

Introducing space and time in local feature-based endomicroscopic image retrieval

Barbara André^{1,2}, Tom Vercauteren¹, Aymeric Perchant¹, Anna M. Buchner³,
Michael B. Wallace³, and Nicholas Ayache²

¹Mauna Kea Technologies (MKT), Paris, France

²INRIA - Asclepios Research Project, Sophia Antipolis, France

³Mayo Clinic, Jacksonville, Florida, USA

Abstract. Interpreting endomicroscopic images is still a significant challenge, especially since one single still image may not always contain enough information to make a robust diagnosis. To aid the physicians, we investigated some local feature-based retrieval methods that provide, given a query image, similar annotated images from a database of endomicroscopic images combined with high-level diagnosis represented as textual information. Local feature-based methods may be limited by the small field of view (FOV) of endomicroscopy and the fact that they do not take into account the spatial relationship between the local features, and the time relationship between successive images of the video sequences. To extract discriminative information over the entire image field, our proposed method collects local features in a dense manner instead of using a standard salient region detector. After the retrieval process, we introduce a verification step driven by the textual information in the database and in which spatial relationship between the local features is used. A spatial criterion is built from the co-occurrence matrix of local features and used to remove outliers by thresholding on this criterion. To overcome the small-FOV problem and take advantage of the video sequence, we propose to combine image retrieval and mosaicing. Mosaicing essentially projects the temporal dimension onto a large field of view image. In this framework, videos, represented by mosaics, and single images can be retrieved with the same tools. With a leave-n-out cross-validation, our results show that taking into account the spatial relationship between local features and the temporal information of endomicroscopic videos by image mosaicing improves the retrieval accuracy.

1 Introduction

With the recent technology of probe-based confocal laser endomicroscopy (pCLE) [1], endoscopists are able to image tissues at microscopic level with a miniprobe, and in real time during ongoing procedure. However, as the acquired pCLE images are relatively new for them, the physicians are still in the process of defining a taxonomy of the pathologies in the images, for instance to differentiate benign tissues and neoplastic, i.e. pathological, tissues of colonic polyps, see Fig. 1 and 2 for an illustrative example of such images. To face this clinical challenge, a

valuable aid to the physician in establishing a diagnosis would be to provide endomicroscopic images that have a similar appearance to the image of interest and that have been previously diagnosed by expert physicians. Knowing that pathological tissue is characterized by some irregularities in the cellular and vascular architecture, we aim at retrieving texture information coupled with shape information by using local operators on pCLE images. To serve that purpose, we decided to investigate a modern method for content-based image retrieval (CBIR), the bag-of-visual words (BVW) method [2]. BVW has been successfully used in many applications of computer vision. For example, on a well-defined non medical application, by using this method on a large variety of images of natural or artificial textures, the authors of [2] obtained excellent recognition results that are close to 98%.

The standard BVW method detects salient regions in the images and extracts information only on these specific regions. However in pCLE images, the discriminative information is distributed over the entire image field. Contrary to classical methods that apply sparse detectors, we use a dense detector to collect densely the local features in the images. This overcomes the information sparseness problem. Moreover, pCLE images contain characteristic pattern at several scales, in particular the microscopic scale of individual cells and the mesoscopic scale of groups of cells. For this reason, we perform a bi-scale description of the collected image regions. Another problem is that the spatial relationship between the local features is lost in the standard BVW representation of an image, whereas the spatial organization of cells is highly discriminative in pCLE images. So we looked at measuring a statistical representation of this spatial geometry. This was achieved by exploiting the co-occurrence matrix of the visual words labeling the local features in the image. After the retrieval process, we introduce the measured spatial criterion in a verification step that allows to remove outliers from the retrieved pCLE images, which are given by the most similar to queried images. Taking into account the spatial relationship between local features is the main contribution of our study, it can be used as a generic tool for many applications of CBIR. Besides, we noticed that the FOV of single still pCLE images may not be large enough for the physicians to see a characteristic global pattern and make a robust diagnosis. As this limitation cannot be solved by the standard methods, we decided to take into account the time information of pCLE video sequences by considering them as objects of interest instead of still images. More precisely, we use image mosaicing [3] to project the temporal dimension of video sequences onto a large FOV image, cf. some resulting mosaics in Fig. 4 and 5. With a leave-n-out cross-validation, classification experiments on the pCLE database serve the validation of the methodology: our method outperforms other methods taken as references, by improving the classification accuracy and by providing more relevant training images among the first retrieved images.




Fig. 1: Typical image retrieval results provided by our method from two benign queries. **B.** indicates Benign and **N.** Neoplastic. From left to right on each row: the queried image, and its k-NNs on the top layer, and their respective colored visual words on the bottom layer. An outlier is indicated by **Outlier** if it has been rejected by the spatial verification process, and by **Error** otherwise. FOV of the images: 240 μm .




Fig. 2: Typical image retrieval results provided by our method from two neoplastic queries. **B.** indicates Benign and **N.** Neoplastic. From left to right on each row: the queried image, and its k-NNs on the top layer, and their respective colored visual words on the bottom layer. An outlier is indicated by **Outlier** if it has been rejected by the spatial verification process, and by **Error** otherwise. FOV of the images: 240 μm .

2 The bag-of-visual words method

As one of the most popular method for image retrieval, the BVW [2] method aims at extracting a local image description that is both efficient to use and invariant with respect to viewpoint changes, e.g., translations, rotations and scaling, and illumination changes, e.g., affine transformation of intensity. Its methodology consists in first finding and describing local features, then in quantizing them into clusters named visual words, and in representing the image by the histogram of these visual words. The BVW retrieval process can thus be decomposed into four steps: detection, description, clustering and similarity measuring, possibly followed by a classification step for image categorization.

The detection step extracts salient regions in the image, i.e. regions containing some local discriminative information. In particular, corners and blobs in the image can be detected by the sparse Harris-Hessian (H-H) operator around key-points with high responses of intensity derivatives. Other sparse detectors like the Intensity-Based Regions (IBR) and the Maximally Stable Extremal Regions (MSER) are also specialized for the extraction of blob features in the images. We refer the interested reader [4] for a survey of these detectors.

Then, each local region can be typically described by the Scale Invariant Feature Transform (SIFT) [5] descriptor. We refer the reader [2] for a survey of this and other powerful descriptors. At the description step, the SIFT descriptor computes, for each salient region, a description vector which is its gradient histogram at the optimal scale provided by the detector, the gradient orientations being normalized with respect to the principal orientation of the salient region. As a result, the image is represented in a high dimensional space by a set of SIFT description vectors that are invariant by translation, rotation and scale.

To reduce the dimension of the description space, the clustering step, for example based on a standard K-Means, builds K clusters, i.e. K visual words, from the union of the description vector sets gathered from all the N images of the training database. Since each description vector counts for one visual word, an image is represented by a signature of size K which is its histogram of visual words, normalized by the number of its salient regions.

Given these image signatures, it is possible to define a distance between two images as the χ^2 distance [4] between their signature and to retrieve the closest training images as the most similar to the image of interest. The relevance of the similarity results can be quantified by a further classification step, for instance based on a standard nearest neighbors procedure that weights the votes of the k -nearest neighbors by the inverse of their χ^2 distance to the signature of the queried image, so that the closest images are the most determinant. Besides, performing image classification is a way to validate a new retrieval method by comparing it with other methods.




Fig. 3: Worst image retrieval results provided by our method. The benign query on the top is a rare benign variety which is not represented in the training dataset. The neoplastic query on the bottom contains on its top left corner a partially visible elongated crypt which could not be totally described. **B.** indicates Benign and **N.** Neoplastic. From left to right on each row: the queried image, and its k-NNs on the top layer, and their respective colored visual words on the bottom layer. An outlier is indicated by *Outlier* if it has been rejected by the spatial verification process, and by *Error* otherwise. FOV of the images: 240 μ m.

3 Including spatial and temporal information

When we applied the standard BVW method on pCLE images [6], we obtained rather poor classification results and the presence of many retrieval outliers. To improve the accuracy of endoscopic image retrieval, we decided to include both spatial and temporal information contained in the pCLE images. While testing on pCLE images the numerous sparse detectors listed by [4], we first

observed that a large number of salient regions sparsely extracted by these standard detectors do not persist between two highly correlated successive images taken from the same video. This is consistent with the rather poor results of the standard sparse detector shown in Section 4. To overcome the persistence problem and take into account all the information in the images, we use a dense detector contrarily to the standard method. Furthermore, a dense detector is relevant because local information appears to be densely distributed over the entire field of the pCLE image. The dense detector is made of overlapping disks localized on a dense regular grid, such that each disk covers a possible image pattern at microscopic level.

We also noticed that the endoscopists establish their diagnosis on pCLE images from the regularity of the cellular architecture in the colonic tissue [7], where goblet cells and crypts are both round-shaped characteristic patterns, but where a crypt has larger size than its surrounding goblet cells so it must not be recognized as the same object. In order to be sensitive to scale changes, our method looked at describing local disk regions at various scales that are not automatically computed, for example by choosing a microscopic scale for individual cell patterns and a mesoscopic scale for larger groups of cells. This leads us to represent an image by several sets of description vectors that are scale-dependent, resulting in several signatures for the image that are then concatenated into one larger signature.

This previous observation also suggests that the spatial organization of the goblet cells must be included in the retrieval process because it is substantial to differentiate benign tissues from neoplastic tissues. The authors of [8] previously proposed adding a geometrical verification to take spatial information into account, however their method is based on the assumption that they want to retrieve images of the exact same scene, which is not the case for our application. In like manner, our idea is to introduce a geometrical verification process after the retrieval process, but based on the assumption that the spatial relationships between the local features are only statistically the same in the images with similar appearance. To introduce spatial information, we took advantage of the dense property to define the adjacency between two visual words as the 8-adjacency between the two disk regions that are labeled by them on the detection grid. Thus, we are able to store in a co-occurrence matrix M of size $K \times K$ the probability for each pair of visual words of being adjacent to each other. In order to best differentiate the images of the benign class from the images of the pathological class, we looked at the most discriminative linear combination W of some elements m of M . This is achieved by a linear discriminant analysis (LDA) which uses the textual diagnostic information in the database. The LDA weights are given by $W = \Sigma^{-1} (\mu_1 - \mu_2)$, where Σ is the covariance matrix of the elements m of M in all training images and μ_i is the mean of the elements m of M in the training image belonging to the class i . From these weights W , we computed the spatial criterion $\alpha = Wm$ for each retrieved image and compared it with the α value of the image of interest. By thresholding the α value during a

verification process, outliers are rejected and the first retrieved training images are more relevant.

Expert physicians pointed out that some characteristic global patterns are too partially visible on single still pCLE images to make a robust diagnosis: two still images may have a very similar appearance but be attached to contradictory diagnoses [9]. Fig. 3 shows some bad retrieval results caused by the small-FOV problem. To address this problem, the time dimension of pCLE videos needs to be exploited, by including in the retrieval process the temporal relationship between successive images from the same video sequence. The study reported by [10] proposes a method for video retrieval using the spatial layout of the image regions, but this method has been designed for object matching, which is not our objective. Since successive frames from pCLE videos are only related by viewpoint changes, our approach uses the image mosaicing of [3] to project the temporal dimension of a video sequence onto one image with a larger FOV and of higher resolution. Thus, mosaics can be queried and retrieved in the same way as still images.

4 Experiments and discussion

At the Mayo Clinic in Jacksonville, the Cellvizio[®] system, MKT, Paris, was used to image colonic polyps during surveillance colonoscopies in 54 patients. On each acquired video sequence, expert physicians established a pCLE diagnosis [7] that differentiates pathological sequences from benign ones. The video sequences contain from 5 to over a thousand frames and each frame is an image of diameter 500 pixels corresponding to a FOV of 240 μm . To build our pCLE database, we considered a subset of these sequences by discarding those whose quality was insufficient to perform a reliable diagnosis. In each of the remaining 52 video sequences, we selected groups of successive frames according to the length of the sequence. The resulting database is composed of $N = 1036$ still pCLE images and $N' = 66$ pCLE mosaics, half of the data coming from benign sequences and half from pathological ones.

In pCLE images, the disk regions containing information at mesoscopic scale have a radius value $\rho_1 = 40$, while the radius value of those containing information at microscopic scale is $\rho_2 = 15$. For the dense detector, we then chose $\delta = 20$ pixels of grid spacing in order to get a reasonable overlap between adjacent regions. Among the values from 30 to 1500 found in the literature for the number K of visual words provided by the K-Means clustering, the value $K = 100$ yielded satisfying classification results. To prevent overfitting, as the size of our pCLE database is still rather small, especially concerning the number of mosaics, the number of LDA weights in the computation of the spatial criterion α had to be restricted. For the elements m of the co-occurrence matrix M , we only considered the K diagonal elements of the matrix M build from the visual words of large radius 40, observing that the overlapping regions of radius 40 have a sufficient spatial correlation, better than those of radius 15. The good values of the threshold θ_α were chosen by analysing the distribution of α across




Fig. 4: Typical mosaic retrieval results provided by our method from one benign query. **B.** indicates Benign and **N.** Neoplastic. From left to right on each row: the queried mosaic, and its k -NNs on the top layer, and their respective colored visual words on the bottom layer. FOV of the mosaics: from 260 μm to 1300 μm .




Fig. 5: Typical mosaic retrieval results provided by our method from one neoplastic query. **B.** indicates Benign and **N.** Neoplastic. From left to right on each row: the queried mosaic, and its k-NNs on the top layer, and their respective colored visual words on the bottom layer. FOV of the mosaics: from 260 μm to 1300 μm .

the benign and pathological images: 2 when retrieving still images from a queried still image and 0.5 when retrieving still images from a queried mosaic.

The classification results of our method are presented in Fig. 6 on the left and compared with the following methods taken as references: the standard sparse scale invariant SIFT method, the statistical approach of Haralick features [11, 12] and the texture retrieval method of Textons [13]. To ensure a non-biased classification, our validation scheme retrieves k nearest images in the training set with training images not belonging to the video sequence of the image being queried, i.e. a leave- n -out cross-validation where n is the number of frames in the video of the queried image. According to the accuracy, sensitivity and specificity rates yielded by each method on the still images of the pCLE database, our retrieval method including spatial information is the most efficient, with an accuracy rate of 78.2% for $k = 22$ neighbors, which is 11.5 points better than the standard SIFT method. The gain of accuracy can be decomposed in 10.2 points for the choice of a dense detector and a bi-scale SIFT description, and 1.3 points for the verification process on the spatial criterion. It is also worth mentioning that with the spatial verification, fewer nearest neighbors are necessary to classify the query at a given accuracy. For $k = 4$ neighbors, some illustrative examples of the image retrieval results are shown in Fig. 1, 2 and 3.

Moreover, when including both spatial and temporal information by querying mosaics, our classification results are much better, see Fig. 6 on the right. Since mosaics contain more information than single images, their content-based neighborhood is more representative of their pathological neighborhood, so they can be better classified by a smaller number k' of nearest neighbors. Indeed, if we retrieve still images for queried mosaics, the classification accuracy is 83.3% for $k' = 10$ neighbors, which demonstrates the robustness of our retrieval method applied on heterogeneous data with different resolution. For the retrieval of still images from queried mosaics, the poor specificity can be explained by the fact that a mosaic annotated as neoplastic may contain some benign patterns which induce the retrieval of single benign images and classify it as benign. However the expert physicians diagnose a pCLE video sequence as neoplastic as soon as it contains neoplastic patterns, even when some benign tissue is imaged. Besides, if we retrieve mosaics for queried mosaics, the classification accuracy is 89%. Thus, even though we only have a small number of mosaics, including time dimension in mosaics provides us proof of concept results for endomicroscopic video retrieval. The relevance of mosaic retrieval can also be qualitatively appreciated in Fig. 4 and 5. For the retrieval of mosaics from queried mosaics, including the spatial information does not improve the classification results because of the overfitting phenomenon: indeed, the number of LDA weights, 100, is bigger than the total number of mosaics, $N' = 66$.




Fig. 6: Left: Classification accuracies yielded by several methods on the still images of the pCLE database, with leave-n-out cross-validation. Right: Results for k nearest neighbors, where **M** means **Mosaic** and **I** means **Image** in the configuration (*Queried, Retrieved*). Our proposed method is referred to as Dense Bi-Scale Verif (**DBV**) if it includes spatial verification and Dense Bi-Scale (**DB**) otherwise.

5 Conclusion

Using visual similarity between a given image and medically interpreted images allowed us to provide the physicians with semantic similarity, and thus could potentially support their diagnostic decision. Although our experiments are focused on a relatively small training dataset, the classification results constitute a validation of our generic methodology. By taking into account the spatio-temporal relationship between the local feature descriptors, the first retrieved endomicroscopic images are much more relevant.

For future work, a larger training database would not only improve the classification results if all the characteristics of the image classes are better represented, but also enable the exploitation of the whole co-occurrence matrix of visual words at several scales. Besides, the learning step of the retrieval process could leverage the textual information of the database and incorporate the spatial information of multi-scale co-occurrence matrices into descriptors. The co-occurrence matrix could also be better analyzed by other tools more generic than linear discriminant analysis. For example, a more complete spatial geometry between local features could be learned by estimating the parameters of Markov Random Fields [14]. We also plan to perform a leave-one-patient-out

cross-validation to ensure fully unbiased retrieval and classification results. As for introducing the temporal information, a more robust approach would not only consider the fused image of a mosaic but the $2D + t$ volume of the registered frames composing the mosaic to work on more accurate visual words and better combine spatial and temporal information.

References

1. Wallace, M., Fockens, P.: Probe-based confocal laser endomicroscopy. *Gastroenterology* **136**(5) (May 2009) 1509–1513
2. Zhang, J., Lazebnik, S., Schmid, C.: Local features and kernels for classification of texture and object categories: a comprehensive study. *International Journal of Computer Vision* **73** (June 2007) 213–238
3. Vercauteren, T., Perchant, A., Malandain, G., Pennec, X., Ayache, N.: Robust mosaicing with correction of motion distortions and tissue deformation for in vivo fibered microscopy. *Medical Image Analysis* **10**(5) (October 2006) 673–692
4. Mikolajczyk, K., Tuytelaars, T., Schmid, C., Zisserman, A., Matas, J., Schaffalitzky, F., Kadir, T., Van Gool, L.: A comparison of affine region detectors. *International Journal of Computer Vision* **65** (November 2005) 43–72
5. Lowe, D.: Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* **60** (November 2004) 91–110
6. André, B., Vercauteren, T., Perchant, A., Wallace, M.B., Buchner, A.M., Ayache, N.: Endomicroscopic image retrieval and classification using invariant visual features. In: *Proceedings of the IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI'09)*. (2009) 346–349
7. Buchner, A., Ghabril, M., Krishna, M., Wolfen, H., Wallace, M.: High-resolution confocal endomicroscopy probe system for in vivo diagnosis of colorectal neoplasia. *Gastroenterology* **135**(1) (July 2008) 295
8. Jegou, H., Douze, M., Schmid, C.: Hamming embedding and weak geometric consistency for large scale image search. In: *Proceedings of the 10th European Conference on Computer Vision (ECCV'08)*. Volume I. (October 2008) 304–317
9. Becker, V., Vercauteren, T., von Weyern, C.H., Prinz, C., Schmid, R.M., Meining, A.: High resolution miniprobe-based confocal microscopy in combination with video-mosaicing. *Gastrointestinal Endoscopy* **66**(5) (November 2007) 1001–1007
10. Sivic, J., Zisserman, A.: Efficient visual search for objects in videos. *Proceedings of the IEEE* **96** (April 2008) 548–566
11. Haralick, R.: Statistical and structural approaches to texture. In: *Proceedings of the IEEE*. Volume 67. (1979) 786–804
12. Srivastava, S., Rodriguez, J., Rouse, A., Brewer, M., Gmitro, A.: Computer-aided identification of ovarian cancer in confocal microendoscope images. *Journal of Biomedical Optics* **13**(2) (March/April 2008) 024021
13. Leung, T., Malik, J.: Representing and recognizing the visual appearance of materials using three-dimensional textons. *International Journal of Computer Vision* **43** (June 2001) 29–44
14. Descombes, X., Morris, R., Zerubia, J., Berthod, M.: Estimation of Markov random field prior parameters using Markov chain Monte Carlo maximum likelihood. *IEEE Transactions on Image Processing* **8**(7) (July 1999) 954–963