Expert Knowledge Based Automatic Regions-of-Interest (ROI) Selection in Scanned Documents for Digital Image Encryption

Alexander Wong
a28wong@engmail.uwaterloo.ca

William Bishop
wdbishop@uwaterloo.ca

Department of Electrical and Computer Engineering
University of Waterloo
Waterloo, Ontario, Canada

Abstract

Conventional image-oriented cryptographic techniques lack the flexibility needed for content-specific security features such as the concealment of confidential information within a portion of a document. Content-specific security is particularly important for digital archival systems that store sensitive documents in the form of digital images. Recently, a novel image encryption scheme utilizing multiple levels of regions-of-interest (ROI) privileges for digital document encryption was developed to address the needs of modern digital document management systems. This image encryption scheme requires the selection of regions-of-interest for encryption. The process of manually selecting regions can be time-consuming. This paper presents an automatic, regions-of-interest selection algorithm that utilizes an expert knowledge learning system to select regions of interest in a scanned document image for the purpose of minimizing human interaction time during the encryption process. Experimental results show that a high level of accuracy and significant timesaving benefits can be achieved using the proposed algorithm.

1. Introduction

Encryption is required to protect digital documents during storage and transmission. While conventional cryptographic techniques are useful for protecting entire files, such techniques are not suitable for encrypting portions of scanned document images. Also, the encryption of entire scanned document images requires a substantial amount of processing time.

Image-oriented algorithms have been developed to provide better overall performance. However, these algorithms lack the level of security provided by conventional techniques and lack the flexibility needed for content-specific security such as the concealment of confidential information within a document. These issues prevent the use of image-oriented encryption techniques for the protection of portions of confidential documents in a digital document archival and distribution system. This is particularly true when paper-based documents are scanned and archived in the form of digital images. Recently, a novel image encryption scheme [18] was developed based on the concept of selected regions-of-interest (ROI). This multi-level digital document encryption scheme addresses the needs of modern digital document management systems. Using this approach, an individual selects the confidential regions of interest to be concealed in an image and the encryption system encrypts these areas based on the level of authority that is required to view portions of the document.

Using a multi-level digital document encryption scheme, an encrypted image of a confidentiality agreement can be viewed with different levels of authority. The general public is able to see the public content of the document to allow for document content browsing and content-based document searching. Individuals with a level-1 security clearance are allowed to access portions of the hand-written areas to validate that the agreement has been completed properly. Individuals with a level-2 security clearance are able to read the entire document including the signature portions of the document to allow for signature verification.

The encryption algorithm requires the selection of ROI for encryption. This process can be time-consuming if done manually without any assistance. The human selection of ROI is generally the bottleneck in the ROI encryption process. The automation of ROI for encryption can yield significant time savings and it is highly desirable.

A number of ROI selection algorithms have been introduced in the area of medical imaging for isolating specific areas in DPA [16], X-ray/CT/PET images [7] [9] [11],

Proceedings of the 3rd Canadian Conference on Computer and Robot Vision (CRV'06)
0-7695-2542-3/06 $20.00 © 2006 IEEE

renograms [1], and mammograms [14] [15]. ROI selection algorithms have also been explored for the analysis of aerial and satellite images [4] [12], content authentication [8] [6], and textual analysis [17]. A number of issues arise when the aforementioned algorithms are viewed in the context of ROI digital document encryption. These algorithms are designed for a specific task and have domain-specific dependencies that make them unsuitable for use in ROI selection for digital document image encryption. Furthermore, a word or phrase may need to be encrypted in one document but not in another document. The process of automatically selecting specific text in a digital document image for encryption seems practically impossible.

However, two important characteristics common to archived digital documents can be exploited to allow for ROI selection in digital document images. First, the digital documents stored in document management systems are typically standardized documents with well-defined document structures. The overall structure of a standardized document can be easily identified and spatial relationships can be established. Second, the digital documents stored in document management systems are typically produced by scanning paper documents using flatbed scanners. Thus, the image disparities between two documents of the same type are primarily Euclidean transformations consisting of minor rotations and translations. These Euclidean transformations can be identified using pattern recognition techniques. Thus, it is possible to implement an effective and practical automatic ROI selection system to help reduce human interaction time.

This paper describes a new automatic, region-of-interest (ROI) selection algorithm that utilizes an expert knowledge learning system to select regions of interest in a standardized, digital document image for the purpose of minimizing human interaction time during the digital document encryption process. In this paper, the proposed selection algorithm is described and explained in detail in Section 2. Experimental results are presented in Section 3 along with an analysis of the accuracy of the proposed algorithm. Section 4 concludes with a discussion of the merits of the proposed algorithm.

2. Proposed Automatic ROI Selection Algorithm

The algorithm proposed in this paper accomplishes automatic ROI selection in a document with the use of an expert knowledge learning system. It is possible to use past knowledge learned from human expert input about similar document images to automatically determine the appropriate set of ROI for an untrained document image given that a sufficient set of representative document image samples are trained on the knowledge learning system. This set of ROI can then be used to accelerate the human interaction process for ROI image encryption algorithm by reducing the amount of interaction needed from a human expert.

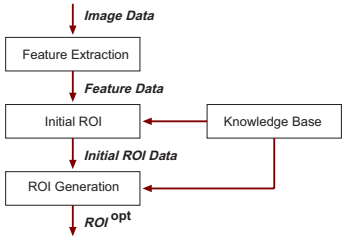

Figure 1. Automatic ROI Selection Algorithm

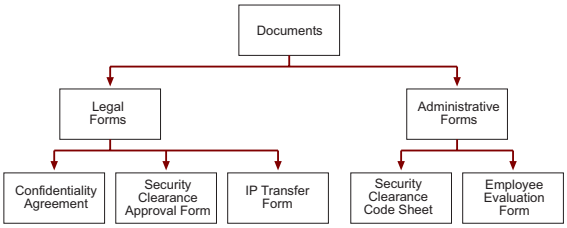

Figure 2. Sample ROI Knowledge Learning System Structure

2.1. Overview

Given an initial uncompressed raster image ($Image$), the proposed algorithm finds an approximate set of ROI (ROI^{opt}) based on prior knowledge instances from a knowledge learning system. A general overview of the automatic ROI selection algorithm is shown in Figure 1.

2.2. ROI Knowledge Learning System

A hierarchical N-ary tree structure acts as the foundation for the ROI knowledge learning system. Learned image instances are categorized by their document structures. A sample of a ROI knowledge learning system structure for a set of documents is shown in Figure 2.

The hierarchical tree structure that serves as the basis of the knowledge learning system architecture permits the rapid identification of ROI since the number of knowledge instances that must be compared against the input image is small. Even for large tree structures, the number of comparisons will be reasonable. Each of the leaf nodes in the knowledge tree corresponds to a set of knowledge instances that fall under a category, known as a knowledge base. Each knowledge instance in the knowledge base represents a separate image that has been "learned" by the knowledge sys-

Table 1. Knowledge Instance Data Fields

id	Unique knowledge instance identifier
$IMAGE_{id}$	Scaled-down, inverted grayscale version of the original document image data
$COARSE_{id}$	Coordinates of the image features used during coarse and fine image registration
$FINE_{id}$	Coordinates of the additional image features used during fine image registration
ROI_{id}	ROI specification data of the image

tem. The data fields associated with knowledge instances are defined in Table 1.

The ROI specification data stored in a knowledge instance consists of the top-left and bottom-right corner coordinates of the rectangular ROI in the image, as well as the corresponding level of authority required to view each region. The image feature information stored in a knowledge instance is used for image registration. This information varies depending on the registration technique used. For standardized documents, two sets of feature points are stored. A small set of corner feature points are used for both coarse-grained image registration and fine-grained image registration. A large set of additional corner feature points are used for fine-grained image registration. Figure 3 illustrates a sample knowledge instance. Corner feature points are used for image matching. Standardized documents consist primarily of lines, boxes, and textual information with well-defined fonts for information clarity. The structural characteristics of these elements can be well represented by their corners. For example, a box element can be represented in a concise manner by its four corners and the letter A can be defined by its five corner points. Therefore, corner feature points can be used effectively for matching such document images.

The knowledge learning system can be viewed as a dynamic rule base consisting of `IF-THEN` rules. These rules have the following form:

$$\text{IF } (x_1 \text{ AND } x_2 \text{ AND } \ldots \text{ AND } x_n \text{ AND } y) \text{ THEN } z$$

The antecedents x_1 to x_n represent the categories in the learning system, the antecedent y represents an image, and the consequent z represents the multiple levels of ROI associated with the image. The training process can be treated as the addition of new rules to the existing rule base, while the initial ROI identification process is simply a search for the best match in the rule base. An example of a "rule" in this system is the following:

$$\text{IF } (\textit{Legal Form}) \text{ AND } (\textit{Confidentiality Agreement}) \text{ AND } (IMAGE_i) \text{ THEN } (ROI_i)$$

id	1
$IMAGE_1$	150×200 grayscale image data

$COARSE_1$	Feature	Coordinates
	1	{(10, 10)}
	2	{(20, 20)}
	3	{(3, 90)}

$FINE_1$	Feature	Coordinates
	1	{(30, 16)}
	2	{(50, 30)}
	3	{(16, 90)}

ROI_1	Region	Authority Level	Coordinates
	1	1	{(5, 6), (10, 15)}
	2	1	{(30, 9), (40, 30)}
	3	2	{(60, 60), (70, 70)}
	4	2	{(75, 75), (90, 90)}
	5	3	{(50, 50), (60, 60)}

Figure 3. Sample Knowledge Instance

The first two antecedents representing categories have binary values of 0 or 1. The last antecedent representing an image as well as the consequent representing set of ROI are fuzzy with values ranging between 0 and 1.

2.3. Training Stage

At this stage, the knowledge learning system is trained through expert interaction. The human expert inputs an image into the knowledge learning system and specifies the ROI within the image, along with their level of authority and the type of document. For example, the human expert may specify a number of ROI with level of authority of 1 around the signature areas in the input digital document image and select the document type as "Confidentiality Agreement". A unique identifier denoted by id is generated for the knowledge instance and ROI specification data ROI_{id} is generated from the image based on the ROI and the level of authority specified by the human expert. A downscaled, inverted grayscale version of the image $IMAGE_{\text{id}}$ of $m_{\text{ref}} \times n_{\text{ref}}$ pixels is generated as well. Image feature data is generated from the image by extracting two sets of corner feature points ($COARSE_{id}$ and $FINE_{id}$) from image using a corner detection technique (such as the Harris corner detector [3] or the SUSAN corner detector [13]). For the purpose of the proposed algorithm, a modified Harris corner detector as described in the thesis by J. Noble [10] was used. In this corner detection algorithm, the corner strengths are calculated using the following formula:

$$C(x,y) = \frac{(\nabla_{x^2}(x,y)\nabla_{y^2}(x,y) - \nabla^2_{xy}(x,y))}{(\nabla_{x^2}(x,y) + \nabla_{y^2}(x,y))} \quad (1)$$

$\nabla_{x^2}(x,y)$ and $\nabla_{y^2}(x,y)$ are Gaussian-smoothed squared derivatives in the x and y directions respectively. $\nabla_{xy}(x,y)$ is the Gaussian-smoothed cross-derivative. The corner positions in the image are then determined by finding points where the corner strength is a local maximum in a $n \times n$ neighborhood and the corner strength is greater than a predefined threshold of t. For the proposed algorithm, a 3×3 neighborhood is used and thresholds of $t = 500$ and $t = 5000$ are used to select corner feature points for fine-grained and coarse-grained image registration respectively. The image feature points are extracted for the entire image in a way that ensures enough feature points for accurate image registration. The id, ROI_{id}, $COARSE_{id}$, and $FINE_{id}$ of the trained image are then stored in a knowledge instance which is then linked to the knowledge base that corresponds to the specified category.

The training algorithm extracts two sets of image feature points from the existing image and stores them along with the user-specified ROI and a downscaled, inverted grayscale version of the image to create the knowledge instance. This creates a reasonably concise knowledge representation of the image in the knowledge base that helps minimize the amount of computations needed during the ROI selection phases of the algorithm.

2.4. Initial ROI Identification Stage

At this stage, the input image under evaluation (and optionally the document category) is entered into the ROI selection system. A downscaled, inverted grayscale version of image ($IMAGE_R$) with a size of $m_{\text{ref}} \times n_{\text{ref}}$ pixels is generated. Two sets of image feature points $COARSE_R$ and $FINE_R$ are generated from the image by extracting the features from $IMAGE_R$ using the same feature detection algorithm as that used in the training stage.

The set of searchable knowledge instances are determined based on the document category that was specified. If no document categories were specified, then all instances in the knowledge learning system are marked as searchable. If a category is specified, then all the instances of all the knowledge bases that fall under that category are marked as searchable. One of the primary drawbacks to the proposed algorithm from a computational performance perspective is that all knowledge instances must be evaluated if the document category is completely unknown. However, this is seldom the case in digital document management systems as the document images are standardized documents of known types. Therefore, the document type of the input image is known to a reasonable extent and therefore reduces the computational requirements to a minimal.

For each knowledge instance i from the set of searchable knowledge instances, a fitness rating is calculated depending on how well $IMAGE_i$ matches $IMAGE_R$. The image correlation technique used is based on a publically available approach [5]. First, the input image and the image from the knowledge instance are smoothed using an 11×11 average filter and subtracted from the original corresponding images to compensate for brightness differences in the images. The normalized cross-correlation between all feature points in $COARSE_R$ from the input image and all feature points in $COARSE_i$ from the knowledge instance i within a search radius of $\frac{1}{3}\min(m_{\text{ref}}, n_{\text{ref}})$ is then calculated using a 11×11 correlation window to obtain the corresponding correlation strengths between two feature points denoted as follows:

$$r(p_1,p_2) = \frac{\Sigma_x \Sigma_y (I_{p_1}(x,y) I_{p_2}(x,y))}{\sqrt{\Sigma_x \Sigma_y (I_{p_2}(x,y))^2}} \quad (2)$$

where p_1 and p_2 are feature points in the input image and the knowledge instance respectively, and $I_{p_1}(x,y)$ and $I_{p_2}(x,y)$ are the grayscale intensity values at (x,y) in the correlation window centered at p_1 and p_2 respectively. The matching feature point pairs are determined by searching for the instances where the most correlated feature points are two-way consistent. For example, if the strongest correlation strength of p_1 corresponds to p_2 and the strongest correlation strength of p_2 corresponds to p_1 then the feature points are considered a matching pair. Finally, a fitness function denoted by F is calculated as the sum of correlation strengths of the matching pairs:

$$F = \Sigma(r(p_1,p_2)), \text{ for all matching pairs} \quad (3)$$

where p_1 and p_2 are the points in a matching pair.

Finally, the knowledge instance from the set of searchable instances with the highest fitness rating is selected as the best match for $IMAGE_R$ and the initial set of ROI is set to that of the matched knowledge instance.

It is important to discuss the significance of using inverted pixel intensities as opposed to the original grayscale intensities for the purpose of image matching. A standardized document image consists primarily of dark textual content on a white background. Given the relatively sparse nature of the textual content as well as the dependence of pixel intensities in the correlation function, the background pixels become the dominant factor in the image matching process if the original pixel intensities are used. This diminishes the discrimination capabilities of the image matching process for such document images. By using the inverse pixel intensities, more emphasis is placed on the actual textual content pixels in the correlation function resulting in better discrimination between different document images.

The ROI identification algorithm uses a number of techniques to improve overall ROI identification performance. First, downscaled grayscale images are used in the document matching process as opposed to the full-size originals. This serves two important purposes:

1. By taking feature points from the downscaled image as opposed to the full-size originals, the number of feature points needed is reduced, as less characteristic feature points found in the original image are lost during the scaling process.

2. Scaling images to the same reference size allows them to be compared directly with other images scaled to the same size. It is often the case where documents are scanned and stored at different resolutions if standard practices are not enforced.

Secondly, the search radius in the normalized correlation calculation process is limited to $\frac{1}{3}\min(m_{\text{ref}}, n_{\text{ref}})$ as opposed to the entire image. This is because the disparities between documents of the same type are relatively minor and therefore the search radius reduction results in a good tradeoff of accuracy for performance. Thirdly, the use of smaller sets of feature points compared to that used in the final ROI generation stage help improve image matching performance, which is important at this stage as the input image can potentially be matched against all knowledge instances in the system if a category is not specified. Finally, the option to specify a category for the image under evaluation reduces the number of knowledge instances that need to be evaluated. Since the document category of the image under evaluation is typically known to a certain extent, this pruning of searchable knowledge instances improves the overall search performance of the system.

2.5. Final ROI Generation Stage

At this stage, the knowledge instance i selected in the previous stage is further refined to create a final set of ROI ROI^{opt} for the $IMAGE_R$. A set of n matching feature point pairs $MP = \{(p_{(1,1)}, p_{(2,1)}), (p_{(1,2)}, p_{(2,2)}), \cdots, (p_{(1,n)}, p_{(2,n)})\}$ is determined by applying the aforementioned normalized correlation matching technique between the set of input feature points ($\{FINE_R,\ COARSE_R\}$) and the set of feature points in knowledge instance i ($\{FINE_i,\ COARSE_i\}$). For this purpose, a correlation window of 11×11 within a search radius of $\frac{1}{3}\min(m_{\text{ref}}, n_{\text{ref}})$ is used. A 3×3 homography T is fitted between the $IMAGE_R$ and $IMAGE_i$ using the Random Sample Consensus (RANSAC) algorithm [2] such that:

$$p_{(1,j)} = T \times p_{(2,j)},\ \text{for all matching pairs} \tag{4}$$

where $p_{(1,j)}$ and $p_{(2,j)}$ are feature points in matching pair j.

In the RANSAC algorithm, a number of data points are selected at random from the sample set and used to fit an estimated model M. For the proposed algorithm, the data points are the feature points and the model is the 2-D homography. The number of data points that fit the estimated model M within a threshold t is then determined. The algorithm is repeated K more times, where K is the number of iterations needed such that there is an estimated probability of ρ that all the available data points fit the estimated model and is calculated using Equation 5. For the purpose of this algorithm, a threshold of $t = 0.002$ and a probability of $\rho = 0.99$ was used.

$$K = \frac{\log(1-\rho)}{\log(1-(n_{\text{inliers}}/N)^s)} \tag{5}$$

where ρ is the desired probability that at least one model exists that fits all available data points within K iterations, n_{inliers} is the number of data points that fit the current estimated model, N is the total number of data points, and s is the minimum number of samples required to fit a model. For the proposed algorithm, the minimum number of samples required is $s = 4$.

The coordinates of the ROI in knowledge instance i are then transformed based on the fitted 2-D homography T:

$$x'_{ROI_i} = \frac{t_0 x_{ROI_i} + t_1 y_{ROI_i} + t_2}{t_6 x_{ROI_i} + t_7 y_{ROI_i} + t_8} \tag{6}$$

$$y'_{ROI_i} = \frac{t_3 x_{ROI_i} + t_4 y_{ROI_i} + t_5}{t_6 x_{ROI_i} + t_7 y_{ROI_i} + t_8} \tag{7}$$

$$\text{where } T = \begin{bmatrix} t_0 & t_1 & t_2 \\ t_3 & t_4 & t_5 \\ t_6 & t_7 & t_8 \end{bmatrix} \tag{8}$$

A rectangular bounding box for ROI_i is constructed with the following coordinates:

$$(x_{\text{top-left}},\ y_{\text{top-left}}) = (\min(x_j),\ \min(y_j))$$

$$(x_{\text{bottom-right}},\ y_{\text{bottom-right}}) = (\max(x_j),\ \max(y_j))$$

The bounding box for a sample ROI_j is shown in Figure 4. The bounding box established is denoted by ROI^{rect} in the figure.

For ROI^{rect}, a four-vector boundary adjustment algorithm is performed for each ROI to form $ROI^{\text{opt\_scaled}}$. Four vectors from four different directions (up, down, left, right) are placed at the four boundaries of each region in ROI^{rect} and are moved inwards 1 pixel at a time for a maximum of 20% of the corresponding ROI dimension in all directions (ROI width for left and right vectors and ROI height for up and down vectors). After each move of a

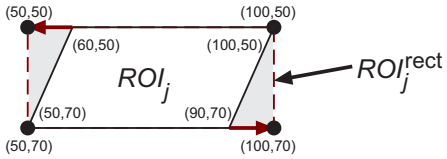

Figure 4. ROI Bounding Box Construction

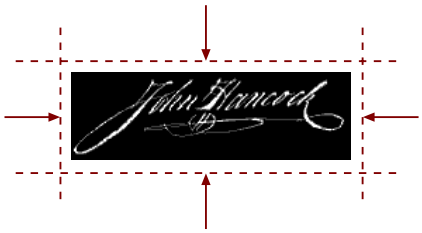

Figure 5. ROI Boundary Adjustment

vector, the boundary is readjusted according to a basic condition: if the number of "white" pixels in the downscaled inverted grayscale version of the image ("white" being any pixel with a grayscale intensity value greater than 10 in a dynamic range of 0 to 255) in the area between the current boundary and the vector is less than 1% of that area, then the current boundary is set to the position of the vector. If the condition is not met, then the current boundary is the final boundary of the ROI. This final boundary readjustment is performed to compensate for some of the irregularities that may occur in the model fitting and rectangular adjustment processes, such as the over-estimation of the ROI. An example of the boundary adjustment process is shown in Figure 5. The number of "white" pixels is used as a condition metric because the downscaled image is an inverted grayscale version of the original image, so these pixels correspond to the textual content of the document image. Finally, $ROI^{\text{opt\_scaled}}$ is scaled to match the dimensions of the original $IMAGE_R$ to form the final ROI set ROI^{opt}.

3. Experimental Results

For the purposes of evaluating the proposed algorithm, the automatic ROI selection system was implemented with a knowledge learning system consisting of 5 types of documents from the University of Waterloo Registrars Office categorized as shown in Figure 6. The system was trained using two sample images of each type.

The system also made use of the computer vision functions provided in [20]. A direct comparison with other established techniques was not possible since no known techniques are comparable in functionality. The main reason for this is that no content-specific digital image encryption

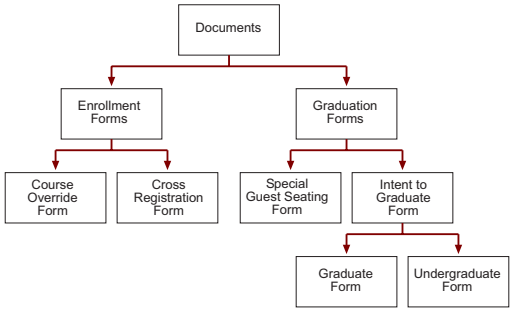

Figure 6. Test Knowledge Learning System

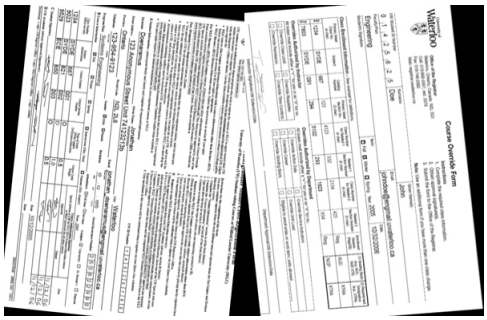

Figure 7. Sample Test Set Documents

algorithm of this kind has been proposed publicly prior to [1]. Therefore, the proposed algorithm is evaluated based on two important criteria:

1. Document matching accuracy at the conclusion of the initial ROI identification stage

2. Generated ROI coverage accuracy at the conclusion of ROI generation

These two criteria are important to evaluating the capabilities of the proposed system. First, the document matching accuracy of the initial ROI identification stage is evaluated. This is important for situations where the document type is unknown and a comprehensive search through the knowledge base is required. To evaluate the document matching accuracy at the initial ROI identification stage, the proposed algorithm is tested against a set of 50 document images to see if they are matched to documents of the same class. Sample document images from the test set are shown in Figure 7.

The test set consists of 10 document images for each document type in the knowledge system, with random rotations between $\pm 10^\circ$ and random translations in the x and y direction of $\pm 5\%$ of the corresponding dimension. These random translations and rotations are representative of com-

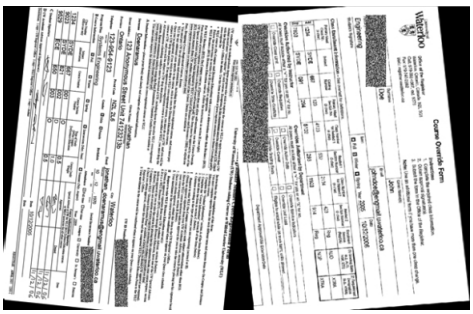

Figure 8. Encrypted Test Set Documents

Table 2. Initial Document Matching Accuracy

Document Type	Documents Matched
Course Override	100 %
Course Registration	100 %
Special Guest Seating	100 %
Intent to Graduate - Undergraduate	100 %
Intent to Graduate - Graduate	100 %
Average	100 %

mon disparities found in scanned document images. A summary of the results is shown in Table 2, and sample output images after encryption are shown in Figure 8. It is observed from the experimental results that a high level of accuracy is achieved at the initial ROI identification stage, with a classification accuracy of 100%. This level of accuracy is important as it reduces the time needed during the human interaction phase of ROI image encryption to correct any misclassification disparities.

While the high level of document matching accuracy at the initial ROI identification stage is beneficial to the system, of greater importance is the generated ROI coverage accuracy of the system at the final ROI generation stage. There are two main reasons for this assertion. First of all, the document type is generally known to a certain extent leading to a higher level of document matching accuracy than can be achieved when the entire system is searched. This in effect diminishes the importance of achieving perfect document matching of unknown documents at the initial ROI identification stage. More importantly, the level of document matching accuracy achieved at the initial stage is meaningless if the set of ROI constructed at the final ROI generation stage does not correspond to the information that needs to be encrypted. This stage requires a higher level of accuracy than the initial ROI identification stage. To evaluate the ROI coverage accuracy at the final ROI generation stage, the proposed algorithm is tested against a set of 10 document images with a set of 2 ROI selected by a human expert for each document image. The test set consists of 2 images of each document type in the system and is used as a gold standard to verify the performance of the proposed algorithm against. Two metrics are used to evaluate the coverage accuracy of the generated ROI. The first metric is a percentage that measures the ratio of computer generated ROI relative to the total area of regions that need to be covered as determined by a human expert. The second metric is a percentage that measures the ratio of ROI uncovered by the generated ROI to the total area of the document image less the total area of the regions that need to be covered as determined by a human expert.

A summary of all results is presented in Table 3. It is observed that a high level of coverage accuracy is achieved at the final ROI generation stage, with an average ROI coverage accuracy of 95.07% and a standard deviation of 4.18%. Furthermore, the average incorrect coverage percentage is only 0.62% with a standard deviation of 0.26%. The coverage inaccuracies that are observed at this stage are the result of ROI readjustment and scaling errors, as well as inaccuracies in the estimated homography used to generate the final set of ROI. The level of accuracy achieved in the final ROI generation stage is important as it minimizes the time needed during the human interaction phase of the ROI image encryption process to correct any content coverage disparities and therefore allows for significant timesaving benefits compared to having to select ROI manually. Better results can be achieved with further refinement of the parameters used in the image correlation and the RANSAC homography model fitting techniques.

4. Conclusions

In this paper, we have proposed a new method for automatic ROI selection using an expert knowledge learning system to reduce the need for human interaction in the ROI image encryption process. The proposed algorithm provides flexible ROI selection capabilities and delivers improved performance. Experimental results show that a high level of accuracy can be achieved for both document matching and final ROI generation. It is our belief that this method can be implemented for ROI image encryption schemes to reduce the amount of human interaction required. This reduction in human interaction leads to significant timesavings.

5. Acknowledgements

This research has been sponsored in part by Epson Canada and the Natural Sciences and Engineering Research Council of Canada.

Table 3. Generated ROI Coverage Accuracy

Document Type	ROI Area	
	Covered †	Uncovered ‡
Course Override	97.08 %	0.88 %
Course Registration	93.73 %	0.44 %
Special Guest Seating	98.46 %	0.93 %
Intent to Graduate - Undergraduate	97.73 %	0.49 %
Intent to Graduate - Graduate	88.33 %	0.37 %
Average Coverage Accuracy	95.07 %	0.62 %
Standard Deviation	4.18 %	0.26 %

The above results are based on the mean average of 5 test trials due to the slight disparities in estimated homographies as a result of the random nature of the RANSAC algorithm.

† Percentage relative to the total area of selected regions. The result is the mean percentage of 2 test images of the corresponding document type.

‡ Percentage relative to the total area of the document excluding the expert selected regions. The result is the mean percentage of 2 test images of the corresponding document type.

References

[1] D. Barber. Automatic Generation of Regions of Interest for Radionuclide Renograms. In *Proceedings of Medical Image Understanding and Analysis*, 2003.

[2] M. A. Fischler and R. C. Bolles. Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography. *Communications of the ACM*, 24:381–395, 1981.

[3] C. Harris and M. S. Plessey. A Combined Corner and Edge Detector. In *Proceedings of 4th Alvey Vision Conference*, pages 147–151, 1988.

[4] R. Heremans, A. Willekens, D. Borghys, B. Verbeeck, J. Valckenborgh, M. Acheroy, and C. Perneel. Automatic Detection of Flooded Areas on ENVISAT/ASAR Images using an Object-Oriented Classification Technique and an Active Contour Algorithm. In *Proceedings of the 31st International Symposium on Remote Sensing of the Environment*, June 2005.

[5] P. Kovesi. MATLAB and Octave Functions for Computer Vision and Image Processing. World Wide Web Document, 2006.
http://www.csse.uwa.edu.au/∼pk/research/matlabfns.

[6] H. Liu, H. Sahbi, L. C. Ferri, and M. Steinebach. Image Authentication using Automatic Detected ROIs. In *5th International Workshop on Image Analysis for Multimedia Interactive Services*, April 2004.

[7] J. Luo, K. J. Parker, and T. S. Huang. A Knowledge-Based Approach to Volumetric Medical Image Segmentation. In *Proceedings of the IEEE International Conference on Image Processing*, volume 3, pages 493 – 497, November 1994.

[8] V. Madasu, M. H. M. Yusof, M. Hanmandlu, and K. Kubik. Automatic Extraction of Signatures from Bank Cheques and Other Documents. In *Proceedings of the VIIth Digital Image Computing: Techniques and Applications*, December 2003.

[9] M. Mancas and B. Gosselin. Towards an Automatic Tumor Segmentation using Iterative Watersheds. In *Proceedings of the Medical Imaging Conference of the International Society for Optical Imaging (SPIE Medical Imaging 2004)*, pages 1598–1608, San Diego, California, May 2004.

[10] J. Noble. *Descriptions of Image Surfaces*. D.Phil. Thesis, Robotics Research Group, Department of Engineering Science, Oxford University, 1996.

[11] M. Park, L. S. Wilson, and J. S. Jin. Automatic Extraction of Lung Boundaries by a Knowledge-Based Bethod. *Visual Information Processing*, 2:14–19, 2001.

[12] F. Rottensteiner and J. Jansa. Automatic Extraction of Buildings from LIDAR Data and Aerial Images. In *Proceedings of the ISPRS Commission IV Symposium*, July 2002.

[13] S. Smith. A New Class of Corner Finder. In *Proceedings of the 3rd British Machine Vision Conference*, pages 139–148, 1992.

[14] G. Torheim, F. Godtliebsen, D. Axelson, K. A. Kvistad, O. Haraldseth, and P. A. Rinck. Feature Extraction and Classification of Dynamic Contrastenhanced T2*-Weighted Breast Image Data. *IEEE Transactions on Medical Imaging*, 20(12):1293–1301, December 2001.

[15] T. Tweed and S. Miguet. Automatic Detection of Regions of Interest in Mammographies Based on a Combined Analysis of Texture and Histogram. In *Proceedings of 16th International Conference on Pattern Recognition*, volume 2, pages 448–452, May 2002.

[16] L. F. Verheij, J. A. K. Blokland, A. M. Vossepoel, R. Valkema, J. A. J. Camps, S. E. Papapoulos, O. L. M. Bijvoet, and E. K. J. Pauwels. Automatic Region of Interest Determination in Dual Photon Absorptiometry of the Lumbar Spine. *IEEE Transactions on Medical Imaging*, 10(2):200–206, June 1991.

[17] R. Wang and Z. Chi. Automatic Segmentation of Chinese Chunks using a Neural Network. In *Proceedings of the 2003 International Conference on Neural Networks and Signal Processing*, volume 1, pages 96–99, December 2003.

[18] A. Wong. Multi-Level Regions-of-Interest (ROI) Image Encryption for Digital Document Archival and Distribution Infrastructures. To Be Submitted, March 2006.