

Automatic identification of ROI in figure images toward improving hybrid (text and image) biomedical document retrieval

Daekeun You^a, Sameer Antani, Dina Demner-Fushman, Md Mahmudur Rahman,
Venu Govindaraju^a, George R. Thoma

National Library of Medicine, National Institutes of Health, Bethesda, MD 20894

^aCUBS, Dept. of Computer Science and Engineering, SUNY at Buffalo, Buffalo, NY 14260

ABSTRACT

Biomedical images are often referenced for clinical decision support (CDS), educational purposes, and research. They appear in specialized databases or in biomedical publications and are not meaningfully retrievable using primarily text-based retrieval systems. The task of automatically finding the images in an article that are most useful for the purpose of determining relevance to a clinical situation is quite challenging. An approach is to automatically annotate images extracted from scientific publications with respect to their usefulness for CDS. As an important step toward achieving the goal, we proposed figure image analysis for localizing pointers (arrows, symbols) to extract regions of interest (ROI) that can then be used to obtain meaningful local image content. Content-based image retrieval (CBIR) techniques can then associate local image ROIs with identified biomedical concepts in figure captions for improved hybrid (text and image) retrieval of biomedical articles.

In this work we present methods that make robust our previous Markov random field (MRF)-based approach for pointer recognition and ROI extraction. These include use of Active Shape Models (ASM) to overcome problems in recognizing distorted pointer shapes and a region segmentation method for ROI extraction.

We measure the performance of our methods on two criteria: (i) effectiveness in recognizing pointers in images, and (ii) improved document retrieval through use of extracted ROIs. Evaluation on three test sets shows 87% accuracy in the first criterion. Further, the quality of document retrieval using local visual features and text is shown to be better than using visual features alone.

Keywords: Biomedical image analysis, biomedical article retrieval, content-based image retrieval, image overlay extraction, pointer symbol extraction, Active Shape Model

1. INTRODUCTION

Biomedical images are frequently used in publications to illustrate the medical concepts or to highlight special cases. They are invaluable in establishing diagnosis, acquiring technical skills, and implementing best practices in many areas of medicine. Conventional approaches for biomedical journal article retrieval have been text-based with little attention devoted to the use of images in the articles. Text-based retrieval uses text information automatically extracted from title, abstract, figure captions, and discussions (mention). It provides fairly good results; however, the relevance quality sometimes is not satisfactory. Content-based image retrieval (CBIR) also has been applied to biomedical image retrieval [1]. However, the retrieval performance is far behind the text-based retrieval due to semantic gap [2]. Low level features such as color, textual, and shape used in CBIR are insufficient to represent medical concepts or meaningful diagnostic information in the images effectively.

To improve the relevance quality of conventional retrieval approaches, we have proposed an approach using hybrid (text and image) features [3]. Information retrieval (IR) techniques are used to identify key textual features in the title, abstract, figure caption, and figure citation (mention) in the article. Structured vocabularies, such as the National Library of Medicine's Unified Medical Language System (UMLS[®]) are used as well to identify the biomedical concepts in the

text [4-7]. Unlike conventional CBIR that uses image features from the entire image, our proposed approach uses a combination of features computed over the entire image and those computed from specific image regions of interest (ROIs). We recognize that authors often use the annotations overlaid on figures and illustrations in the articles in the form of pointers or symbols to highlight regions of interest. Recognizing these image annotations may be useful to identify specified local regions. These annotations are also often referenced in the figure meta-text (captions, mentions) in the article. We hypothesize that correlating biomedical concepts from the figure meta-text with image features computed on the image regions identified by the pointers may improve the quality of biomedical document retrieval when using both images and text features in the query.

This article presents our efforts to improve our prior work [3-4] on pointer recognition and ROI extraction to achieve better relevance quality in the proposed multimodal biomedical article retrieval. Additional pointer segmentation and ROI extraction methods were developed based on region growing method. A new pointer recognizer based on Active Shape Model (ASM) was developed to complement our Markov random field (MRF)-based pointer recognizer proposed in [3]. Biomedical image retrieval tests utilizing several components we have developed have been performed and some initial test results and performance analysis are discussed as well.

2. PRIOR WORK

We actively participated in the ImageCLEF 2009¹ medical retrieval track and submitted various retrieval runs based on textual, image, and multimodal (combinations of text- and content-based approaches) features [8]. Our group and several runs were ranked on top among 17 research groups and over 100 submitted runs for image- and case-based topics [9]. Our multimodal relevance feedback and visual-based retrieval approaches were ranked 1st in interactive and visual runs, respectively. Our case-based runs also were ranked in 1st and 2nd among all case-based run submissions. Text features were extracted from image captions provided in the ImageCLEFmed'09 collection and several automatically extracted search areas such as title of the article in which the image appears, the article's abstract, a brief mention of the image from the article's full text, and the Medical Subject Headings (MeSH terms), which is a controlled vocabulary created by NLM to index biomedical articles, assigned to the article. Image feature vectors were obtained from both visual concept-based feature based on a "bag of concepts" model comprising color and texture patches from local image regions [10] and various low-level global features including color, edge, and texture.

Our previous approach to locating and recognizing pointers in biomedical images proposed a Markov random field (MRF)-based recognition scheme to add robustness to our first approach discussed in [4]. 43 labels were defined from boundary parts frequently seen in commonly used arrow type pointers overlaid on biomedical images. MRF theory is applied to label line segments extracted from a pointer boundary. Dynamic programming (DP) technique was applied to the line segment labeling results to select the best label for each line segment and find an optimal configuration for the boundary. Hidden Markov model (HMM)-based classifier following the MRF labeling was applied to classify a labeling configuration into three pointer classes. Our test results showed that the proposed method can recognize almost all arrow type pointers used in biomedical images and is less affected by the large variation in pointer shape. 82% success rate in pointer recognition was reported on a pointer image set manually cropped from ImageCLEFmed'09 image set.

In addition to the pointer recognition test, we performed biomedical image retrieval test utilizing the pointer recognition method and local image analysis. Image regions of interest (ROI) indicated by the pointers were localized and image features were computed from the ROIs. Two image-based topics from ImageCLEFmed'08 were selected. Initial results retrieved by text retrieval were re-ranked based on the comparison results between the ROIs from the images and sample ROIs of each topic. We achieved mixed retrieval results, some of which were promising. Several potential solutions to improve the quality of retrieval were discussed in [3].

3. METHODS

Three methods are proposed as solutions to the drawbacks identified in [3]. These are improved pointer segmentation, improved pointer recognition, and variable-size ROI extraction. These methods are described in this section.

¹ <http://imageclef.org>

3.1 Pointer segmentation

Edge-based segmentation is simple but powerful and has been widely used for object segmentation [11]. Object boundary can be segmented by edge detection followed by binarization of the edge image. A drawback of edge-based segmentation caused by weak edges was described in [4]. Recently we identified another drawback causing errors in small pointers recognition. Figure 1 shows an example of the case. Small pointers frequently lose their tail due to the smoothing effect of edge detection and binarization, which shrink the region of segmented pointers. Figure 1(a) shows extracted boundaries by Roberts edge detection followed by adaptive thresholding. The boundaries are overlaid on the pointers. They were recognized as arrowhead or noise.

Among many edge detection operators, Roberts edge detection was applied in our prior work. Compared to other famous edge detection methods such as Sobel and Canny, Roberts operator generates less noise boundaries than Sobel but more than Canny. Figure 2 shows edge detection results of the three methods on the entire image containing Figure 1(a). Canny operator generates least noise boundaries; however, it does not guarantee a closed boundary of a pointer. In order to form a closed boundary, a simple morphological operation can be applied to the edge image. One serious problem of morphological operation is that the shrinking discussed above is worse than the result shown Figure 1(a). Figure 1(c) shows extracted pointer boundaries by Canny edge detection followed by dilation. The tails of all arrows disappeared and the arrowheads are smaller than the result in Figure 1(a).

A pointer segmentation method based on Canny edge detection and region growing was developed to take the benefit and solve the drawback of Canny edge detection in pointer recognition. The segmentation method has two steps, viz., histogram analysis and region growing. In the histogram analysis step, interior region of detected pointer boundary is used as seed region and all gray intensity levels of the pixels are collected. The *min* and *max* of all the intensity levels are obtained and used in region growing step. In region growing step, only pixels outside of the boundary are examined and merged based on the *min* and *max*. The region growing can be used with any edge detection method and can segment almost entire region of target pointers (compare (a) and (c) with (b) and (d), respectively). Figure 1(b) and (d) show pointer boundaries extracted by region growing with Roberts and Canny edge detection methods, respectively.

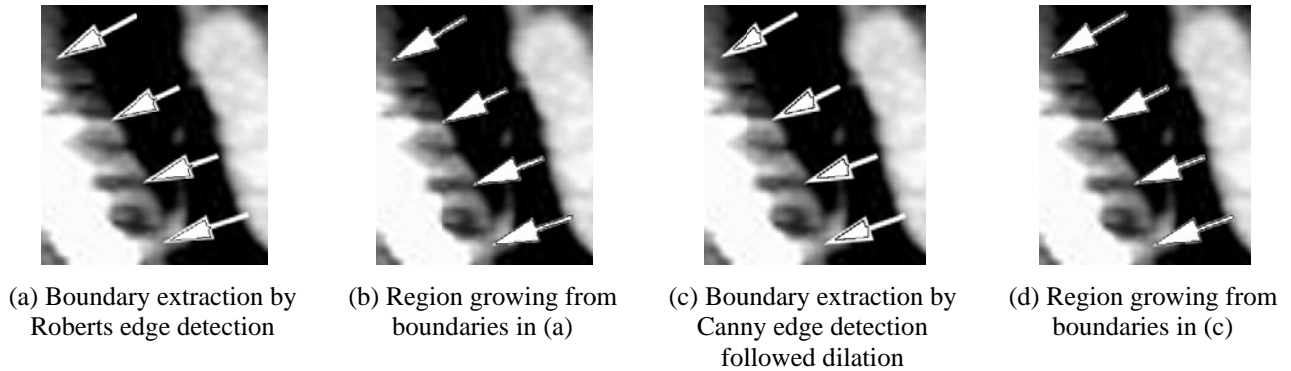


Figure 1. Edge detection and pointer segmentation results




Figure 2. Comparison of three edge detection methods

3.2 Active Shape Model (ASM) pointer recognizer

Our Markov random field (MRF)-based pointer recognition is capable of recognizing various shapes of arrow type pointers with no specific pointer models used. General model based approaches define object models, train the models, and recognize objects by comparing unknown object with the models. Difficulties in applying model-based approaches to pointer recognition were discussed in our prior work [3]. One drawback of the MRF-based recognizer is that it can not recognize pointers with some degree of distortion in their boundaries. The MRF recognizer can handle minor distortion and over-segmented line segments; however, severe distortion as shown in Figure 3 can not be handled even though the pointer is perfectly segmented from the background. The boundary and its line segment approximation result have some distortion and the distortion caused incorrect line segment labeling and recognition result as well. Other approaches rather than boundary-based method may recognize the segmented pointer since the overall shape of the pointer is fine enough and the pointer region is well-separated from background. Figure 3(d) shows result of ASM recognizer, which is discussed below.

A recognizer based on Active Shape Model (ASM) was developed to handle pointers with distortion in their boundary and line segment approximation results. The ASM was developed by Cootes and Taylor in 1995 and has been successfully applied to object detection and recognition problems such as face detection, hand detection, and medical applications [12]. Unlike other applications where rotation of object is not explicitly considered, pointer recognition should consider pointer rotation since pointers point to arbitrary direction and they need to be aligned with model pointers to get correct similarity score. To address this problem, a segmented pointer is rotated and mirrored several times and then each transformed pointer is matched with model pointers. The eigen analysis used in [4] is applied to obtain initial rotated pointer. Four different pointers including the initial pointer are generated and then compared to models. Pixels with same gray intensity level (black-black or white-white) from two overlapped pointer regions, one is unknown pointer and the other is interior region of fitted ASM model, are counted and then divided by $width \times height$ of the segmented pointer image. The ratio is used as similarity score to choose the best matching model.

Figure 4 shows types of recognizable pointers by ASM recognizer. The numbers shown in parentheses are the number of landmark points in each pointer model. Only several pointers frequently used in biomedical images are selected. Compared to MRF recognizer, only several types of pointers used in model generation can be recognized and that is the biggest drawback of the ASM recognizer. Figure 5 shows an example pointer and its transformed images in matching step. The curved arrow is initially rotated by a rotation angle calculated by eigen analysis (Figure 5(b)) and three more pointers (Figure 5(c)~(e)) are generated by mirroring the pointer in previous step. The pointer in Figure 5(d) is the best matching mirror image and the pointer model (Figure 4(b)) generates the largest similarity score and hence is selected as recognition result.

Pointers are recognized by MRF recognizer first and then ASM pointer recognizer only if the similarity score from MRF recognizer is smaller than a threshold. Boundaries are assumed as noise if the similarity score from ASM recognizer is less than a threshold.




Figure 3. A pointer that MRF recognizer cannot recognize due to distortion in its boundary

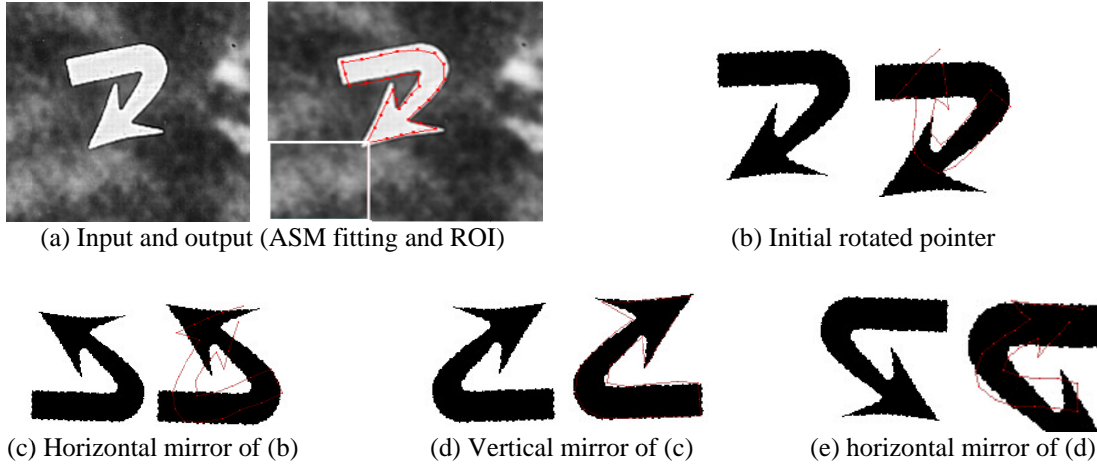
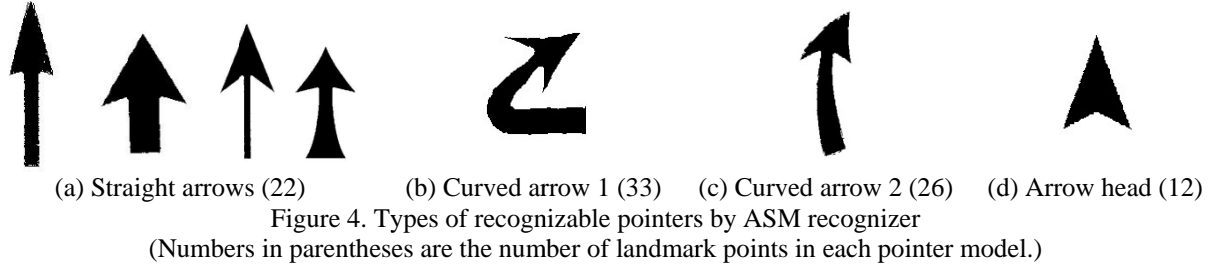


Figure 5. Example of ASM recognition
(Left: input to ASM, right: final fitting result with overlaid fitted model)




Figure 6. ASM recognition results and ROIs

Figure 6 shows some recognition results of ASM recognizer. Solid lines and points overlapped on each pointer show the best fitted model pointer and its landmark points. The rectangle is the region of interest (ROI) pointed by each pointer.

3.3 ROI extraction

In our prior work [3-4], a fixed size rectangle ROI (e.g., 200×200) is localized based on some information from pointer recognition result. The fixed size ROI is easy to localize; however, it is sometimes not proper to contain an entire concept in the local region. A variable-size ROI extraction method has been developed similar to the pointer segmentation method described in section 3.1.

The ROI extraction method has two steps: (i) seed region histogram analysis and (ii) region growing. In the seed region histogram analysis step, a 100×100 rectangle ROI as before is used as seed region (see Figure 7). Gray intensity histogram is obtained from the seed region and some representative gray intensity levels are extracted by detecting dominant peaks in the histogram. In the region growing step, pixels with any of the representative gray intensity levels are selected as seed pixel for region growing. From a seed pixel in the seed region, pixels with gray intensity of $seed\ intensity \pm 10$ are merged. Two measurements are calculated from a grown region to decide termination of the growing process. They are the centroid of a region bounding box and distances between two reference lines and the centroid. Growing process is terminated when the distances are larger than some thresholds or the region bounding box becomes larger than a threshold. Figure 7 describes the seed region, a region being grown, and the measurements for the decision of growing process termination. Figure 8 shows all local regions grown from two representative intensity levels (67 and 107). The final ROI is a rectangle containing all the rectangles of grown regions. As seen in the result, ROI extracted by region growing is more precise than the ROI obtained by 200×200 fixed-size rectangle shown in (a).

One benefit of the ROI extraction method is that it enables extraction of large ROI pointed by several identical pointers. ROI in Figure 9(a) is indicated by six black arrows and the arrows are located near the boundary of the interesting content. None of the 200×200 rectangle ROIs contains the entire ROI and hence these ROIs should be combined in some ways to extract the entire ROI. The right result of (a) shows a ROI extracted by region growing starting from one arrow. Other ROIs by the other five arrows are similar with the result.




Figure 7. ROI extraction by region growing




Figure 8. Region growing ROI extraction result




Figure 9. Variable-size ROI extraction results

(In each pair left one is 200×200 fixed ROI and right one is variable-size ROI.)

3.4 Multimodal biomedical image retrieval

In our proposed multimodal biomedical image retrieval, we are interested in local image features rather than features computed from entire images. Features from image ROIs can be used in various multimodal retrieval scenarios. Text- or image-based retrieval results can be re-ranked by comparing ROIs extracted from initial retrieved images with sample ROIs. The sample ROIs can be identified by analyzing text query and then retrieved from an image DB or can be a user-marked ROI if the query includes an image. Then images with pointers and containing similar local regions with the query, which are expected to be more relevant, could be ranked on top of the new result.

For hybrid queries (text and image), utilizing features computed from ROIs could provide better results than using image features computed from entire images. Text retrieval is efficient to reduce the search space; however, the results may not be satisfactory since it does not use any image features and hence may not find images containing similar contents that the user-provided query image has. CBIR by features computed from entire images could retrieve images that look like the query image. Those features, however, may be inadequate to represent specific concepts in local regions or concepts that users provide by text queries. ROI analysis can be placed between text retrieval and CBIR and may provide useful information complementing both text retrieval and CBIR. Specific medical concepts associated with the ROIs when the images are annotated also could improve relevance quality. The extracted text concepts from the ROIs in the retrieved images can be compared with the concepts extracted from the text query. The initial results could be simply re-ranked by the local image features or text concepts. Additional relevant images could be retrieved by another text retrieval using the concepts identified through the ROI analysis.

Multimodal features may be the most potential method to obtain better results over conventional retrievals. Using entire images, however, may not be helpful in multimodal retrieval. Results by text- or image-only are expected to be fairly relevant to the queries. Features computed from entire images may no longer be able to provide detailed and meaningful information for better relevance quality over the initial results due to semantic gap. The pointer and overlay finding results is expected to improve the specificity in image retrieval. If a pointer has been detected in an image through the use of text clues indicating their presence or directly by the MRF pointer recognizer, the image region pointed to can provide greater specificity on the image content. We expect this to significantly improve retrieval quality.

4. EXPERIMENTS

4.1 Test setup

We performed biomedical image retrieval test as discussed in [3]. ImageCLEFmed'09 data set was used and three topics out of 25 topics were selected based on the number of ground truth relevant images to the topics. The three topics are i) *Topic 2: Breast cancer mammogram*, ii) *Topic 12: Radiographic findings of osteomyelitis*, and iii) *Topic 21: Osteoporotic bone*. Figure 10 shows sample images relevant to each topic. Two retrieval runs, text and image modes, submitted to the competition were selected for the initial text- and image-based retrieval results [9]. The initial results were re-ranked by comparing ROIs extracted from images of initial result with sample ROIs for each topic. TRECEVAL package [13] was used to analyze and obtain common evaluation measures such as precision and recall.




Figure 10. Sample relevant images to each topic

4.2 Evaluation results

Through this retrieval test, we evaluated two components in our proposed multimodal retrieval approach; i) pointer recognition and ii) retrieval performance using local image features. Table 1 and 2 shows some useful statistics on rate of pointer presence in retrieved images and ground truth images of the three topics. *Total ret* denotes the number of total retrieved images by each retrieval mode for each topic and the next three columns (w/ pointer, w/o pointer, and rel-ret) show percentages out of the *Total ret* in each row. *rel-ret w/ pointer* means images that contain pointers and are relevant to the topic among the retrieved images. The numbers show percentages out of the *rel-ret* (relevant and retrieved). For example, 294 (29%) out of the 1,000 images retrieved for topic 2 were judged relevant and 171 (58%) images among the 294 *rel-ret* images were containing pointer(s). From Table 1 we can notice that approximately 60% of retrieved images by any retrieval mode have pointers and the other 40% images have no pointers. Similar percentages can be found from ground truth of each topic shown in Table 2.

Table 1. Pointer presence rate in retrieved images

	Text-based retrieval					Image-based retrieval				
	Total ret	w/ pointer	w/o pointer	rel-ret	rel-ret w/ pointer	Total ret	w/ pointer	w/o pointer	rel-ret	rel-ret w/ pointer
Topic 2	1,000	58%	42%	29%	58%	1,000	67%	33%	4%	64%
Topic 12	847	64%	36%	23%	64%	843	59%	41%	4%	46%
Topic 21	1,000	61%	39%	6%	39%	996	60%	40%	1%	36%

Table 2. Pointer presence rate in ground truth images

	# of Qrel relevant	w/ pointers	w/o pointers
Topic 2	444	53%	47%
Topic 12	236	60%	40%
Topic 21	133	44%	56%

Table 3. Pointer recognition performance

	Text-based retrieval			Image-based retrieval		
	Success	False alarm	Missed detection	Success	False alarm	Missed detection
Topic 2	82%	10%	8%	89%	2%	9%
Topic 12	83%	4%	13%	87%	3%	10%
Topic 21	82%	7%	11%	85%	5%	10%

Table 3 shows pointer recognition evaluation results. Three measures, *success*, *false alarm*, and *missed detection*, are considered to evaluate the performance. The *success* includes images that i) have pointer(s) and algorithm detected all or some of them (w/ or w/o some noise pointers) and ii) have no pointers and algorithms detected no pointers. The *false alarm* includes images that have no pointers but algorithm detected some noise pointers. The *missed detection* counts images that have pointer(s) but algorithm detected nothing or some noise pointers.

Our pointer recognizer is a combination of the MRF and ASM recognizers. The ASM recognizer alone can not achieve better recognition performance than the MRF recognizer due to the drawback discussed in Section 3.2. Adding and training more pointer models may solve the drawback. However, the processing time increases according to the number of pointer models and the increased running time may not be acceptable.

Average success rates are 82% and 87% for images retrieved by text and image modes, respectively. Most images retrieved by image mode have less complicated background than those retrieved by text mode. They have similar texture

and color with the sample images shown in Figure 10 since they are retrieved by CBIR engine. Images similar with the sample images have fewer noise boundaries and hence produced lower false alarm rate.

Figure 11 shows Recall-Precision graph and precision graph after N images retrieved. Only relevant images with pointers on the ground truth list were considered as relevant images since we re-rank the initial retrieved result based on presence of pointers and ROI analysis. The graphs show that text mode retrieval shows the best performance in all cases. *Visual+ROI* and *Visual* show very poor performance since the initial retrieval results (by image mode) contained fewer relevant images than text mode results.

Our retrieval tests show some promising results on use of pointer localization and local image analysis for improved biomedical image retrieval. Our ultimate goal is to obtain better performance graph than the *Text* results in the graphs. In that case more relevant images with pointers can be ranked on top and hence the relevance quality could be improved.

To achieve improvement on our initial test, we need to obtain list of relevant images to a certain query topic judged by local regions pointed by pointers. The relevant images in this test were selected by observing the entire image, not the ROIs, and hence some ROIs may not be relevant to the query topic even though they are extracted from relevant images. Another improvement can be made by enhancing pointer recognition performance. Achieving higher success rate is necessary to extract more precise ROI.

5. CONCLUSION

Local image region in biomedical images may have more meaningful information and may be more relevant than other region in an image for biomedical image retrieval. Authors frequently use pointers and symbols to highlight specific local regions and mention them in figure captions and text discussions. Localizing those pointers can help extract specific local regions of interest (ROIs) and using the ROIs could improve relevance quality of conventional retrieval approaches by combining textual and image features from local regions.

In this article we present our research effort to enhance our prior work on pointer recognition and ROI extraction. Region growing technique was applied to improve pointer segmentation and ROI extraction performance. Active Shape Model (ASM)-based pointer recognizer was developed to handle pointers that can not be recognized by the MRF recognizer due to some distortion in their boundary. Average 87% success rate on pointer recognition was achieved.

This article also presents preliminary retrieval test results. In order to verify effectiveness of our retrieval approach, it is necessary to consider several important issues and appropriate evaluation methods, viz., (i) accurate ROI identification and extraction, (ii) feature selection for image ROI comparison, (iii) database generation of ground truth ROIs (image patches) for query topics identified from text, and (iv) scheme of utilizing local image comparison results to obtain better retrieval relevance quality.

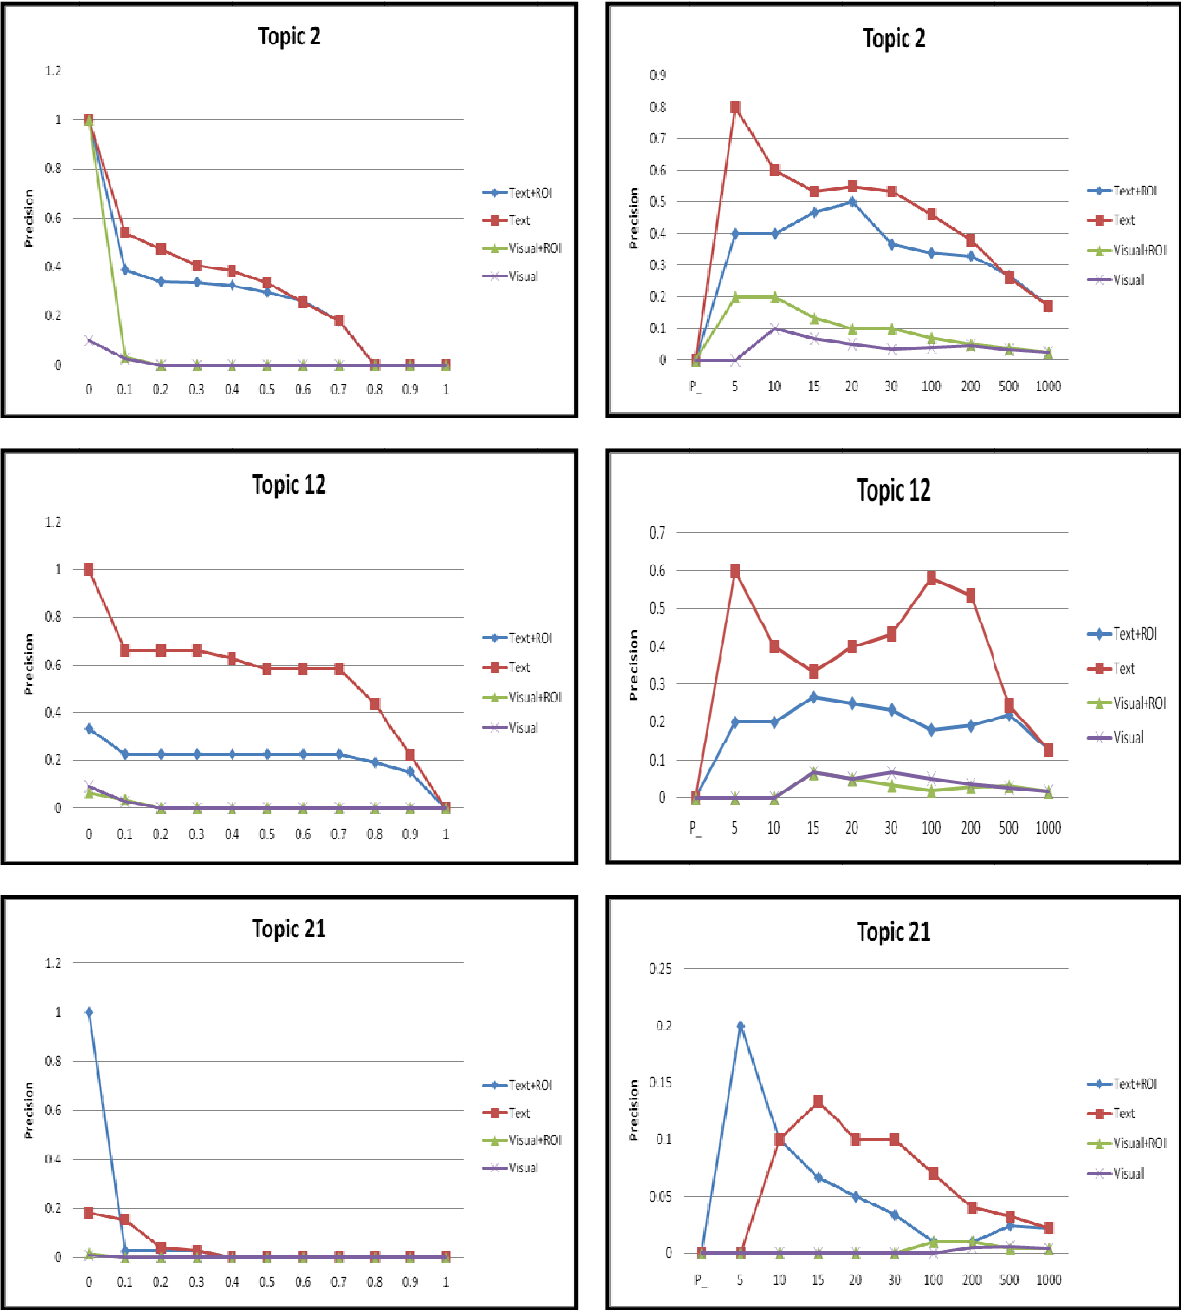
One of interesting future work is identifying ROIs in images that are mentioned in text such as figure captions and discussions. Several image features such as color and shape of pointers are available by pointer recognition. More specific features are available from text such as size (small, large), location (left, right, etc.), and plurals (single or multiple pointers pointing to an ROI). Numerical features of pointers obtained from image (e.g., gray intensity level, length of contour boundary, location coordinates, etc.) can be more specific by mapping related text features to them. ROIs also could be identified and extracted more accurately by utilizing the text features and image features as well. Fixed-size ROI may provide us easy way of extracting and utilizing local image regions in our retrieval approach. However, our ultimate goal is to extract and use variable-size ROI since it fits better to desired ROIs.

Acknowledgement

This research was supported by the Intramural Research Program of the Lister Hill National Center for Biomedical Communications, an R&D division of the National Library of Medicine, at the National Institutes of Health, U.S. Department of Health and Human Services. We would like to thank the imageCLEF [14, 15] organizers and the Radiological Society of North America (RSNA), publisher of Radiology and RadioGraphics², for making the database

² <http://www.rsna.org/publications>

available for the experiments under the ImageCLEFmed medical image retrieval task. Also we would like to thank Matthew Simpson (NLM) for processing text analysis results.



(a) Recall-Precision graph (b) Precision after N images retrieved
Figure 11. Retrieval result graphs

REFERENCES

- [1] Müller H, Michoux N, Bandon D, Geissbuhler A, "A review of content-based image retrieval systems in medical applications - clinical benefits and future directions," *Int. J. of Medical Informatics*. 73(1):1-23 (2003)
- [2] Deserno TM, Antani S, Long R, "Ontology of gaps in content-based image retrieval," *Journal of Digital Imaging*. 22(2):202-15, April (2009).
- [3] You D, Antani SK, Demner-Fushman D, Rahman MM, Govindaraju V, Thoma GR, "Biomedical article retrieval using multimodal features and image annotations in region-based CBIR," *Proc. SPIE-IS&T Electronic Imaging*. San Jose, CA. January 2010;7534:75340V(1-12).
- [4] You D, Apostolova E, Antani SK, Demner-Fushman D, Thoma GR, "Figure content analysis for improved biomedical article retrieval," *Proc. SPIE-IS&T Electronic Imaging*. San Jose, CA. January 2009;7247:72470V(1-10).
- [5] Demner-Fushman D, Antani SK, Thoma GR, "Automatically Finding Images for Clinical Decision Support," *Proceedings of Workshop on Data Mining in Medicine*, 7th IEEE Intl Conf on Data Mining 2007:139-44.
- [6] Demner-Fushman D, Antani SK, Simpson M, Thoma GR, "Combining Medical Domain Ontological Knowledge and Low-level Image Features for Multimedia Indexing," *Proc. 2nd International "Language Resources for Content-Based Image Retrieval" Workshop (OntoImage 2008)*, part of 6th Language Resources and Evaluation Conference (LREC 2008). 2008;CDROM Proceedings.
- [7] Antani SK, Demner-Fushman D, Li J, Srinivasan BV, Thoma GR, "Exploring use of images in clinical articles for decision support in Evidence-Based Medicine," *Proc. SPIE-IS&T Electronic Imaging*. San Jose, CA. January 2008;6815:68150Q(1-10).
- [8] Müller H, Kalpathy-Cramer J, Eggel I, Bedrick S, Radhouani S, Bakke B, Kahn Jr. CE, Hersh W, "Overview of the CLEF 2009 medical image retrieval track," *Working Notes of CLEF 2009 (Cross Language Evaluation Forum)*.
- [9] Simpson M, Rahman MM, Demner-Fushman D, Antani S, Thoma GR, "Text- and content-based approaches to image retrieval for the ImageCLEF 2009 medical retrieval track," *Working Notes for the CLEF 2009 Workshop*, Corfu, Greece. September (2009).
- [10] Rahman MM, Antani SK, Thoma GR, "A medical image retrieval framework in correlation enhanced visual concept feature space," 22nd IEEE International Symposium on Computer-Based Medical Symposium (CBMS). Albuquerque, NM. August (2009).
- [11] Sonka M, Hlavac V, Boyle R, [Image Processing, Analysis, and Machine Vision], Thomson-Engineering (2007).
- [12] Cootes TF, Taylor CJ, Cooper DH, Graham J, "Active shape models - their training and application," *Computer Vision and Image Understanding*, 61(1):38-59 (1995).
- [13] National Institutes of Standards and Technology (NIST) – Text Retrieval Evaluation Conference (TREC) <http://trec.nist.gov/>
- [14] Müller, H., Kalpathy-Cramer, J., Kahn, Jr. C. E., Hatt, W., Bedrick, S., and Hersh, W. 2008. Overview of the ImageCLEFmed 2008 Medical Image Retrieval Task. In *Proceedings of the LNCS, 5706*. 9th Workshop of the Cross-Language Evaluation Forum, CLEF 2008, (Aarhus, Denmark, Sep. 17-19, 2008). Revised Selected Papers 2009.
- [15] Müller, H., Kalpathy-Cramer, Eggel, I., Bedrick, S., Radhouani, S., Bakke, B., Kahn, Jr. C. E., and Hersh, W. 2009. Overview of the CLEF 2009 Medical Image Retrieval Track. *Working Notes for the CLEF 2009 Workshop*, (Corfu, Greece, Sep. 30-Oct. 2, 2009).