

Chapter 14

Text- and Content-Based Medical Image Retrieval in the VISCERAL Retrieval Benchmark

Fan Zhang, Yang Song, Weidong Cai, Adrien Depeursinge and Henning Müller

Abstract Text- and content-based retrieval are the most widely used approaches for medical image retrieval. They capture the similarity between the images from different perspectives: text-based methods rely on manual textual annotations or captions associated with images; content-based approaches are based on the visual content of the images themselves such as colours and textures. Text-based retrieval can better meet the high-level expectations of humans but is limited by the time-consuming annotations. Content-based retrieval can automatically extract the visual features for high-throughput processing; however, its performance is less favourable than the text-based approaches due to the gap between low-level visual features and high-level human expectations. In this chapter, we present the participation from our joint research team of USYD/HES-SO in the VISCERAL retrieval task. Five different methods are introduced, of which two are based on the anatomy–pathology terms, two are based on the visual image content and the last one is based on the fusion of the aforementioned methods. The comparison results, given the different methods indicated that the text-based methods outperformed the content-based retrieval and the fusion of text and visual contents, generated the best performance overall.

F. Zhang (✉) · Y. Song · W. Cai
Biomedical and Multimedia Information Technology (BMIT) Research Group,
School of Information Technologies, University of Sydney, Sydney, NSW, Australia
e-mail: fzha8048@uni.sydney.edu.au

W. Cai
e-mail: tom.cai@sydney.edu.au

A. Depeursinge · H. Müller
University of Applied Sciences Western Switzerland (HES-SO), Sierre, Switzerland
e-mail: adrien.depeursinge@hevs.ch

H. Müller
e-mail: henning.mueller@hevs.ch

14.1 Introduction

Medical image data produced has been growing rapidly in quantity, content and dimension, due to an enormous increase in the number of diverse clinical examinations performed in digital form and to the large range of image modalities and protocols available [1–5]. Retrieving a set of images that are clinically relevant to the query from a large image database has been the focus of medical research and clinical practice [6–9]. The relevance between the images is normally computed in two manners, i.e. text and content based. The text-based approach is performed given the manual clinical / pathological descriptions, which require that the experts manually index the images with alphanumeric keywords if no text is already available with the images. The content-based retrieval is based on the image visual content information, which automatically extracts the rich visual properties / features to characterize the images [10–12]. While the text-based retrieval is the more common method, the content-based approach is attracting more interest due to the fact that medical image data have expanded rapidly in the past two decades [13, 15–17]. The combination of the two approaches suggests a potential direction of medical image retrieval for performance improvement [18].

In the VISCERAL Retrieval Benchmark,¹ we conducted medical image retrieval based on multimodal and multidimensional data [20]. The similarities between medical cases are computed based on the extracts of the medical records, radiology images and radiology reports. The VISCERAL Retrieval dataset consists of 2311 volumes originated from three different modalities of CT, MRT1 and MRT2. The volumes are from different human body regions such as the abdomen, thorax and the whole body. Within the whole dataset, 1815 volumes are provided with anatomy–pathology terms extracted from the radiology reports. A total of 10 topics with diagnosis and case description were used as queries (see [20] for details). Each of them was annotated with the 3D bounding box of the region of interest (ROI), binary mask of the main organ affected and the corresponding anatomy–pathology terms. A brief introduction of our participation has been reported in [19] and more on the VISCERAL data in general, and the evaluation approach can be found in [20]. Our experimental results are reported with text-based retrieval that utilized the anatomy–pathology terms, with visual content-based retrieval that made use of the visual content features, and with information fusion that combined the above results.

The structure is as follows: in Sect. 14.2, we introduce the text, visual content and fusion retrieval methods that were used in our participation; in Sect. 14.3, we present the experimental results and discussion; and we provide the conclusion in Sect. 14.4.

¹<http://www.visceral.eu/benchmarks/retrieval2-benchmark/>.

14.2 Methods

A general framework of image retrieval consists of the following steps [13, 14]: feature extraction, similarity calculation and relevance feedback, as illustrated in Fig. 14.1. For our methods, the feature extraction is conducted by analysing the anatomy–pathology term (Sects. 14.2.1 and 14.2.2) and the image content information (Sect. 14.2.3). The similarity is computed by measuring the Euclidean distance between the feature vectors. The relevance feedback is extracted based on the neighbourhood information among the cases for retrieval result refinement (Sect. 14.2.4).

14.2.1 Term Weighting Retrieval

Medical image retrieval is conventionally performed with text-based approaches, which rely on manual annotation with alphanumeric keywords. The anatomy–pathology term files provided in the VISCERAL Retrieval Benchmark list the pathology terms and affected anatomies that were extracted from the German radiology reports and mapped to RadLex. The co-occurrence of different anatomy–pathology terms on the same cases can be used to evaluate the terms’ effectiveness of finding the similarity between subjects, for example, of some “stop words” that occur widely but have little influence on describing the similarities. Our text-based methods are based on the co-occurrence matrix between the terms and cases.

For our first text-based method, we used term frequency inverse document (case) frequency (TF-IDF) [21] to weight the terms for each case. TF-IDF can find the rare terms that carry more information than the frequent ones and is thus widely applied in term weighting problems. Formally, a case-term co-occurrence matrix $OCC_{NT \times NC}$ is constructed according to the anatomy–pathology terms on different cases, where the element $occ(t, c)$ refers to the number of occurrences of term T_t on case C_c ,

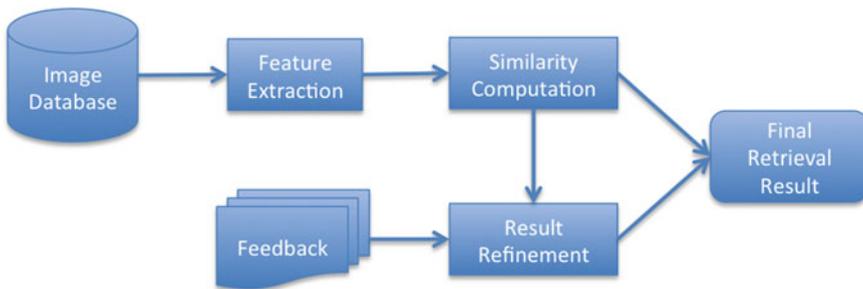


Fig. 14.1 Image retrieval pipeline in this study: (1) feature extraction from the anatomy–pathology terms and image content information; (2) similarity computation to measure the similarity between the cases in terms of feature vectors; (3) result refinement to rerank the candidate cases according to the feedbacks extracted from the neighbourhood information among the cases

NC is the number of cases and NT is the number of terms. Term frequency $TF(t, c)$ evaluates the frequency of the term T_t occurred on the case C_c , which is

$$TF(t, c) = \frac{occ(t, c)}{\sum_{t \in [1, NT]} occ(t, c)}. \quad (14.1)$$

Inverse document (case) frequency $IDF(t)$ indicates whether the term T_t is common or rare across all cases, which is

$$IDF(t) = \log\left(\frac{\sum_{c \in [1, NC]} occ(t, c)}{1 + occ(t, c)}\right). \quad (14.2)$$

TF-IDF measure of T_t for C_c is then computed as

$$TF-IDF(t, c) = TF(t, c) \times IDF(t). \quad (14.3)$$

Case C_c is finally formulated as a vector of TF-IDF measures of all terms as

$$V_{TF-IDF}(c) = (TF-IDF(1, c), \dots, TF-IDF(NT, c)). \quad (14.4)$$

The Euclidean distance between the vectors is then computed. We conducted a k-NN method for retrieval, which means selecting the cases that have the closest feature vectors to the one of the queries $V_{TF-IDF}(q)$ in terms of Euclidean distance.

14.2.2 Semantics Retrieval

While the TF-IDF method merely utilizes the direct co-occurrence relationship between the terms and cases, this relationship can be further used to infer the semantic information and can provide a more discriminative description of these terms for similarity computation. The latent semantic topic model is one of the most representative methods that can automatically extract the semantic information based on the co-occurrence relationship. It assumes that each image can be considered as a mixture of latent topics, and the latent topic is a probability distribution of terms. In this study, we applied probabilistic latent semantic analysis (pLSA) [22], which is a widely used latent topic extraction technique, for learning the latent semantics.

The schema of pLSA is shown in Fig. 14.2. pLSA considers that the observed probability of a term T_t occurring on a case C_c can be expressed with a latent or unobserved set of latent topics $Z = \{z_h | h \in [1, H]\}$, where H is the number of latent topics, as:

$$P(t|c) = \sum_h P(t|z_h) \cdot P(z_h|c). \quad (14.5)$$

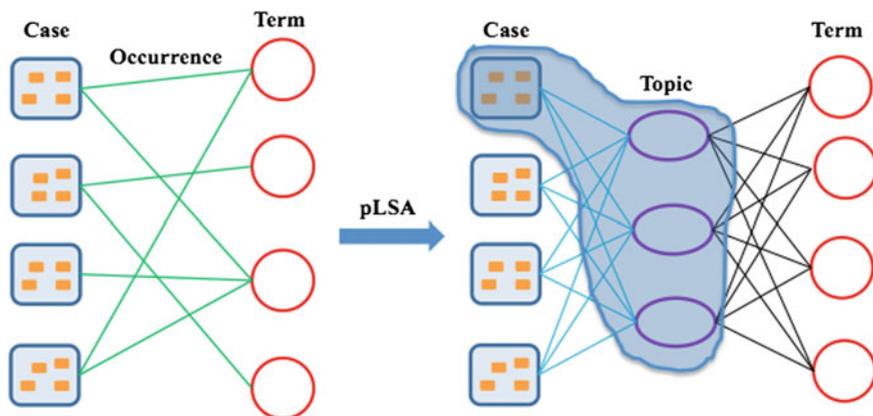


Fig. 14.2 Latent topic generation with pLSA

The probability $P(z_h|c)$ describes the distribution of latent topics given a certain case. The latent topics Z can be learnt by fitting the model with the expectation-maximization (EM) [25] algorithm that maximizes the likelihood function L :

$$L = \prod_t \prod_c P(t|c)^{occ(t,c)}. \quad (14.6)$$

After the latent topic extraction, each case is represented as the probability vector of the extracted latent topics,

$$V_{pLSA}(c) = (P(z_1|c), \dots, P(z_H|c)), \quad (14.7)$$

where each element is the probability of the latent topic given this case. The similarity between different cases is then measured by the Euclidean distance between the probability vectors, followed by the k-NN method for retrieval as introduced in Sect. 14.2.1. During the experiments, we empirically fixed the number of latent topics to 20, i.e. $H = 20$.

14.2.3 BoVW Retrieval

Unlike the aforementioned text-based methods, the visual content-based retrieval computes the similarity between the images based on their visual characteristics, such as the texture and colour. In the literature, there are many methods that can automatically extract the visual features to characterize the medical images [2, 26–28]. The Bag of Visual Words (BoVW) [29, 30] method, which is one of the popular methods for visual content-based image retrieval, is applied as our first content-based

retrieval method. The BoVW model represents an image with a visual word frequency histogram that is obtained by assigning the local visual features to the closest visual words in the dictionary. Rather than matching the visual feature descriptors directly, the BoVW-based approaches compare the images according to the visual words that are assumed to have higher discriminative power [29].

Specifically, the scale invariant feature transform (SIFT) [31] descriptors are extracted from the image to obtain a collection of local patch features for each image/case. The entire patch feature set computed from all images in the database is then grouped into clusters, e.g. with k-means method. Each cluster is regarded as a visual word W , and the whole cluster collection is considered as the visual dictionary $D = \{W_d | d \in [1, ND]\}$, where ND is the size of dictionary. Following that, all patch features in one image are assigned to the visual words, generating a visual word frequency histogram to represent this image (case) as,

$$V_{BoVW}(c) = (fre(1, c), \dots, fre(ND, c)), \quad (14.8)$$

where $fre(d, c)$ is the frequency of visual word W_d on case C_c . Finally, the similarity between images is computed based on these frequency histograms for retrieval.

In our experiments, the SIFT [31] descriptors were extracted from each scan of the 3D volume from the axial view. A visual dictionary of size 100, i.e. $ND = 100$ that could be sufficient for capturing local visual details and does not introduce too much noise based on our previous study in medical image analysis [23, 24], was computed with k-means. During the retrieval, given the ROI of a query case, we traversed all possible subregions (of the same size as the ROI in terms of the pixels) in a candidate volume in sliding window manner. Two subregions can be overlapped with an interval of 10 pixels at X/Y/Z directions. The subregion that has the smallest Euclidean distance from the query ROI in terms of visual word frequency histograms was regarded as the most similar area of the candidate to the query ROI, while the other subregions were not used. The distance between the two regions represented the similarity between the query and candidate images in our study. The k-NN method was applied for retrieval considering the obtained similarities.

14.2.4 Retrieval Result Refinement

While the first two steps form a basic retrieval process, relevance feedback refines the retrieval results if the top-ranked items are not fully satisfactory. Relevance feedback is based on the preferences upon the initial retrieval results, which can be provided by the users. However, providing manual feedback can be quite challenging due to the huge amount of image data. The relevance can also be affected since manual interpretation sometimes could be error-prone. The neighbourhood among images on the other hand can be used as a form of relevance feedback and is expected to be beneficial for image retrieval.

Algorithm 1 Pseudo code for preference and relativity computation

Input: Number of iterations T , neighborhood matrix A .
Output: Preference and relativity values.
1: **initialize** $rel_0 = 1$ and $pref_0 = 1$.
2: **for** each it in $[1, IT]$ **do**
3: **for** each C_{cr} **do**
4: Compute $pref_{it}(C_{cr})$ based on $rel_{it-1}(C_{cc})$ using Eq.(14.9);
5: **end for**;
6: **for** each C_{cc} **do**
7: Compute $rel_{it}(C_{cc})$ based on $pref_{it}(C_{cr})$ using Eq.(14.10);
8: **end for**;
9: L2-normalize $pref_{it}$ of all retrieved items.
10: L2-normalize rel_{it} of all candidates.
11: **end for**;
12: **return** $pref_{it}$ and rel_{it} .

Based on the results of the BoVW method, we further conducted a retrieval result refinement process based on our recent work [32]. In our method, we assume that the similarity relationship between the initial retrieved results and the remaining candidates can be used as a relevance feedback for result refinement. For a given query image, we first get a ranked list of initial retrieval results based on the BoVW model. Then, the similarities between the retrieved items and all candidates are used to evaluate their *preference* and *relativity*.

Formally, a preference score $pref(C_{cr})$ for the retrieved item C_{cr} is defined to evaluate the preference upon C_{cr} with regard to the query, i.e. relevance and irrelevance. A relativity score $rel(C_{cc})$ is appointed to the candidate image C_{cc} , indicating the similarity of C_{cc} to the query. The two values are computed conditioned on each other regarding the query case C_{cq} : the relativity score $rel(C_{cc})$ of C_{cc} would be high if it is similar to the highly preferred retrieved item C_{cr} , and the preference score $pref(C_{cr})$ of C_{cr} would be high if it is close to the more relevant candidate C_{cc} . The relativity score of C_{cc} is formulated as the sum of preference scores of its neighbouring retrieved items, similar to the preference score of C_{cr} . Denoting rel and $pref$ as the vectors of relativity and preference scores, we have the following formulations:

$$pref(C_{cr}) = \sum_{C_{cc}:A(C_{cc},C_{cr})=1} rel(C_{cc}), \quad (14.9)$$

$$rel(C_{cc}) = \sum_{C_{cr}:A(C_{cc},C_{cr})=1} pref(C_{cr}), \quad (14.10)$$

where A is a matrix indicating the bipartite neighbourhood relationship between the retrieved items and the candidates, i.e. $A(C_{cc}, C_{cr}) = 1$ if C_{cc} is the neighbour of C_{cr} ; otherwise, $A(C_{cc}, C_{cr}) = 0$. Equation (14.9) and (14.10) can be alternatively solved iteratively to obtain the relativity and preference scores as shown in Algorithm 1.

We then ranked all candidate images based on their relativity scores, in which the top-ranked ones were regarded as the most similar cases to the query. For our experiments, we selected the top 30 volumes based on the BoVW outputs as the initial results C_{cr} . Then, a bipartite relationship between the initial results C_{cr} and the remaining candidates C_{cc} , which represented the neighbourhood, was constructed by keeping the top 30 candidates for each initial result. The iterative ranking method [32] was applied to recompute the similarity score of each candidate with an iteration number IT of 20, after which the relativity and preference scores tended to be stable and have insignificant influence on the ranking orders of the candidates.

14.2.5 Fusion Retrieval

It is often suggested that the combination of textual and visual features can improve the retrieval performance [18]. Many fusion strategies have been proposed in the past such as maximum combination [34], sum combination [34] and Condorcet fusion [35].

Given the results from the text- and content-based retrievals, we conducted the fusion retrieval by using the sum combination method, which has been effective for textual and visual feature fusion [33]. To do this, a normalization step was firstly incorporated to normalize the similarity scores obtained from the aforementioned results, as:

$$S' = \frac{S - S_{min}}{S_{max} - S_{min}}, \quad (14.11)$$

where S_{min} and S_{max} are the lowest and highest similarity scores obtained within a certain method. The sum combination was then adopted to compute a fusion score for each candidate, as:

$$S_F = \sum_{r \in [1,4]} S'_r, \quad (14.12)$$

where $r \in [1, 4]$ represents the first four methods. The ones with the higher scores were for the results of fusion retrieval.

14.3 Results and Discussion

To evaluate the performance of retrieval results, medical experts were invited to perform relevance assessment of the top-ranked cases for each run. Various evaluation measures were used considering the top-ranked X cases, including the precision for top-ranked 10 and 30 cases (P@10, P@30), mean uninterpolated average precision (MAP), bpref measure and the R-precision.

Figure 14.3 displays the retrieval result for each of the topics given the aforementioned measures. The performances were diverse across the cases. It can be generally observed that better results were obtained for topics 1 and 7 when compared to the

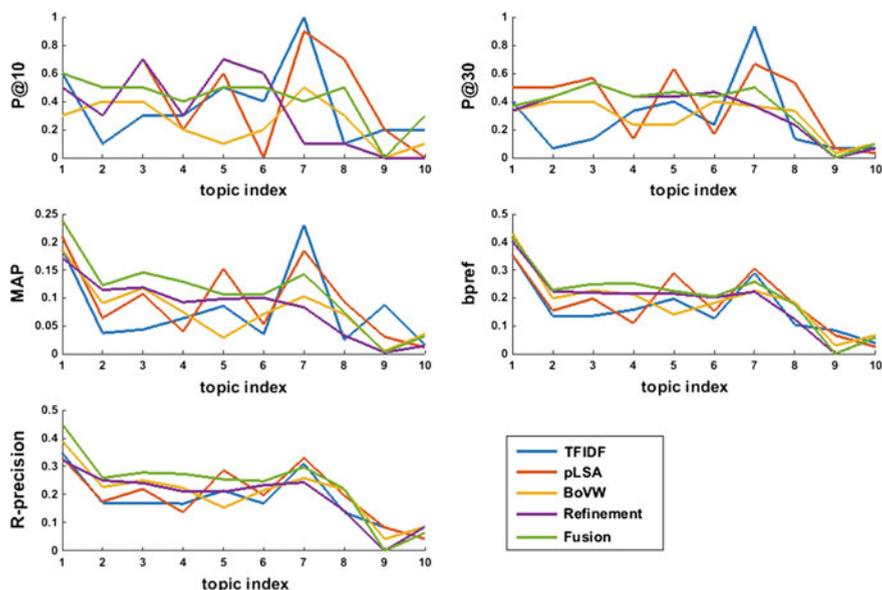


Fig. 14.3 Retrieval results of the 10 topics given different evaluation measures

Table 14.1 Average results of the different measures across the 10 queries

	P@10	P@30	MAP	bpref
TFIDF	0.370	0.277	0.081	0.162
pLSA	0.410	0.380	0.094	0.183
BoVW	0.250	0.283	0.078	0.190
Refinement	0.330	0.330	0.083	0.188
Fusion	0.420	0.353	0.110	0.207
Text	0.570	0.497	0.194	0.322
Image	0.330	0.330	0.083	0.188
Mixed	0.688	0.638	0.283	0.340

other topics, but the results for topics 9 and 10 were unfavourable. The differences were due to the different affected regions. Our methods computed the similarity between cases using the entire volumes, instead of focusing on the local details. Therefore, for cases that have relatively smaller annotated regions (the 3D bounding box of the ROI) compared to the others, e.g. case 10, the retrieval performance tended to be less favourable.

Table 14.1 shows the average results of the measures across the 10 queries, with the first five rows from our results and the last three rows showing the best results from all participants of the VISCERAL retrieval benchmark. Within our text-based approaches, pLSA generated better performance when compared to the

TF-IDF method, by further using the latent semantic information inferred from the co-occurrence relationship between cases and terms. Regarding the content-based retrieval, we obtained better results when applying the result refinement. Across the four methods, better performance was obtained from the text-based retrieval when compared to the content-based retrieval. The content-based methods use the visual content characteristics that may have large variation between the relevant cases but small difference between the irrelevant ones. The SIFT feature used in our experiments is widely known for capturing the local image content information, but it sometimes can be hard for SIFT to recognize the subtle visual difference between different images. In addition, while the size of the dictionary was set to 100 in our experiments, it can be varied for different datasets and potentially affect the retrieval performance. The text-based approaches on the other hand compare the different cases directly based on the pathology terms and affected anatomies. Thus, the text-based retrieval obtained the more favourable retrieval results. While the anatomy–pathology terms provide an overall description for the similarity computation, the visual content feature can better capture the local anatomical differences between cases. Therefore, the fusion approach achieved the overall best result, which is in accordance with the findings in the literature. Regarding the comparisons across all VISCERAL Retrieval Benchmark participations, we had the best performance with the result refinement among all image-based methods. The results from the text and fusion methods were less favourable since only co-occurrence information between the terms were used. Further analysis of the terms in the benchmark relating to the entire anatomy–pathology RadLex term collection would be helpful for retrieval improvements.

14.4 Conclusion

In this chapter, we introduced the approaches from our joint research team of USYD/HES-SO to address the VISCERAL Retrieval Benchmark, including the TF-IDF and pLSA methods for text-based retrieval, the BoVW and its result refinement for content-based retrieval, and the fusion retrieval of the above methods. The experimental results are in accordance with the findings in the literature, i.e. the text-based approaches typically perform better than purely visual content-based methods, and the combination of text- and content-based retrieval can achieve improved retrieval performance.

A further potential exploration could be the parameter selection. In this study, we empirically selected the settings of the parameters based on our previous work on other medical image retrieval tasks, such as the number of topics in the semantic retrieval, the size of dictionary in the BoVW retrieval and the number of initial retrieved items in the retrieval result refinement. It would be interesting to learn the parameters within the VISCERAL Retrieval Benchmark dataset but can be difficult due to the large amount of image data and the current lack of ground truth annotations. Another direction can be investigating a better way to combine the textual and image

content information. While the fusion retrieval tended to generate better performance in our study in general, we can also observe that the semantic retrieval overperformed the fusion method, e.g. the precision for top-ranked 30 cases (P@30). We expected a better performance if the feature extraction could utilize both textual- and image-content information rather than analysing them individually.

References

1. Doi K (2006) Diagnostic imaging over the last 50 years: research and development in medical imaging science and technology. *Phys Med Biol* 51:R5–R27
2. Müller H, Michoux N, Bandon D, Geissbuhler A (2004) A review of content-based image retrieval systems in medicine—clinical benefits and future directions. *Int J Med Inf* 73:1–23
3. Cai W, Feng D, Fulton R (2000) Content-based retrieval of dynamic PET functional images. *IEEE Trans Inf Technol Biomed* 4(2):152–158
4. Song Y, Cai W, Eberl S, Fulham MJ, Feng D (2011) Thoracic image case retrieval with spatial and contextual information. In: *IEEE international symposium biomedical imaging (ISBI)*, pp 1885–1888
5. Kumar A, Kim J, Cai W, Fulham M, Feng D (2013) Content-based medical image retrieval: a survey of applications to multidimensional and multimodality data. *J Digit Imaging* 26(6):1025–1039
6. Zhang S, Yang M, Cour T, Yu K, Metaxas D (2014) Query specific rank fusion for image retrieval. *IEEE Trans Pattern Anal Mach Intell* 47(4):803–815
7. Müller H, Antoine R, Arnaud G, Jean-Paul V, Antoine G (2005) Benefits of content-based visual data access in radiology 1. *Radiographics* 25(3):849–858
8. Song Y, Cai W, Zhou Y, Wen L, Feng D (2013) Pathology-centric medical image retrieval with hierarchical contextual spatial descriptor. In: *IEEE international symposium on biomedical imaging (ISBI)*, pp 202–205
9. Song Y, Cai W, Eberl S, Fulham MJ, Feng D (2010) A content-based image retrieval framework for multi-modality lung images. In: *IEEE international symposium on computer-based medical systems (CBMS)*, pp 285–290
10. El-Naqa I, Yang Y, Galatsanos NP, Nishikawa RM, Wernick MN (2004) A similarity learning approach to content-based image retrieval: application to digital mammography. *IEEE Trans Med Imaging* 23:1233–1244
11. Zhang F, Song Y, Cai W, Lee M-Z, Zhou Y, Huang H, Shan S, Fulham MJ, Feng D (2014) Lung nodule classification with multi-level patch-based context analysis. *IEEE Trans Biomed Eng* 61(4):1155–1166
12. Foncubierta-Rodríguez A, Depeursinge A, Müller H (2012) Using multiscale visual words for lung texture classification and retrieval. In: Müller H, Greenspan H, Syeda-Mahmood T (eds) *MCBR-CDS 2011*. LNCS, vol 7075. Springer, Heidelberg, pp 69–79. doi:[10.1007/978-3-642-28460-1_7](https://doi.org/10.1007/978-3-642-28460-1_7)
13. Cai W, Kim J, Feng D (2008) Content-based medical image retrieval. Elsevier book section 4:83–113
14. Squire D, Müller W, Müller H, Raki J (1999) Content-based query of image databases, inspirations from text retrieval: inverted files, frequency-based weights and relevance feedback. In: *The 11th scandinavian conference on image analysis*, pp 143–149
15. Müller H, Deserno TM (2011) Content-based medical image retrieval. In: Deserno TM (ed) *Biomedical image processing*. Springer, Berlin, pp 471–494
16. Haas S, Donner R, Burner A, Holzer M, Langs G (2012) Superpixel-based interest points for effective bags of visual words medical image retrieval. In: Müller H, Greenspan H, Syeda-Mahmood T (eds) *MCBR-CDS 2011*. LNCS, vol 7075. Springer, Heidelberg, pp 58–68. doi:[10.1007/978-3-642-28460-1_6](https://doi.org/10.1007/978-3-642-28460-1_6)

17. Zhang F, Song Y, Cai W, Hauptmann AG, Liu S, Liu SQ, Feng DD, Chen M (2015) Ranking-based vocabulary pruning in bag-of-features for image retrieval. In: Chalup SK, Blair AD, Randall M (eds) ACALCI 2015. LNCS (LNAI), vol 8955. Springer, Cham, pp 436–445. doi:[10.1007/978-3-319-14803-8_34](https://doi.org/10.1007/978-3-319-14803-8_34)
18. Müller H, Kalpathy-Cramer J (2010) The ImageCLEF medical retrieval task at ICPR 2010 — information fusion to combine visual and textual information. In: Ünay D, Çataltepe Z, Aksoy S (eds) ICPR 2010. LNCS, vol 6388. Springer, Heidelberg, pp 99–108. doi:[10.1007/978-3-642-17711-8_11](https://doi.org/10.1007/978-3-642-17711-8_11)
19. Zhang F, Song Y, Cai W, Depeursinge A, Müller H (2015) USYD/HES-SO in the VISCERAL retrieval benchmark. In: Müller H, Jimenez del Toro OA, Hanbury A, Langs G, Foncubierta Rodríguez A (eds) Multimodal retrieval in the medical domain. LNCS, vol 9059. Springer, Cham, pp 139–143. doi:[10.1007/978-3-319-24471-6_13](https://doi.org/10.1007/978-3-319-24471-6_13)
20. Hanbury A, Müller H, Langs G, Weber MA, Menze BH, Fernandez TS (2012) Bringing the algorithms to the data: cloud-based benchmarking for medical image analysis. In: Catarci T, Forner P, Hiemstra D, Peñas A, Santucci G (eds) CLEF 2012. LNCS, vol 7488. Springer, Heidelberg, pp 24–29. doi:[10.1007/978-3-642-33247-0_3](https://doi.org/10.1007/978-3-642-33247-0_3)
21. Jones KS (1972) A statistical interpretation of term specificity and its application in retrieval. *J Doc* 28:11–21
22. Hofmann T (2001) Unsupervised learning by probabilistic latent semantic analysis. *Mach Learn* 42:177–196
23. Zhang F, Song Y, Cai W, Liu S, Liu S, Pujol S, Kikinis R (2016) ADNI: pairwise latent semantic association for similarity computation in medical imaging. *IEEE Trans Biomed Eng* 63(5):1058–1069
24. Zhang F, Song Y, Cai W, Hauptmann AG, Liu S, Pujol S, Kikinis R, Chen M (2016) Dictionary pruning with visual word significance for medical image retrieval. *Neurocomputing* 177:75–88
25. Heinrich G (2005) Parameter estimation for text analysis. Technical report
26. Zhang X, Liu W, Dundar M, Badve S, Zhang S (2015) Towards large scale histopathological image analysis: hashing-based image retrieval. *IEEE Trans Med Imaging* 34:496–506
27. Yang W, Lu Z, Yu M, Huang M, Feng Q, Chen W (2012) Content-based retrieval of focal liver lesions using bag-of-visual-words representations of single- and multiphase contrast-enhanced CT images. *J Digit Imaging* 25:708–719
28. Song Y, Cai W, Eberl S, Fulham MJ, Feng D (2011) Thoracic image matching with appearance and spatial distribution. In: The 33rd annual international conference of the IEEE engineering in medicine and biology society (EMBC), pp 4469–4472
29. Sivic J, Zisserman A (2003) Video Google: a text retrieval approach to object matching in videos. In: IEEE international conference on computer vision (ICCV), pp 1470–1477
30. Liu S, Cai W, Song Y, Pujol S, Kikinis R, Feng D (2013) A bag of semantic words model for medical content-based retrieval. In: The 16th international conference on MICCAI workshop on medical content-based retrieval for clinical decision support
31. Lowe DG (1999) Object recognition from local scale-invariant features. In: IEEE international conference on computer vision (ICCV), pp 1150–1157
32. Cai W, Zhang F, Song Y, Liu S, Wen L, Eberl S, Fulham M, Feng D (2014) Automated feedback extraction for medical imaging retrieval. In: IEEE international symposium on biomedical imaging (ISBI), pp 907–910
33. Zhou X, Depeursinge A, Müller H (2010) Information fusion for combining visual and textual image retrieval. In: International conference on pattern recognition (ICPR), pp 1590–1593
34. Fox EA, Shaw JA (1993) Combination of multiple searches. Text retrieval conference, pp 243–252
35. Montague M, Aslam JA (2002) Condorcet fusion for improved retrieval. In: Proceedings of the eleventh international conference on information and knowledge management (CIKM), pp 538–548

Open Access This chapter is licensed under the terms of the Creative Commons Attribution- Non-Commercial 2.5 International License (<http://creativecommons.org/licenses/by-nc/2.5/>), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

