

I-FGM as a Real Time Information Retrieval Tool for E-Governance

Eugene Santos, Jr., Dartmouth College, USA

Eunice E. Santos, Virginia Polytechnic Institute & State University, USA

Hien Nguyen, University of Wisconsin, USA

Long Pan, Virginia Polytechnic Institute & State University, USA

John Korah, Virginia Polytechnic Institute & State University, USA

Huadong Xia, Virginia Polytechnic Institute & State University, USA

ABSTRACT

Homeland security and disaster relief are some of the critical areas of E-governance that have to deal with vast amounts of dynamic heterogeneous data. Providing rapid real-time search capabilities for such applications is a challenge. Intelligent Foraging, Gathering, and Matching (I-FGM) is an established framework developed to assist users to find information quickly and effectively by incrementally collecting, processing and matching information nuggets. This framework has been successfully used to develop a distributed, unstructured text retrieval application. In this paper, we apply the I-FGM framework to image collections by using a concept-based image retrieval method. We approach this by incrementally processing images, extracting low-level features and mapping them to higher level concepts. Our empirical evaluation shows that our approach performs competitively compared to some existing approaches in terms of retrieving relevant images while offering the speed advantages of distributed and incremental process and unified framework between text and images

Keywords: agent technology; CBIR; distributed systems; heterogeneous databases; image retrieval; information search and retrieval; intelligent agents; IS performance evaluation; real-time IS

INTRODUCTION

One of the main challenges in e-governance is to effectively and efficiently find relevant information from vast amounts of dynamic heterogeneous sources quickly under the pressures and limitations of time, supporting tools, and

resources. For instance, when natural disasters such as Hurricane Katrina (2005) or the Asian Tsunami of 2004 happen, we need to quickly locate the areas that are most affected and collect information in order to estimate the needed aid for items such as medicines, foods, and

drinking water. Unfortunately, in such a situation, frontline communications are typically chaotic (or non-existent) and/or there are too many channels of information from different sources that make the retrieval of relevant pieces of information a lot harder. For "hot spots" such as disaster relief areas, combat zones, and so forth information is changing rapidly and as such there is only a small window of time for information to remain valid. Additionally, various types of data representation are used such as images, blogs, maps, news reports, audios, and videos. Each type of data format contains important and indispensable information for the various governmental agencies. Therefore, in order to better assist these agencies in addressing these challenges, there is a clear and urgent need to develop a system that rapidly provides real-time retrieval capabilities of heterogeneous sources of information. There are three main issues that we need to address: (1) how to gather and retrieve information quickly in a real-time setting given the limitations of resources and time; (2) how to address the problem of heterogeneous data; and, (3) how to improve retrieval success.

We address the aforementioned issues by developing a framework for intelligent foraging, gathering, and matching (I-FGM) that incrementally and distributively gathers, processes, and matches information nuggets to assist users at finding information quickly and effectively. In our previous work (Santos, Santos, Nguyen, Pan, & Korah, 2005, Santos et al., 2006), I-FGM has been empirically demonstrated to be an effective tool for text retrieval on large and dynamic search spaces. Even though unstructured text is a typical format for most databases/sources, images are also popular with significant support from commercialized search engines such as Google, Yahoo!, and MSN. In order to demonstrate that I-FGM is a general framework for information retrieval, it is necessary to study the system's ability at effectively handling such heterogeneous data which contains at least text and images. In this article, we apply the I-FGM framework on image collections by using a concept-based

image retrieval method. We approach this by incrementally processing the images, extracting low-level features, and then mapping them to higher level concepts. The novelties of our approach lie with the distributed storage and incremental processing and matching of information nuggets extracted from a region-based wavelet image retrieval scheme. We deploy a concept-based image retrieval algorithm that maps low-level features of the images to high-level concepts. In this way, we are also able to translate the visual information of images into document graphs (Santos et al., 2005), which are used in I-FGM as a common representation of information for heterogeneous datatypes. Thus, I-FGM provides a seamless integration of text and image through a single unifying semantic representation of content. By implementing and testing our image retrieval algorithm in I-FGM, we can validate the I-FGM framework as a method for future unified rankings of heterogeneous documents.

This article is organized as follows: In the methodology section, we first provide a brief background on current image retrieval systems. Next, we give an overview of the image retrieval algorithm that was implemented in I-FGM. We then describe the architecture of I-FGM and the implementation details of the current prototype. This section is followed by a description of our system evaluation procedure along with simulation results. We validate the performance of the image algorithm and show that it can be used in I-FGM for faster retrieval of images. Finally, we will present our conclusions and future work.

METHODOLOGY

Background and Related Work

The goal of the I-FGM framework is to facilitate real time searches in large and dynamic search spaces. In particular, I-FGM aims to integrate the retrieval of multiple data types; to enable a unified framework to compare heterogeneous documents; and to employ resource allocation policies facilitated by partial processing

to conduct efficient anytime searches. In this article, we integrate image retrieval in I-FGM by employing an algorithm based on techniques from Content Based Image Retrieval (CBIR). We now present the current state of research in CBIR and the approaches that have been leveraged in I-FGM.

CBIR is an important sub-area of image retrieval that aims to classify, identify, and retrieve images based on its visual content rather than using human generated information such as meta-tags and annotations. Through the previous decade, there has been significant progress in this field, documented in various surveys (Rui, Huang, & Chang, 1999; Smeulders, Worring, Santini, Gupta, & Jain, 2000; Velkamp & Tanase, 2000). Although, a general solution for image retrieval still eludes researchers due to the difficulties posed by the well-documented phenomena of sensory gap and semantic gap (Enser & Sandom, 2003), CBIR techniques have shown great promise in automating the process of interpreting images. We decided to employ CBIR in I-FGM because other methods require extra-image information such as meta-tags or HTML text from Web pages containing the image. In the CBIR methods, the visual content or low-level features of the images have to be converted into an intermediate form or signature before it can be interpreted. Our choice of low-level features and feature signature are important as it has a direct influence on the performance of the system. Examples of low-level features in images are color, texture, and contrast.

Region-based methods are a widely used type of CBIR. In these methods (Carson, Belongie, Greenspan, & Malik, 2002; Li, Wang, & Wiederhold, 2000; Ma & Manjunath, 1997; Natsev, 2001) images are divided into regions, each of which has similar pixel characteristics. This collection of regions forms a signature for the image. The premise behind these methods is that similar images will have similar regions. Methods such as WALRUS (Natsv, Rastogi, & Shim, 2004) and SIMPLcity (Wang, Li, & Wiederhold, 2001) use wavelets (Daubechies, 1992) to represent the feature signature of regions. Methods such as (Natsv et al., 2004;

Wang et al., 2001) generate regions by calculating a feature signature for a part of the image (window) at a time. The window is slid across the image such that every pixel of the image is contained in at least one window. These windows are then clustered based on a distance measure. Each cluster is then considered to be a region. These methods work well even with complex images containing multiple concepts or objects. One of the drawbacks in these methods is that they take query images instead of a text query. This is also one of the reasons preventing its direct usage within the I-FGM framework.

We will use the image segmentation and region clustering techniques from WALRUS in our image retrieval algorithm. The feature vector that we use to represent the low-level features of the image is similar to the one used in SIMPLcity. Although SIMPLcity is a well known regions method, it does only coarse-grained classification of images into semantic classes such as "outdoor," "indoor," "texture," "non-textured," and "photographs." I-FGM on the other hand deals with more fine-grained classifications. Our algorithm seeks to map low-level features to high-level concepts by forming regions of similar pixel characteristics and then mapping them to the concepts. Although this is similar to the CAMEL (Natsv, Chadha, Soetarm, & Vitter, 2001) algorithm (a variant of the WALRUS method), there are significant differences. The concepts used in CAMEL are single object concepts such as "apple." It compares the feature vectors for this concept with incoming images and tries to determine if the concept is contained in the image using a distance-based similarity measure. We found that CAMEL does not deliver the required performance when implemented with the images that we typically deal with. The main reason is that it is difficult to accurately identify single-object concepts in images. Another important reason why we cannot use this method is that it uses a distance-based similarity measure that is adversely affected by the presence of non-relevant features in the images. A better option is to use machine learning techniques to match the low-level features

of concepts and images. ALIPR (Li & Wang, 2006) is another automatic annotation method that tries to map low-level features to keywords. ALIPR is an established method that has given good performance. It has also been used commercially. Unfortunately, it cannot be used with the partial processing paradigm of I-FGM. In I-FGM, we incrementally process an image, using each partial step to gain an insight into the relevancy of image and proportionately grant resources. The information gathered from one step has to be reused in successive partial steps so as to minimize the overhead. In ALIPR, the probability distributions associated with its feature vectors cannot be calculated with the incomplete image information available at each partial step in I-FGM.

Another technology that we will leverage in I-FGM is the prevalent WWW image search engine. I-FGM uses the Internet as a search space in this article since it is a rich source of images and dynamic in nature. We will use image search engines such as Google and MSN to do a first order filtering on the search space similar to the Internet and narrow down the search space. These search engines cannot be used for real-time retrieval as they use static indexes that are built off-line (Kherfi, Ziou, & Bernardi, 2004). Also, they require human-generated information in the form of the html text surrounding the images (Tsymbalenko & Munson, 2001) or meta-tags. On the other hand, analysts using I-FGM must deal with images from various sources: satellite imagery, aerial reconnaissance images, and so forth, which may not be accompanied by annotations.

In this article, we seek to incorporate image search by implementing a concept-based image retrieval algorithm. Our algorithm will extract low-level features of images and map to predefined concepts. We will show how we can convert images into document graphs, which are the common knowledge representation in I-FGM. By doing so, we will validate our framework for unifying heterogeneous data types.

Image Retrieval for I-FGM

As mentioned earlier, we have validated I-FGM as a credible framework to provide real-time search in large and dynamic free-text databases. Here, we incorporate image retrieval into the framework. As such, our goal is not about presenting a new method for image retrieval but rather a method that is built on principles that have been tested and validated by the image retrieval research community that can be used in the I-FGM framework. We will use our image retrieval algorithm to validate the I-FGM framework on image collections. We also intend to show how smart resource allocation strategies can reduce computational costs and save time when dealing with large and dynamic databases. Furthermore, by representing the information in images as a document graph, we will have demonstrated how a unified representation can be applied to multimedia (text and/or image) documents.

As we mentioned earlier, the third point on heterogeneity is a particularly important issue in the databases that I-FGM expects to deal with in support of governmental agencies. Our databases may contain documents of different data type (images, text, multi-media files). Within each data type, another level of heterogeneity can occur. For example, images can occur in different forms such as photographs, satellite imagery, maps, and so forth. A straightforward method of dealing with this is to use separate retrieval methods for each data type and somehow combine their ranking results. This is a very ineffective method since the ranking measures used by the separate methods may share little or no semantic relationships. For example, it is unclear whether an image of similarity x based on texture and color values is equally relevant to users, or a text of the same similarity value based on keyword matches. We resolve this problem by extracting the information from a document using a given retrieval method and then representing the semantic contents of the document in a common graphical structure called the document graph. A document graph (example in Figure 1) consists of nodes representing concepts, and edges representing

relations between nodes. This representation extracts the main information contained in any given document.

The image retrieval method that we use with I-FGM requires certain characteristics as we alluded to in the previous section. First, it should be amenable to partial processing where information gleaned from each partial step is re-useable by the next partial step. Next, the similarity measure calculated from each partial step should be a good indicator of the final similarity value, which helps in designing resource allocation strategies. Also, the time taken to process an image with partial processing should not be significantly larger than one-step

processing. Finally, the method should also be able to convert the visual content of images into document graphs. All of this will allow us to accept natural language queries so as to be consistent with the text retrieval methodology already deployed within I-FGM.

A step-by-step representation of the retrieval process is depicted in Figure 2. When a new image arrives in the system, it is incrementally processed to extract its low-level features in the images. These features are compared with those images found in the concept library and matching concepts are then noted. Here, the concept library is a collection of well-known concepts along with its annotation and an image feature vector representing its characteristics. The concept feature vector is modeled by a machine learning algorithm using a training set. The text annotations of these concepts are then combined to form the textual description of the image. The description is then converted into a document graph and compared with the document graph representation of the natural language query (query graph) to compute a similarity value. We now describe the steps in the image retrieval in detail:

1. **Low-level feature extraction:** Each image is represented by a 16-dimensional feature

Figure 1. Document graph for the text "Workers at coal mine on strike"


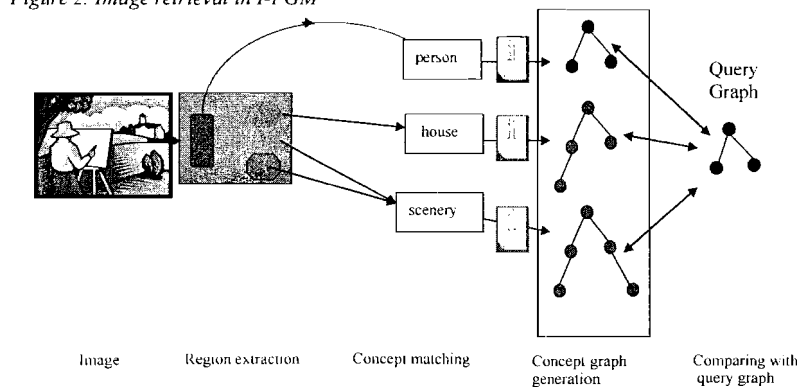


Figure 2. Image retrieval in I-FGM



vector with each dimension representing a region of image. An image is said to have a maximum of 16 regions. This is borrowed from SIMPLcity where it was found that 16 regions give good performance. Since we also use the same feature vector, we decided to use this value. Each region is represented by a six-dimensional vector representing the color and texture information of the images. In order to extract the regions from an image, the image is incrementally processed with a sliding window. We use the LUV color space. For each window, a signature is calculated from the average values of the color and the 2x2 wavelet transform on the L component. Haar wavelets are used as they are computationally efficient and have good performance (Natsev et al., 2004). The first three components of the signature are the average values of L, U, and V values of the pixels in the window, respectively. The high-frequency components of the wavelet transform form the remaining three components of the window signature. The signature of a window w is of the form:

$$h_w = \{l_w, u_w, v_w, m^1_w, m^2_w, m^3_w\}$$

where l_w, u_w, v_w are the average values of the pixel color in windows w and m^1_w, m^2_w, m^3_w are the high-frequency components of the Haar wavelet transform of the L values in w .

Based on their signatures, the windows are clustered together using a clustering algorithm. We use the BIRCH (Zhang, Ramakrishnan, & Livny, 1996) clustering algorithm as it is one of the most efficient clustering algorithms available. Each of the clusters form a region represented by its centroid. The area of each region is also calculated simultaneously. The regions are sorted according to its area and the largest region's signature form the first dimension of the image feature vector, the second largest region form the second dimension, and

so on. Hence the image feature vector of an image x is represented by the vector:

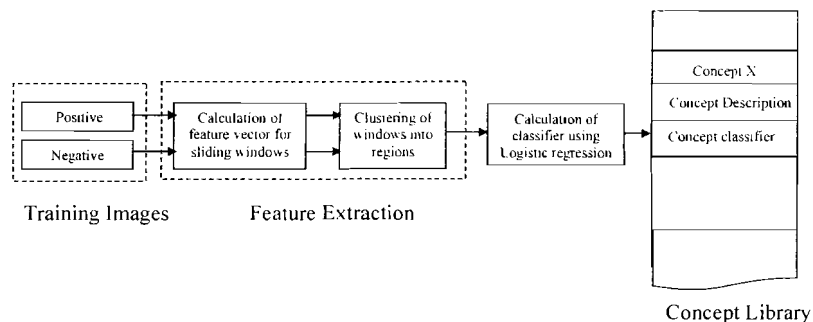
$$f_x = \{l_{x,1}, u_{x,1}, v_{x,1}, m^1_{x,1}, m^2_{x,1}, m^3_{x,1}, \dots, l_{x,16}, u_{x,16}, v_{x,16}, m^1_{x,16}, m^2_{x,16}, m^3_{x,16}\}$$

where $l_{x,i}, u_{x,i}, v_{x,i}$ are the average color values in the centroid for region i and $m^1_{x,i}, m^2_{x,i}, m^3_{x,i}$ are the wavelet transforms for the L color space in the centroid for region i .

2. **Creation of the concept library:** The library is created off-line and this process is depicted in Figure 3. It consists of a set of concepts that are commonly found in the image databases that I-FGM is searching. Each concept is accompanied by a short textual description, which is generated manually. In our method the low-level features of images are mapped to concepts in the library. The descriptions of the relevant concepts are then used to automatically generate the annotation for the images. This method of automatic annotation of images has less manual intervention than most other methods that generally require human generated metadata for each image. A classifier for each concept is generated by a machine learning algorithm based on logistic regression. A detailed description of the machine learning algorithm is presented later on. A set of training images are used to create the classifier. The training images are of two types: images that represent the concept (positive training) and images that do not represent the images (negative training). The image feature vector is extracted for each of these images and fed into the machine learning technique to create the classifier.

3. **Image matching:** This process matches low-level features of the images to high-level concepts (Figure 4). For each image, we incrementally compute its feature vector. During each incremental step, the image feature vector for the portion of the image that has been processed is calculated and

Figure 3. Off-line creation of concept library



compared against the concept classifiers in the library. A machine learning algorithm determines if a particular concept is relevant to the image or not. If a concept is relevant, its concept annotation is stored. At the end of the matching phase, the annotation of all matched concepts are combined together to form the image description. The description is converted into a document graph and compared with the document graph of the query where the similarity is calculated.

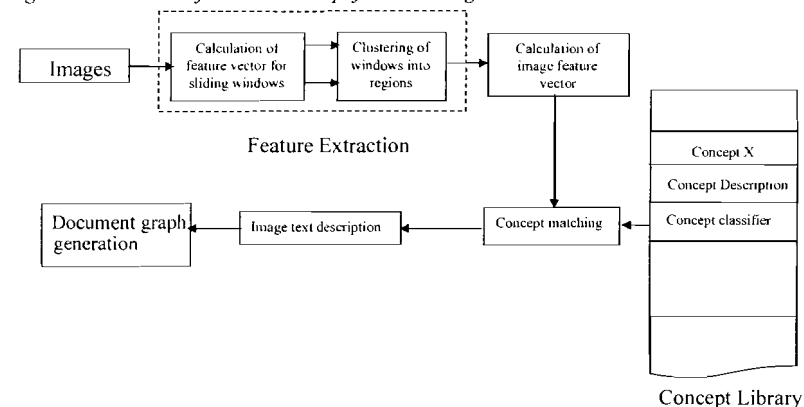
4. **Machine learning algorithm:** Logistic regression (LR) is appropriate for matching concepts in the concept library with images because it can be modeled as a binomial classification problem. LR has been found to be appropriate for binary classification and can be improved to achieve fast speed and reliability (Komarek, 2004). Some studies have compared LR with some other popular classifiers used in machine learning such as tree induction (Perlich, Provost, & Simonoff, 2003), Naïve Bayesian classifier (Mitchell, 2005), and support vector machine (Vapnick, 1999). It has been found (Perlich et al., 2003) that LR performs better with small training data set while tree induction is better with a large data set. Naïve Bayesian classifiers require conditional independence among variables. Therefore, its application is

limited, and in situations where such an assumption does not hold, LR is found to perform more accurately than Naïve Bayes (Mitchell, 2005). The loss function used in LR has been shown to be approximated by a support vector machine (Vapnick, 1999; Zhang, Jin, Yang, & Hauptmann, 2003). Algorithms based on LR have been used in CBIR (Ksantini, Ziou, Colin, & Dubeau, 2007) and in relevance feedback (Caenen & Pauwels, 2002) for image retrieval. Each concept is either relevant to the given image or not. The dependant variable (relevancy of the concept) can have two values: 0-not relevant and 1-relevant. The explanatory variables are the centroids of the regions in the image. The classifier for a particular concept is calculated by using positive and negative images. As mentioned before, the image vectors of each of these images are used to train the classifier. The training inputs are of the form:

$$t_j = \{l_{i,1}, u_{i,1}, v_{i,1}, m_{i,1}^1, m_{i,1}^2, m_{i,1}^3, \dots, l_{i,16}, u_{i,16}, v_{i,16}, m_{i,16}^1, m_{i,16}^2, m_{i,16}^3, r_j\}$$

where $l_{i,r}$, $u_{i,r}$, $v_{i,r}$ are the average color values in the centroid for region i in the training image j . $m_{i,r}^1$, $m_{i,r}^2$, $m_{i,r}^3$ are the wavelet transforms for the L color space in the centroid for region i in the training image j .

Figure 4. Extraction of relevant concept from the image



$$r_j = \begin{cases} 0 & \text{for negative images} \\ 1 & \text{for positive images} \end{cases}$$

The logit function for an image vector x is given as

$$g(x) = \alpha + \beta_1 l_{x,1} + \beta_2 u_{x,1} + \beta_3 v_{x,1} + \beta_4 m_{x,1}^1 + \beta_5 m_{x,1}^2 + \beta_6 m_{x,1}^3 + \dots + \beta_{91} l_{x,16} + \beta_{92} u_{x,16} + \beta_{93} v_{x,16} + \beta_{94} m_{x,16}^1 + \beta_{95} m_{x,16}^2 + \beta_{96} m_{x,16}^3$$

where α , β_i ($1 \leq i \leq 96$) are weights calculated by the maximum-likelihood method from the training images. The LR algorithm is implemented using the WEKA software (Witten & Frank, 2005).

System Architecture

I-FGM has a flexible architecture based on a multi-agent system that enables it to quickly deploy and redeploy computational resources. It also provides plug-and-play facilities for rapid deployment of any new retrieval method. Based on functionality, I-FGM can be decomposed into the following components:

1. I-Forager
2. glG-Soup
3. glG-Builder
4. I-Matcher
5. Blackboard

A detailed description of the various components of I-FGM can be found in Santos et al. (2005, 2006):

- **I-Forager:** This component uses third-party search tools to perform a first-order filtering of the documents in the search space. In our prototype, three I-foragers are used to download potentially relevant images into the glG Soup. Google, MSN, and Yahoo are used in the I-Foragers. The query is sent to each of the I-Forager and the top 50 results are downloaded into the glG Soup. I-Forager calculates the first order similarity that is based on the rank returned by the search engines.
- **glG-Soup:** This component is the central data repository of I-FGM that holds the images as they are processed by I-FGM. The glG-Soup is implemented using a NFS share directory to hold the image documents and a MySQL database that

stores the document details such as image file name, similarity values, and first order similarity. MySQL tables are also used to synchronize the working of the different components of I-FGM.

- **gIG-Builder:** This component processes the document in the gIG-Soup into a document graph that is compared with the query graph to calculate its similarity. The gIG-Builders incrementally process the images and calculate the image feature vector. The feature vector is compared with classifiers in the concept library and matching concepts selected. The descriptions of the matched concepts are combined to form the text description of the concepts. Now that we have a text description of images, we can convert it into a document graph using the same method that we use for a text document (Santos et al., 2005, 2006). The image retrieval algorithm was described in detail in the previous sub-section. gIG-Builders select an image based on the priority values and calculate the signature of a certain number of windows, which is again dependant on the image priority. Pixel windows (8x8) and sliding distance of four pixels are used. The feature vectors of the windows processed until that moment are clustered using the BIRCH clustering algorithm with a threshold value of 70. The centroids and areas of the clusters or regions are stored. The feature vector of the image is updated and compared with the classifiers in the library using the LR algorithm. The text description of the matched concepts are combined together to provide the text representation of the image.

- **I-Matcher:** After the images have been converted into a document graph, it is compared with the query graph. The query graph is produced when the query is converted into a document graph. After comparison of the two graphs, a similarity measure is calculated. It is a numerical measure of how close the image is to the query. The similarity is given by the formula:

$$\text{sim}(q, d_i) = \left(\frac{n}{2 * N} + \frac{m}{2 * M} \right) * \rho$$

$$\rho = \begin{cases} 1 - \frac{C_1 - C_2}{L} & \text{where } C_1 \geq C_2 \\ 1 - \frac{C_1}{L} & \text{where } C_1 < C_2 \end{cases}$$

Where q is the query graph, d_i is the document graph, n and m are the number of concepts and relation nodes respectively, of the query graph found in the document graph, and N and M are the total numbers of concept and relation nodes of the query graph. Two relation nodes are matched if and only if at least their parent and their child are matched; ρ is a probability measure of our confidence in the similarity; C_1 is the number of concepts matched to the image and C_2 is the number of possible relevant concepts. We calculate C_2 from the number of concepts in the query graph. Since the concepts that are matched to the image are not always correct, we use ρ as a weight to the similarity value. For some images, the matched concepts are varied; for example “*dinosaur*,” “*snowstorm*,” and “*fruits*.” These images happen to have pixel characteristics found in all these concepts. By using a confidence measure, we are quantifying how sure the machine learning algorithms is about its matches. If the matches are small in number, that images have certain unique pixel characteristics that are found in only certain concepts. The confidence measure will be high for these images. For images that match a large number of concepts, we say that the machine learning algorithm is not so confident about its matches.

The priority value for the image is also calculated after each partial step. The priority function is used to determine the order in which the documents in the gIG Soup will be processed. This guides the allocation of computing resources. The amount of processing done during a particular step is also determined by the priority. The prior-

ity formula used is identical to the one in Santos et al. (2006) except for the weights that were calculated for images, are used in the formula.

- **Blackboard:** Since the documents are partially processed, its similarity values changes. In addition to this, the search space is dynamic. The documents having the top similarity values will change over time. Blackboard will display the most recent results in a rolling fashion. Such a display is essential when real time search results are required.

DISCUSSION

Results and Analysis

We validate the performance of our image retrieval algorithm in three steps. In the first set of experiments, we validate the ability of the LR-based machine learning technique to match image features to relevant concepts. We do this by comparing with another machine learning technique called C5.0 (Quinlan, 1993). C5.0 is an established data mining tool for discovering patterns that delineate categories, assembling them into classifiers, and using them to make predictions. It does automatic classification using decision trees. In the field of image classification, it has been used in classification of satellite imagery (Hurd, Civco, Gilmore, Prislac, & Wilson, 2006) and protein crystallization (Zhu, Sun, Cheng, & Bern, 2004) images.

After we show the LR method to be an effective one in matching low-level features in images, we proceed to the second step in our validation, that is, validating the whole image retrieval algorithm. In our algorithm, we map regions in the image to known concepts. We now test if this matching is done effectively. We compare its performance with a previous prototype of I-FGM that implements WALRUS only. WALRUS is an established method and we have used WALRUS in I-FGM (Santos et al., 2007) to get good retrieval results. It may be noted that the concept-based I-FGM is built on top of the WALRUS sub-system. If we can show

that the precision of the concept-based I-FGM is similar or better than the WALRUS based I-FGM, we can then conclude that matching of concepts and low-level features is effective and thus validate our algorithm.

Finally, we have to validate whether the image retrieval algorithm coupled with the distributed processing approach of I-FGM helps in fast retrieval of images. We do this by comparing its performance with two control systems that represent traditional distributed methods.

Creation of Testbed

For conducting the experiments, an image testbed containing a sizeable number of images and a set of queries is created. We chose a scenario titled “Natural disasters” as disaster relief is one of the areas of e-government where quick retrievals are required in large and dynamic databases. We select a set of five queries from this scenario. They are:

1. Building damaged by hurricane Katrina,
2. Firefighters fight wildfires,
3. Heavy snow storms in the winter,
4. Houses damaged by tornado, and
5. Houses damaged by tsunami.

For each query, a testbed is created by running the three I-Foragers and downloading the top 50 results from each search engine.

Creation of Concept Library

Before the concept-based I-FGM can be run, a concept library has to be constructed. Building a comprehensive library is difficult and time consuming. Instead, we tailor our library to contain concepts that are related to the “Natural disasters” scenario.

The library must also contain unrelated concepts to provide a fair assessment of our image retrieval algorithm. We use a set of 22 concepts listed in Figure 5. Each concept in the library consists of a brief concept description and the concept classifier. Examples of concept descriptions are shown in Figure 6. For each concept, two sets of images (20 each) repre-

Figure 5. List of concepts in the concept library

Beach	Damaged buildings	Cold weather	Dinosaur	Firefighter
Fires	Fruits	Tribal People	Snow Storm	Mud Slide
Mountains	Trees	Floods	Forest fires	Rainstorm
Tornado	Elephants	Flowers	Food	Tsunami
Hurricane	Katrina			

Figure 6. Examples of concepts and concept description in the concept library

Concept	Concept description
1. Beach	Beach has sand. Beach has lines with coconut trees. Beach has a blue ocean. People sunbathe on the beach. Waves can be seen on the beach. People lie on the beach. People watch boats on the beach. There are boats on the beach. It is windy on the beach. People relax on the beach. People play on the beach.
2. Damaged buildings	Damaged buildings are caused by natural disasters. Damaged buildings are caused by flood, storm, and hurricane. Damaged buildings have damaged walls, leaking roofs, broken doors, broken windows, and so forth. Damaged buildings can also be caused by terrorism attack. For example, the World Trade Center was damaged by terrorist attack.




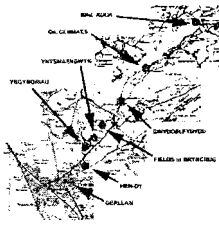
senting the positive and negative images are used to train the classifier. These images are processed to extract its feature vector consisting of the centroids of its constituent regions. The feature vectors are used to train a classifier in LR. Details about the LR technique are given in the previous sub-section. The concept library is stored in the memory of the IIG-Builders at the beginning of the experiment.

Validation of the Machine Learning Technique

In this section we validate the performance of LR-based machine learning technique implemented in the concept based I-FGM system by comparing it with C5.0. The existing studies in the machine learning community (Lim, Loh, & Shih, 2000; Perlich et al., 2003) have shown that LR performs competitively with other classifiers and performs better with small data sets than C4.5 (earlier version of C5.0). We would like to verify this finding on our own data set. The C5.0 algorithm was trained using positive

and negative images for each concept in the concept library and classifiers were created. Our LR technique and the C5.0 algorithm was used on the image testbed created for query 4: "Houses damaged by tornado." For each image the algorithm delivered matching concepts from the concept library. The relevant concepts for this query are "Damaged Buildings" and "Tornado." When the matched concepts of the two algorithms were compared, we observe that the LR technique matched the relevant concepts "Damaged Houses" and "Tornado" for more relevant images than the C5.0 algorithm. This can be observed in the results tabulated in Figure 7. Due to space limitation, results for only four randomly picked images are displayed. Three of the images (1, 2, & 3) are relevant and the last one (4) is not relevant. For all the images, we see that the LR method retrieves more matched concepts (marked in red in Figure 7) than C5.0. It may be observed that most of the matched concepts of LR are also semantically similar. For example, for image

Figure 7. Concepts matched by the logistic regression and C5.0 algorithms

Image	Matching concepts	
	Logistic regression	C5.0
1 	Damaged buildings, Dinosaur, Fruits, Snowstorm, Mountains, Flood, Tornado, Elephants, Flowers	Damaged- building, Dinosaur, Elephant, Firefighter, Flood, Fruit, Rainstorm
2 	Damaged buildings, cold weather, dinosaur, fruits, tribal people, food, mountains, trees, forest fires, rainstorm, tornado, flowers	Damaged- building, Dinosaur, Elephants, Firefighter, Flood, Fruit, Rainstorm
3 	Beach, Flood, Damaged buildings, Cold weather, Snowstorm, Mountains, Rainstorm, Tornado, Elephants	Beach, Damaged- building, Elephants, Flood, Mountains, Rainstorm
4 	Fruits, Trees, Tribal people, Snowstorm, Mountains, Forest fires, Elephants	Beach, Tree, Damaged-building, Elephants, Flood, Rainstorm, Snowstorm

1 in Figure 7, two of the matched concepts are flood and rainstorm. Though non-relevant, they are semantically related to Tornado as they all belong to the class of natural disasters. For the relevant images 1, 2, and 3, the LR algorithm consistently matches the relevant concepts while

C5.0 fails to match the concept "Tornado" to these images. Thus, we see that LR algorithm has a better performance than C5.0. In future work, user-based relevance feedback will be used to refine the matches made by the machine learning algorithm. Thus we have validated the

performance of the LR-based machine learning technique implemented in our image retrieval algorithm.

Validation of Image Retrieval Algorithm

In this section, we will validate the image retrieval algorithm implemented in I-FGM. In the algorithm we map the image features to high-level concepts and obtain a textual description of the images. The document graphs are then generated from the text descriptions. The document graphs are then compared against the text query to determine the similarity measures. (The image algorithm was described in detail earlier.) We will compare the performance of the concept-based I-FGM prototype with a purely WALRUS-based I-FGM prototype that we developed in Santos et al. (2007). By getting comparable or better results in concept-based I-FGM prototypes, we will show that using the concept mapping on top of WALRUS does not have an adverse effect on performance and we get just as effective results.

The WALRUS-based I-FGM prototype was the first step towards deploying image retrieval in I-FGM (Santos et al., 2007). It was developed to show that a regions-based image retrieval algorithm like WALRUS can be easily incorporated into the I-FGM. The prototype has the same system architecture as the concept-based I-FGM prototype. The image retrieval algorithm used in the WALRUS-based I-FGM is identical to the concept-based I-FGM except for two differences: (1) the WALRUS system uses query images instead of text queries, and (2) the WALRUS system compares the regions of the query image with that of the search space images, instead of extracting concepts from images. I-matcher extracts the image feature vector from the query image and compares it with image feature vector of the search space image. The similarity measure used is

$$\text{Similarity}(P, Q) = \frac{\text{area}(\cup_{i=1}^n (P_i)) + \text{area}(\cup_{j=1}^m (Q_j))}{\text{area}(P) + \text{area}(Q)}$$

Here P and Q represent the retrieved image and the query image respectively. The set of ordered pairs $\{(P_1, Q_1), \dots, (P_n, Q_n)\}$ form a similarity region pair set for P and Q , where P_i is similar to Q_j and for $i \neq j$, $P_i \neq P_j$, $Q_i \neq Q_j$.

Both WALRUS-based and concept-based I-FGM prototypes are run with the testbed of five queries. Precision (Salton & McGill, 1983) is the performance metric that we use to compare the performance of the concept-based and WALRUS-based I-FGM system. It is a commonly used metric in information retrieval and is defined as the ratio of the relevant documents (images) retrieved to the total documents (images) retrieved. In this article we use a slightly different definition: ratio of the relevant documents (images) among the top n retrieved documents (images). The two systems are run on each of the query testbed and the top 15 images for each query are analyzed to determine their relevance and finally calculate the precision.

Analysis of the Validation Results

The top 15 images retrieved by the WALRUS-based I-FGM prototype for each query are analyzed to determine if they are relevant or not. We do this using certain criteria. For example, in query 1: "Building damage by hurricane Katrina," we only consider images that depict damaged buildings, especially on the coast as relevant images. The precision results obtained by the two systems for all the queries are tabulated in Figure 8. For lack of space we present the top 15 image results for only query 1 (Figures 9a and 9b). Each image is tagged as $x.y$ where x is the name of the search engine that downloaded the image and y is the rank provided by it. The images deemed relevant are indicated by a check mark in Figure 9. A quick glance shows that the performances of the systems are comparable as they retrieve around the same number of relevant documents in each query. The concept-based and WALRUS-based systems win in two queries each (queries 1 & 3 for the concept-based system and queries 2 & 4 for the WALRUS-based system) and have a tie for query 5. From these results, we can

come to the conclusion that our concept-based image retrieval algorithm performs as well as WALRUS. It may be noted here that the concept mapping part of our current prototype is built on top of the WALRUS feature extraction subsystem. Relevant concepts are matched after the WALRUS sub-system extracts the regions in the images. The document graphs are then generated and compared with the query. We use this to provide a unified framework for heterogeneous data. Similar performance results for both WALRUS and concept-based I-FGM show that the performance is not degraded when concepts matching are used with the feature extraction. This also leads us to conclude that the regions-based feature extraction methodology can indeed be used to effectively map low-level image features to higher level concepts.

Evaluation of Image Retrieval in I-FGM

In this section we will validate the effectiveness of I-FGM for image retrieval. We achieve this by comparing its performance with two other control systems that represent the traditional information and distributed processing paradigms.

The control systems differ from I-FGM only in the way the documents are chosen for processing. The control systems are baseline and partially intelligent systems. In the baseline system, the documents are chosen at random for processing. Each document is given the exact

same chunk of time for processing in each step. In the partially intelligent system, the priority of a document is static and is equivalent to the first-order similarity. We use the image testbed created previously to conduct the simulation experiments. The images are processed by the gIG builders and the final similarity measure recorded. The top n documents are selected as the target document set for each query.

The control systems and I-FGM systems are run on the testbed for each query. Two performance metrics calculated during each run are: (1) recall vs. time, and (2) document waiting times. Recall (Salton & McGill, 1983) is defined as the ratio of the relevant documents retrieved to the total number of relevant documents. We designate the top x documents selected by our image retrieval method as the set of relevant documents. We store the recall values for each control system at regular intervals throughout the period of simulation. The waiting time for a document is defined as the time spent in the gIG Soup before it is displayed on the blackboard. Since one of the goals of I-FGM is to get to the relevant documents as quickly as possible, the waiting time is an important metric.

RESULTS AND ANALYSIS

As mentioned before, the control systems and I-FGM are run on the test beds of the five queries. The waiting times of relevant documents and recall values are monitored during the simulations.

Figure 8. Precision values attained by WALRUS-based and concept-based I-FGM

	Precision	
	WALRUS-based I-FGM	Concept-based I-FGM
Query 1	9/15	10/15
Query 2	10/15	9/15
Query 3	9/15	11/15
Query 4	13/15	12/15
Query 5	10/15	10/15

Figure 9a. Precision results for query1 using WALRUS-based I-FGM



It has the shortest waiting time for the most documents in queries 4 and 5. For query 3, it ties with baseline system to get the most number of documents in the shortest time. Overall, I-FGM is best in three queries, baseline in two queries, and partial in one query.

We now carefully analyze the results of each query. In query 1 (Figure 10a), baseline system is the best performing system getting 6 out of the 11 documents the fastest. Partial comes

second getting four documents and I-FGM is third with two documents. When we analyze the documents closely, we see that I-FGM performs poorly for this query for one primary reason. The similarity measure fluctuates, decreasing in one step and increasing in another.

In I-FGM, a priority function (used for selecting from the gIG-Soup) decreases the priority of an image when its similarity at a particular step does not increase proportional

Figure 9b. Precision results for query1 using concept-based I-FGM



to the processing time in that particular step or when the similarity decreases. Documents with rapid fluctuations in similarity do not work well with this priority function. Although we used an initial set of experimental results to determine the weights in the priority function, more work needs to be done in formulating a better priority function. As we obtain more experimental results, the priority function will be refined and this will be part of future work. From the precision graph (Figure 10b), we see that I-FGM has a better recall for part of the simulation time. For query 2 (Figures 11a and

11b), partial system has lowest waiting times for four documents while I-FGM and baseline system comes second with three documents each. We can easily see why this is the reason. The relevant documents are highly ranked by the I-Foragers, that is, they have a high expected first order similarity. Since the partial system processes the documents based only on the expected first order similarity, such documents are most likely to be picked by it sooner than the other systems. This is further helped by the fact that documents such as MSN.0 and Google.0 are small (around 2,000 windows). But it may

Figure 10a. Document waiting time—Query 1




Figure 10b. Recall vs. time—Query 1

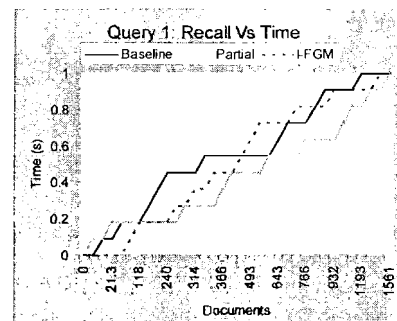


Figure 11a. Document waiting time—Query 2


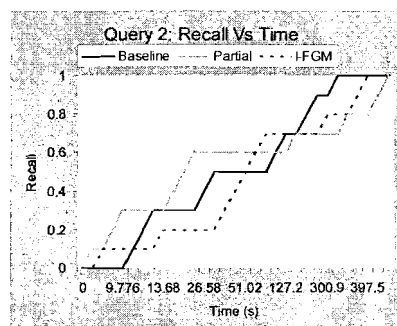


Figure 11b. Recall vs. time—Query 2



be noted that the performance of I-FGM is not much worse than partial as it gets only one document less than the partial systems.

For query 3 (Figures 12a and 12b), both I-FGM and baseline systems perform the best, getting four documents each. Partial system comes in last with one document. When we analyze three documents (Yahoo.41, Google.53, Google.35), which is picked by baseline system the earliest, we see that these documents have a high similarity value after the first partial step itself. Since they are lowly ranked by the I-foragers, they have a low initial priority in the I-FGM system. Therefore I-FGM gets these documents only after a certain amount

of simulation time has passed. Baseline system gets lucky by selecting these documents early. But since these documents give a high similarity value in the first processing step, they appear in the blackboard after the first processing step. I-FGM has better performance than baseline and I-FGM for queries 4 (Figures 13a and 13b) and 5 (Figures 14a and 14b). For query 4, I-FGM gets 5 out of the 8 relevant documents. Baseline system comes in second with three documents. In the recall graph also, I-FGM has a higher recall value than the other systems for most of the simulation time. For query 5, I-FGM gets 5 documents out of 11 target documents. Partial and baseline tie for the second place with 4 each.

Figure 12a. Document waiting time—Query 3




Figure 12b. Recall vs. time—Query 3




Figure 13a. Document waiting time—Query 4




Figure 13b. Recall vs. time—Query 4




Figure 14a. Document waiting time—Query 5





Figure 14b. Recall vs. time—Query 5



In the recall graph (Figure 15b), I-FGM has a better recall than other systems over the range (0.2-0.7). The similarity values for images can decrease over time (which does not happen in text retrieval) and fluctuate rapidly. These factors make the modeling of priority function for images a difficult task. Despite this, we have shown I-FGM to be better in the majority of queries, thus proving its feasibility as a quick and reliable retrieval method for images.

CONCLUSION

In this article we designed and implemented an image retrieval algorithm based on mapping low-level features of images to higher level concepts. We validated this algorithm on I-FGM and showed that it can be used to represent the visual information of images as concepts and relations. Consequently, we also validated the methodology in I-FGM of providing a unified ranking for heterogeneous data types. I-FGM was also compared against two control systems to show that the image retrieval algorithm works well with partial processing and helps in faster retrieval of images.

The simulations in this article were conducted on a static testbed. In the future, this will be extended to include simulation studies for different dynamic conditions in the search space. Furthermore, with our single unifying representation of semantic content in text and image via our document graphs, I-FGM will be readily modified for a mixed document search space containing both text and image documents. Documents that contain both text and image data will also be considered. Currently, I-FGM does not request any user input except the user's queries. We plan to incorporate relevance feedback to allow users to give their inputs on the retrieval quality. Users can be asked explicitly to indicate which images are most similar to what they are looking for. We may also observe the user's interactions with our system to infer implicitly which images might be of their interests. These retrieved relevant images then will be used to modify the user's original queries by adding both textual and image information to them. This process will

help guide the retrieval task so that it returns more relevant images to the user.

ACKNOWLEDGMENT

The work was supported in part by the National Geospatial Intelligence Agency Grant Nos. HM1582-04-I-2027 and HM1582-05-I-2042. We would also like to thank Dr. Qunhua Zhao for his various comments and advice in this project.

REFERENCES

Caenen, G., & Pauwels, E. J. (2002). *Logistic regression models for relevance feedback in content-based image retrieval*. San Jose, CA: Storage and Retrieval for Media Databases.

Carson, C., Belongie, S., Greenspan, H., & Malik, J. (2002). Blobworld: Image segmentation using expectation-maximization and its application to image querying. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(8), 1026-1038.

Daubechies, I. (1992). *Ten lectures on wavelets*. Philadelphia: SIAM.

Enser, P., & Sandom, C. (2003). Towards a comprehensive survey of the semantic gap in visual image retrieval. (LNCS 2728, pp. 279-287).

Hurd, J. D., Civco, D. L., Gilmore, M. S., Prisloe, S., & Wilson, E. H. (2006, May 1-5). *Tidal wetland classification from landsat imagery using an integrated pixel-based and object-based classification approach*. Paper presented at the ASPRS 2006 Annual Conference, Reno, NV.

Kherfi, M. L., Ziou, D., & Bernardi, A. (2004). Image retrieval from the World Wide Web: Issues, techniques, and systems. *ACM Computing Surveys*, 36(1), 35-67.

Komarek, P. (2004, May). *Logistic regression for data mining and high-dimensional classification* (Tech. Rep. No. CMU-RI-TR-04-34). Pittsburgh, PA: Carnegie Mellon University, Robotics Institute.

Ksantini, R., Ziou, D., Colin, B., & Dubeau, F. (2007). *Logistic regression models for a fast CBIR method based on feature selection*. Paper presented at the International Joint Conference on Artificial Intelligence.

Li, J., & Wang, J. Z. (2006). *Real-time computerized annotation of pictures*. Paper presented at the ACM Multimedia Conference. Santa Barbara, CA.

Li, J., Wang, J. Z., & Wiederhold, G. (2000). *IRM: Integrated region matching for image retrieval*. Paper presented at the ACM Multimedia Conference, Los Angeles, CA.

Lim, T. S., Loh, W. Y., & Shih, Y. S. (2000). A comparison of prediction accuracy, complexity, and training time for thirty-three old and new classification algorithms. *Machine Learning*, 40, 203-228.

Ma, W. Y., & Manjunath, B. (1997). *NaTra: A toolbox for navigating large image databases*. Paper presented at the IEEE International Conference on Image Processing.

Mitchell, T. (2005). *Generative and discriminative classifiers: Naïve Bayes and logistic regression*. Retrieved from <http://www.es.cmu.edu/~tom/ml-book/NBayerLogReg.pdf>

Natsev, A. (2001). *Multimedia retrieval by regions, concepts, and constraints*. Unpublished PhD thesis, Duke University, Durham, NC.

Natsev, A., Chadha, A., Soctarman, B., & Vitter, J. S. (2001). *CAMEL: Concept annotated image libraries*. Paper presented at the Storage and Retrieval for Image and Video Databases, SPIE, San Jose, CA.

Natsev, A., Rastogi, R., & Shim, K. (2004). WALRUS: A similarity retrieval algorithm for image databases. *IEEE Transactions on Knowledge and Data Engineering*, 16(3), 301-316.

Perlich, C., Provost, F., & Simonoff, J. (2003). Tree induction vs. logistic regression: A learning-curve analysis. *Journal of Machine Learning Research*, 4, 211-255.

Quinlan, J. R. (1993). *C4.5: Programs for machine learning*. San Francisco: Morgan Kaufmann.

Rui, Y., Uluang, T., & Chang, S. (1999). Image retrieval: Current techniques, promising directions and open issues. *Journal of Visual Communication and Image Representation*, 10(4), 39-62.

Salton, G., & McGill, M. (1983). *Introduction to modern information retrieval*. McGraw-Hill.

Santos, E., Jr., Santos, E., Nguyen, H., Pan, L., & Korah, J. (2005). Large-scale distributed foraging, gathering, and matching for information retrieval:

Assisting the geospatial intelligent analyst. In *Proceedings of the SPIE* (Vol. 5803, pp. 66-77).

Santos, E., Jr., Santos, E., Nguyen, H., Pan, L., Korah, J., Zhao, Q., et al. I-FGM information retrieval in highly dynamic search spaces. In *Proceedings of SPIE* (Vol. 6229, pp. 1-12).

Santos, E., Jr., Santos, E., Nguyen, H., Pan, L., Korah, J., Zhao, Q., et al. (2007). Applying I-FGM to image retrieval and an I-FGM system performance analyses. In *Proceedings of the SPIE: Defense & Security Symposium* (Vol. 6560), Orlando, FL.

Smeulders, A. W. M., Worring, M., Santini, S., Gupta, A., & Jain, R. (2000). Content-based image retrieval at the end of the early years. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(12), 1349-1380.

Tsybalenko, Y., & Munson, E. V. (2001). *Using HTML metadata to find relevant images on the Web*. Paper presented at the Internet Computing, Las Vegas, NV.

Vapnik, V. (1999). *The nature of statistical learning theory* (2nd ed.). Springer Verlag.

Veltkamp, R. C., & Tanase, M. (2000). *Content-based image retrieval systems: A survey* (Tech. Rep. No. UU-CS-2000-34). The Netherlands: Utrecht University, Department of Computer Science.

Wang, J. Z., Li, J., & Wiederhold, G. (2001). SIMPLcity: Semantics-sensitive integrated matching for picture libraries. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(9), 947-963.

Witten, I. H., & Frank, E. (2005, June). *Data mining: Practical machine learning tools and techniques* (2nd ed.). Morgan Kaufmann.

Zhang, J., Jin, R., Yang, Y., & Hauptmann, A. G. (2003). Modified logistic regression: An approximation to SVM and its application in large-scale text categorization. In *Proceedings of the Twentieth International Conference on Machine Learning 2003*, Washington DC.

Zhang, T., Ramakrishnan, R., & Livny, M. (1996). *BIRCH: An efficient data clustering method for very large databases*. Paper presented at the ACM SIGMOD International Conference on Management of Data, Montreal, Canada.

Zhu, X., Sun, S., Cheng, S. E., & Bern, M. (2004). *Classification of protein crystallization imagery*. Paper presented at the 26th Annual International Conference of IEEE Engineering in Medicine and Biology Society, San Francisco.

Eugene Santos, Jr. is a professor of engineering in the Thayer School of Engineering at Dartmouth College. He received his BS ('85) in mathematics and computer science from Youngstown State University, a MS ('86) in mathematics (specializing in numerical analysis) from Youngstown State University, as well as ScM ('88) and PhD ('92) degrees in computer science from Brown University. His areas of research interest include artificial intelligence, intent inferencing, information retrieval, automated reasoning, decision science, adversarial reasoning, user modeling, natural language processing, probabilistic reasoning, and knowledge engineering, verification and validation, protein folding, load balancing, virtual reality, and active user interfaces. He has served on many major conference program committees from intelligent agents to evolutionary computing. He is currently an associate editor for the IEEE Transactions on Systems, Man, and Cybernetics and for the International Journal of Image and Graphics, and is also on the editorial advisory board for System and Information Sciences Notes.

Eunice E. Santos has BS and MS degrees in both mathematics and computer science, and received a PhD in 1995 from the University of California, Berkeley in computer science. Dr. Santos is the director of the laboratory for Computation, Information & Distributed Processing, and an associate professor in the Department of Computer Science at Virginia Polytechnic Institute & State University. Her research interests include parallel and distributed processing, computational science and engineering, intelligent systems, networking, and modeling & simulation. Dr. Santos has been awarded an NSF CAREER grant, the Spira Award for Excellence in Teaching in Engineering, and the Robinson Faculty Award. She is subject area editor for the Journal of Supercomputing and a member of the DARPA/IDA Defense Science Study Group.

Hien Nguyen is an assistant professor at the Department of Mathematical and Computer Sciences at the University of Wisconsin-Whitewater. She received a PhD in computer science from the University of Connecticut in 2005. Her research interests include user modeling, information retrieval, collaborative information retrieval, recommender systems with a current focus on hybrid user model for improving a user's performance in information retrieval. Professional services and committee work include program committee for 2007 and 2006 IEEE International Conference on Systems, Man, and Cybernetics (IEEE SMC2007 and SMC2006). She also is a member of User Modeling-User Adapted Interactions journal Special Reviewers Board. She also serves as a reviewer for User Modeling conference (2007), IEEE Transactions on Systems, Man, and Cybernetics, Part B, Information Processing and Management, and AAAI 2006 (Poster).

Long Pan is a doctoral candidate in the LCID (Laboratory for Computation, Information & Distributed Processing) of the Department of Computer Science at the Virginia Polytechnic Institute & State University. He received his BS and MS degrees from the automation department of the Tsinghua University. His research interests are focused on SNA (Social Network Analysis) on large and dynamic networks, SNA with culture infusion, parallel/distributed computing, data mining, and information retrieval. He serves as a reviewer for Journal of Supercomputing, and the International Conference on Parallel Processing (2007).

John Korah is a doctoral candidate in the Department of Computer Science at the Virginia Polytechnic Institute & State University. He received his BE degree in electronics and instrumentation engineering from the Govt. College of Technology, Coimbatore, India. He has an master's degree in electrical engineering from Virginia Tech. His research interests are focused on parallel/distributed computing, numerical algorithms, large scale network modeling/simulation and information retrieval. He also serves as a reviewer for the Journal of Supercomputing.

Huadong Xia is a graduate student in the Department of Computer Science at the Virginia Polytechnic Institute & State University. He received his BS degree in information science from Peking University and got his ME degree in computer science from the Chinese Academy of Sciences. His research interests include distributed computing, information retrieval and social networks analysis. He also serves as a reviewer for the Journal of Supercomputing.