by comparing the results provided by our implementations with those available in a well-known commercial software package: the Environment for Visualizing Images (ENVI) package distributed by ITT Visual Information Solutions. In our tests, conducted across four different locations around the world, we adopt exactly the same parameters when running our implementations and those available in the ENVI package, comparing the results in terms of the overall accuracy (OA) resulting from the confusion matrix between the classification maps provided by ENVI and by our tool.

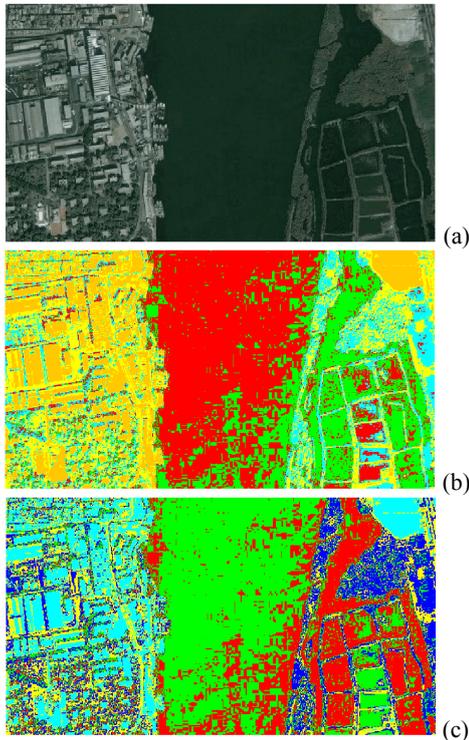


Figure 5. (a) Satellite image collected over a stretch of the Nile river in Egypt (the image is available online through Google Maps engine). (b) Unsupervised classification result provided by our implementation of ISODATA. (c) Unsupervised classification result provided by ENVI's implementation of ISODATA.

In the following, we present the obtained results in a specific case study focused on classification of satellite images available from the Nile river in the north of Cairo (Egypt) (see Fig. 5). The resolution of the image is quite high, with approximately 5 meters per pixel. As show by Fig. 5, the color labels for our implementation and the one available in ENVI are different, but the classification maps are very similar. In both cases, the parameters for both algorithms have been set to exactly the same values, with the number of clusters set to 5. For illustrative purposes, Table 1 reports the OA (in percentage) measured after comparing our ISODATA classification map with the one obtained by ENVI. The similarity between both classification maps is very high, achieving an OA close to 90% assuming the ENVI ISODATA map as the ground reference.

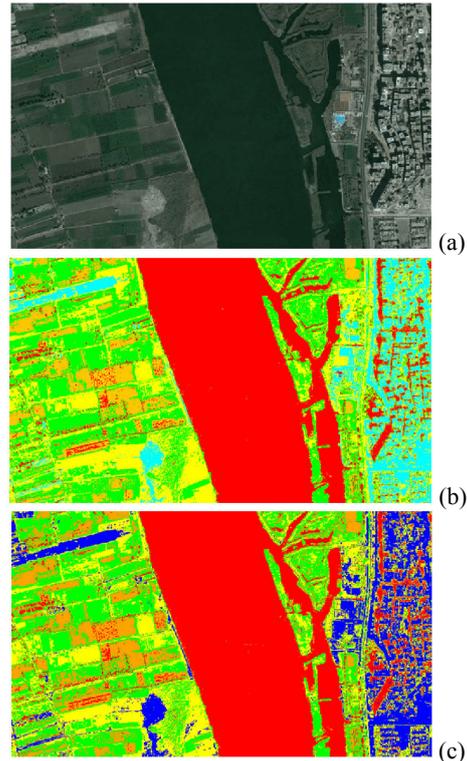


Figure 6. (a) Satellite image over a different location in the Nile river. (b) Supervised classification result provided by our implementation of maximum likelihood (trained with the areas derived from other zone by ISODATA). (c) Supervised classification result provided by ENVI's implementation of maximum likelihood (trained with the areas derived from other zone by ISODATA).

In a second experiment, we select a larger image over a different location centered over the Nile river in Egypt (see Fig. 6). The spatial resolution of this image is approximately 5 meters per pixel. In order to

classify this scene in supervised fashion, we have selected the areas resulting from an unsupervised classification (using ISODATA) and used those areas to train our maximum likelihood classifier using 5 clusters. In order to obtain this map, the maximum likelihood algorithm in ENVI was trained using exactly the same areas and with the same parameters as those used in our implementation.

Soil #1 (blue)	Water#1 (green)	Urban #1 (orange)	Water (red)	Soil #2 (yellow)	Overall accuracy
63.89%	96.98%	99.74%	89.01%	94.59%	88.47%

Table 1. Overall (OA) and individual accuracies (percentage) after comparing our ISODATA classification map with the one obtained by ENVI for the scene in Fig. 5(a).

Urban areas (blue)	Vegeta- tion #3 (green)	Vegeta- tion #1 (orange)	Water (red)	Vegeta- tion #2 (yellow)	Overall accuracy
79.19%	99.68%	75.58%	99.99%	98.18%	90.52%

Table 2. Overall (OA) and individual accuracies (in percentage) after comparing our maximum likelihood classification map with the one obtained by ENVI for the image in Fig. 6(a).

For illustrative purposes, Table 2 reports the OA (in percentage) measured after comparing our maximum likelihood classification map with the one obtained by ENVI. The similarity between both classification maps is very high, achieving on OA above 90% assuming the ENVI maximum likelihood map as the ground reference. These results indicate a high degree of similarity between our implementations of available classification algorithms and those available in ENVI. The main contribution of this work is the possibility of integrating these classifiers with the satellite images derived from Google Maps via the developed tool.

To conclude this section, Fig. 7 shows different views of the developed tool. As shown by the figure, the tool allows selecting an area to be classified, obtaining classification results both in unsupervised and supervised fashion, retaining the classified area at different zoom levels (although the classification is obtained at the maximum zoom level), and other functionalities not illustrated in Fig. 7 such as spatial post-processing of obtained results for increased spatial consistency, managing of the resulting classification and extracted satellite images, loading/storing of results via file logs which can be saved in a database, automatic positioning in any latitude and longitude coordinates in the entire Google Maps database, overlaying of classification results with different views (satellite, map, hybrid), etc. Overall, we feel that the developed tool incorporates

interesting additional functionalities to the Google Maps engine (particularly in the possibility of better exploiting the satellite images available from this tool in different application domains).

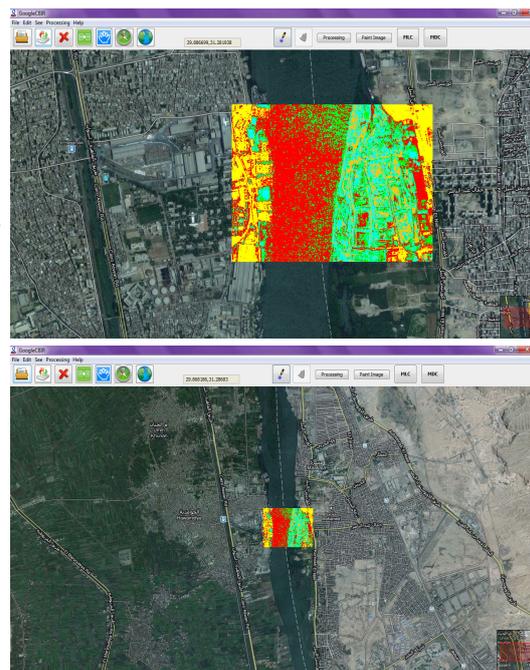


Figure 7. Different views of the developed tool. Unsupervised classification result provided by our implementation of ISODATA superimposed on the tool (top). Zoom reduction retaining the classified area (bottom).

6. CONCLUSION AND FUTURE LINES

This paper has described a new application for unsupervised and supervised classification of satellite images from Google Maps. The system has been developed using the SwingX-WS library, and incorporates functionalities such as unsupervised classification of image portions selected by the user (at the desired zoom level) using ISODATA and k-Means, and supervised classification using the Minimum Distance and Maximum Likelihood, followed by spatial post-processing based on majority voting. Our experimental results reveal that the proposed tool provides classification maps of high similarity with regards to those provided by ENVI for the same satellite imagery, but with the possibility to perform classification of any image portion available in Google Maps engine, both in unsupervised and supervised fashion.

In future developments, we plan to extend the developed tool with the incorporation of content-based image retrieval functionalities. For that purpose, the strategy will be based on a query system linked to

feature extraction from an image repository (Google Maps). The retrieved features (which will comprise shape descriptors, texture features, etc.) will be stored in a database of features and used to compare the feature vector of the input query with those recorded in the database by means of a similarity function, which will provide a result to the end-user in the form of image portions (across different locations) with sufficient similarity with regards to the features in the input query.

Finally, another drawback of the proposed tool is the high computational cost of extracting and processing large images (especially at the highest zoom level, which results in large image sizes). To address this issue, we are currently experimenting with different forms of high performance computing architectures (Plaza et al., 2006; Plaza and Chang, 2007; Plaza and Chang, 2008). The most promising strategy seems to be the parallelization of the different software modules for efficient processing in multiple cores (generally available in modern desktop PCs) as well as in graphics processing units (GPUs) of NVidia type (Setoain et al., 2008; Tarabalka et al., 2009). Both mechanisms allow incorporation of high performance computing capabilities at relatively low cost in order to speed-up the computations performed by the developed tool, which mainly comprise regular image processing-type computations which are quite appealing for parallel implementation.

ACKNOWLEDGEMENT

This work has been supported by the European Community's Marie Curie Research Training Networks Programme under reference MRTN-CT-2006-035927 (HYPER-I-NET). Funding from the Spanish Ministry of Science and Innovation (HYPERCOMP/EODIX project, reference AYA2008-05965-C04-02) is gratefully acknowledged.

REFERENCES

- Ball, G., and Hall, D., 1965, ISODATA, a novel method of data analysis and classification. *Technical Report AD-699616*, Stanford University.
- Benediktsson, J. A., Palmason, J. A., and Sveinsson, J. R., 2005, Classification of hyperspectral data from urban areas based on extended morphological profiles. *IEEE Transactions on Geoscience and Remote Sensing*, 42, 480–491.
- Bruzzone, L., Chi, M., and Marconcini, M., 2006, A novel transductive SVM for the semisupervised classification of remote sensing images. *IEEE Transactions on Geoscience and Remote Sensing*, 44, 3363–3373.
- Gamba, P., Dell'Acqua, F., Ferrari, A., Palmason, J. A., and Benediktsson, J. A., 2004, Exploiting spectral and spatial information in hyperspectral urban data with high resolution. *IEEE Geoscience and Remote Sensing Letters*, 1, 322–326.
- Hartigan, J. A., and Wong, M. A., 1979, Algorithm AS 136: A k-means clustering algorithm. *Journal of the Royal Statistical Society, Series C*, 28, 100–108.
- Landgrebe, D. A., 2003, *Signal Theory Methods in Multispectral Remote Sensing* (Hoboken, NJ: Wiley).
- Plaza, A., and Chang, C.-I., 2007, *High Performance Computing in Remote Sensing* (Boca Raton, FL: CRC Press).
- Plaza, A., and Chang, C.-I., 2008, Clusters versus FPGA for parallel processing of hyperspectral imagery. *International Journal of High Performance Computing Applications*, 22, 366–385.
- Plaza, A., Valencia, D., Plaza, J., and Martinez, P., 2006, Commodity cluster-based parallel processing of hyperspectral imagery. *Journal of Parallel and Distributed Computing*, 66, 345–358.
- Richards, J. A., 2005, Analysis of remotely sensed data: the formative decades and the future. *IEEE Transactions on Geoscience and Remote Sensing*, 43, 422–432.
- Setoain, J., Prieto, M., Tenllado, C., and Tirado, F., 2008, GPU for parallel on-board hyperspectral image processing. *International Journal of High Performance Computing Applications*, 22, 424–437.
- Soille, P., 2003, *Morphological Image Analysis: Principles and Applications* (Berlin, Germany: Springer-Verlag).
- Tarabalka, Y., Haavardsholm, T. V., Kasen, I., and Skauli, T., 2009, "Real-time anomaly detection in hyperspectral images using multivariate normal mixture models and GPU processing. *Journal of Real-Time Image Processing*, 4, 1–14.