# Immediate ROI Search for 3-D Medical Images

Karen Simonyan[1], Marc Modat[2], Sebastien Ourselin[2], David Cash[2], Antonio Criminisi[3], Andrew Zisserman[1]

[1]University of Oxford, UK
{karen,az}@robots.ox.ac.uk
[2]Centre for Medical Image Computing, University College London, UK
{m.modat,s.ourselin,d.cash}@ucl.ac.uk
[3]Microsoft Research, Cambridge, UK
antcrim@microsoft.com

**Abstract.** The objective of this work is a scalable, real-time, visual search engine for 3-D medical images, where a user is able to select a query Region Of Interest (ROI) and automatically detect the corresponding regions within all returned images.
We make three contributions: (i) we show that with appropriate off-line processing, images can be retrieved and ROIs registered in real time; (ii) we propose and evaluate a number of scalable exemplar-based image registration schemes; (iii) we propose a discriminative method for learning to rank the returned images based on the content of the ROI. The retrieval system is demonstrated on MRI data from the ADNI dataset [9], and it is shown that the learnt ranking function outperforms the baseline.

**Keywords:** Immediate structured search, visual search, ROI, exemplar-based registration, learning to rank

## 1 Introduction

Throughout the last decade there has been a rapid growth of medical image repositories. Medical images and corresponding clinical cases, stored in these large collections, capture a wide range of disease population variability due to numerous covariates (diagnosis, age, co-morbidities, etc). Instant image retrieval from such repositories could be of great value for clinical practice, e.g. by providing a "second opinion" based on the corresponding diagnostic information or course of treatment. Apart from the processing speed, another important aspect of a practical retrieval system is the ability to focus the search on a particular part (structure) of the image which is of most interest.

This paper addresses the problem of immediate structured image retrieval in large repositories of 3-D medical images. Given a query 3-D image (e.g. from a new patient we wish to diagnose) and a user-drawn Region Of Interest (ROI) in it, we seek to retrieve repository images with the ROI automatically located, and rank them based on a clinically relevant score, driven by the content of the ROI. Instant ROI localisation in large repositories is achieved by off-line preprocessing of the repository based on fast image registration. Figure 1 shows screenshots of our brain MRI retrieval system.

The contribution of this paper is three-fold. First, we show that 90 (and potentially more) images can be retrieved and ROIs registered in real time. Second, we present and evaluate several modifications of scalable exemplar-based image registration. Finally, we propose a technique for learning to rank the retrieved ROI. We envisage a number of applications of the proposed framework, and discuss three of them below.

**Atrophy-Aware Brain MRI Retrieval.** Structural MRI data has been shown to provide reliable quantification of the atrophy process in the brain caused by Alzheimer's disease (AD) [5] or other neurodegenerative disorders. There are numerous natural history studies, the Alzheimer's Disease Neuroimaging Initiative (ADNI) [9], launched in 2003, being the most prominent. The immediate visual search engine can aid in differential diagnosis, as there are discriminating patterns between numerous forms of dementia. The ability to focus the search on specific anatomical regions, that have been identified as being sensitive to the disease, is an important advantage of structured (ROI-level) visual search as opposed to retrieval based on global cues. For example, the hippocampal deterioration is increasingly being considered as a way of identifying subjects who have a higher risk of developing AD. Providing the images with relevant ROI and their respective diagnosis to clinicians will aid in their decision process. We give an implementation of a search-engine for this application in this paper.

**Lesion Retrieval in CT Scans.** The wide application of computed tomography to lesion detection (e.g. liver or kidney lesions) has led to the collection of large quantities of imaging data together with corresponding clinical reports. Recently, image retrieval frameworks [11] have been proposed, which can help clinicians to search for similar lesions in image repositories. However, such methods do not take into account the relative location of the lesion inside the liver, which can be an important search criterion (e.g. the query "find all visually similar lesions in the same part of liver"). To process such queries, the geometrical correspondence of a query ROI in target images should be quickly obtained, which can be done using exemplar-based registration employed in the proposed framework.

**Image Quality Control.** Another application of the visual search engine is the quality control of incoming images for research studies and clinical trials. This predominantly manual task is one of the most time consuming areas of the processing pipeline. Even though the failure rate is low, all data needs to be reviewed by radiologists and images with poor quality must be excluded from the analysis. Typically, this task consists of a careful qualitative review of each image independently. If the reviewers were provided with a visual search engine to retrieve similar (on ROI-level) images from the repository and the outcomes of their image quality review, this would speed up the process dramatically.

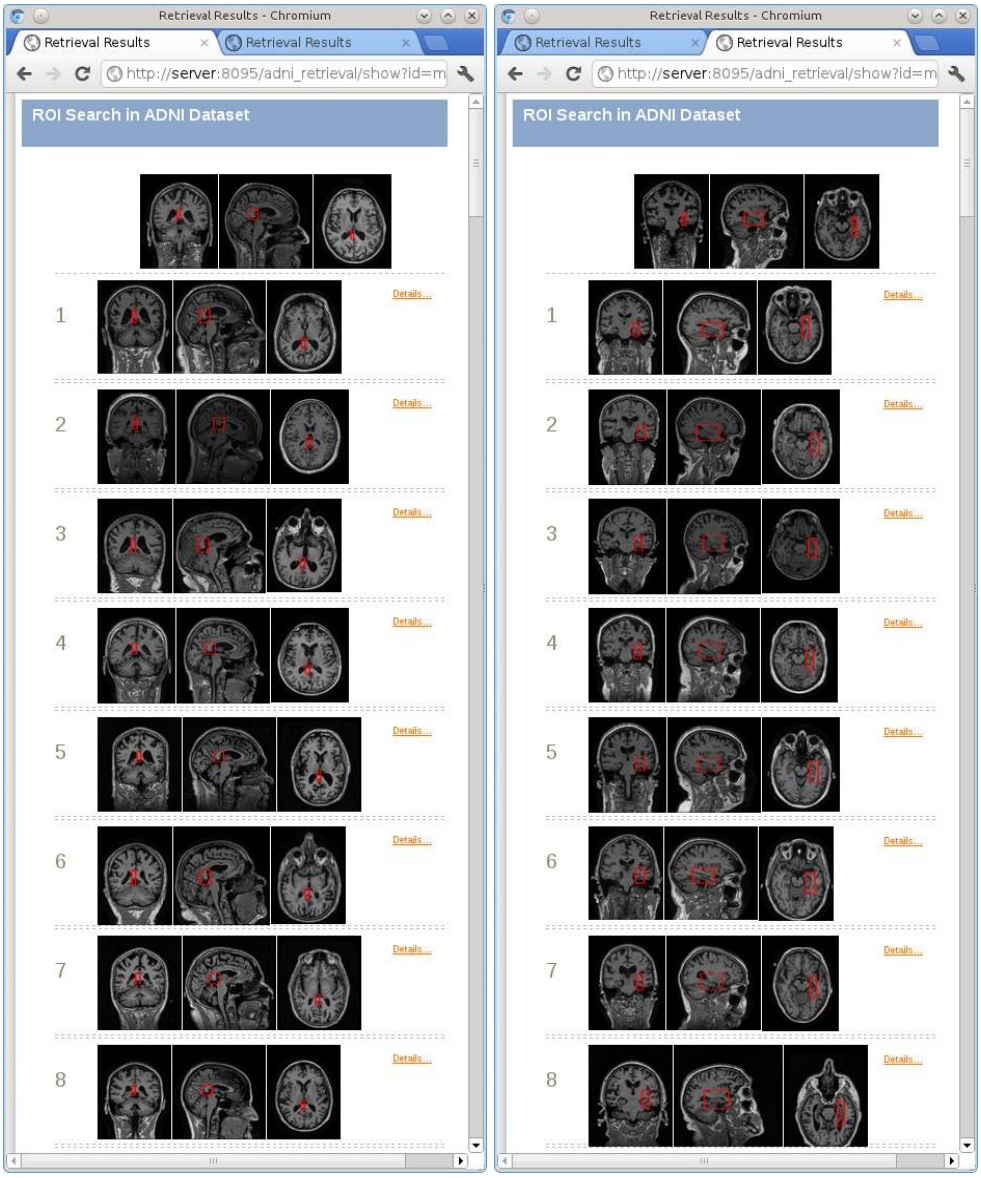


**Fig. 1.** Screenshots of our brain MRI retrieval system for two different queries and the top eight retrieval results. The system is accessed via a conventional Web browser. The top of the page shows the three orthogonal views of the user-specified query volume and axis-aligned ROI outlined in red. The high ranked retrieved volumes with the corresponding ROI (in red) are shown below.

### 1.1 Related Work

The problem of content-based medical image retrieval has a vast literature. Most conventional approaches [10] consist in retrieving images that are *globally* sim-

ilar to the query image. Recently, the problem of ROI-level search has been addressed in [1, 2, 14]. In [1, 2], an approach related to that of [15] was employed, which consists in computing an ROI descriptor vector (known as a bag of visual words), and retrieving images based on the distance between query and target ROI descriptors. It should be noted that such a technique discards information about the spatial location of the query ROI in an image. While this can be an advantage when searching in natural images or videos [15], in the case of medical images this can result in the retrieved ROIs lying in completely different anatomical locations, which is often undesirable. To circumvent this, augmenting of visual desriptors with their spatial location was proposed in [1]. A more principled approach is taken in the framework of [14], which makes use of the fact that medical images are usually acquired under standardised protocols with a fixed viewpoint, field of view, etc. It allows one to quickly compute registrations between query and repository images, making the target ROI detection trivial. We build upon their framework and review it in more detail below.

**Outline.** The rest of the paper is organised as follows. Section 2.1 contains the details of our brain MRI dataset and the retrieval system implementation. In Sect. 3 we present several exemplar-based registration techniques and evaluate them. In Sect. 4 we describe a learning to rank framework.

## 2 ROI Retrieval Framework

To enable immediate ROI retrieval at run time, processing is divided into off-line and on-line parts, as summarised in Fig. 2. A similar approach was previously applied to 2-D X-ray images retrieval [14]. The key idea is to pre-compute registrations between repository images off-line, so that at run time the correspondences of a query ROI in target images can be obtained immediately if the query image is taken from the repository. Once the regions of interest have been aligned in repository images, they can be ranked based on an application-specific clinically relevant score.

If a query image is not in the repository, it is added there by registering it with repository images. This brings up the issue of computational efficiency in the case of large datasets. In [14] an efficient exemplar-based registration technique was proposed to solve this problem. It requires only a small fixed number of registrations to exemplar images to be computed. The transformations to the rest of the repository is then obtained by computationally cheap transform composition. It should be noted that registration with exemplar images can be performed using *any* off-the-shelf method suitable for a particular type of images. In this paper, we apply the framework to a large dataset of brain MRI scans, where images exhibit similar fields of view.

### 2.1 Dataset and Implementation Details

Our dataset consists of 90 brain MRI scans randomly selected from the ADNI dataset [9] (`http://www.loni.ucla.edu/ADNI/Data/`). The subset contains an

1. **On-line (given a user-specified query volume and ROI bounding box)**
   - Using the pre-computed registration and transform composition (Sect. 3), compute the ROIs corresponding to the query ROI in all repository images.
   - Rank the retrieved ROIs using a clinically meaningful ranking function of choice (Sect. 4).
2. **Off-line (pre-processing)**
   - Compute registration between exemplar images and all other images (Sect. 3)

**Fig. 2.** The retrieval algorithm outline.

equal number of images (30) of each of the three subject groups: Alzheimer's disease, control, and MCI (mild cognitive impairment). For the evaluation of the methods proposed in the paper, for each of these images we computed the "gold standard" parcellation and pairwise registrations. We note that this is not required for the functioning of the proposed search engine. The parcellation into 83 brain anatomical structures was performed using the method of [4]. The non-rigid registration has been performed using the Free-Form Deformation approach [12]. Briefly, it consists of a cubic B-Spline parametrisation model where the Normalised Mutual Information (NMI) is used as a measure of similarity. We used an efficient implementation [8] that is freely available as a part of the NiftyReg package. We provide a default ranking function based on the contents of the ROI (the ranking function can also be learnt, Sect. 4). For the default, we measure the $\chi^2$ distance between the brain tissue type distributions in query and target ROI. The distributions were computed using the GMM-based probabilistic segmentation algorithm [3].

Our retrieval system is implemented as a Web-based application, which can be accessed from any device equipped with a Web browser (thin client paradigm). The system is split into a front-end and back-end. The front-end, implemented in Python and JavaScript, allows a user to select a query volume, specify arbitrary axis-aligned ROI in it, and explore the retrieval results. The back-end is currently implemented in unoptimised Python, leaving a lot of room for potential speed-up. The average ROI registration time using five exemplar images (Sect. 3) on a single CPU core is 0.06s per image, which allows for retrieval of hundreds of MRI volumes under 1s when rolled out on a multi-core server.

In certain use cases, using multiple query ROI can be beneficial, as it would allow one to select several relevant areas in a query image. Here we consider a single query ROI, the extension to multiple ROI being rather straightforward. We also restrict the ROI to be an axis-aligned 3-D bounding box, but in general any ROI shape is possible.

## 3 Exemplar-Based Registration

Carrying out non-rigid registration of the query image with each of the target images scales badly with the number of repository images as 3-D image reg-

istration is computationally complex, and the number of registrations equals the number of images. Moreover, storing all pairwise registrations is prohibitive due to high storage requirements of non-rigid transforms (e.g. B-spline warps computed over a dense 3-D grid).

The key idea behind scalable exemplar-based registration is that instead of registering a query image with each of the repository images by pairwise registration, the query is registered with only a few *fixed* images (called exemplars), which effectively define several reference spaces. The remaining repository images are already pre-registered with exemplars, so they can be registered with the query by composing the two transforms. Finally, to obtain a single correspondence from several exemplars, the composed transforms are aggregated. The exemplar-based registration is schematically illustrated in Fig. 3 (left).

More formally, for a dataset of $N$ images, a query image $I_q$ is registered with only a subset of $K = \text{const}$ exemplar images, which results in $K$ transforms $T_{q,k},\ k = 1 \dots K$. The transformations $T_{k,t}$ between an exemplar $I_k$ and each of the remaining repository images $I_t$ are pre-computed. Then the transformation between images $I_q$ and $I_t$ can be obtained by composition of transforms (computed using different exemplars) followed by aggregation:

$$T_{q,t}(\mathbf{x}) = f\left(\{T_{k,t} \circ T_{q,k}\}\right)(\mathbf{x}) \tag{1}$$

where $\mathbf{x}$ is a point in the query image and $f$ is the aggregation function.

The advantage of exemplar-based registration scheme is that for a query image only $K \ll N$ registrations should be computed, and the transform composition complexity is negligible. Thus pairwise registrations between all images can be computed in $O(KN)$ rather than $O(N^2)$. The same estimates apply to the storage requirements for the computed registrations, which allows them to be stored in RAM for fast access. Compared to group-wise registration algorithms, transform composition does not rely on the computation of a group mean model, and is scalable in the case of rapidly growing datasets. Additionally, the use of several transformations instead of one improves the registration robustness.

There are two important choices to make: how to select the exemplars and how to define the function $f$, aggregating the transforms obtained using different exemplars. In [14] the exemplars were selected randomly, and the aggregation was performed by taking a median. In both cases the accuracy of registration is not taken into account, which can potentially lead to the selection of exemplars which can not be accurately registered with other images. If the ratio of such exemplars is large, the median filter will not be able to recover the correspondence.

Here we investigate different ways of exemplar selection and transform aggregation. We assume that the registration error $d_{ij} = d(I_i, I_j)$ for two images $I_i$ and $I_j$ belongs to the range $[0, 1]$ with 0 corresponding to a perfect registration. In general, the error can be computed using different cues (intensity, deformation field smoothness, re-projection error, etc.). In our experiments, we employed inverse normalised mutual information, rescaled to the $[0, 1]$ range.

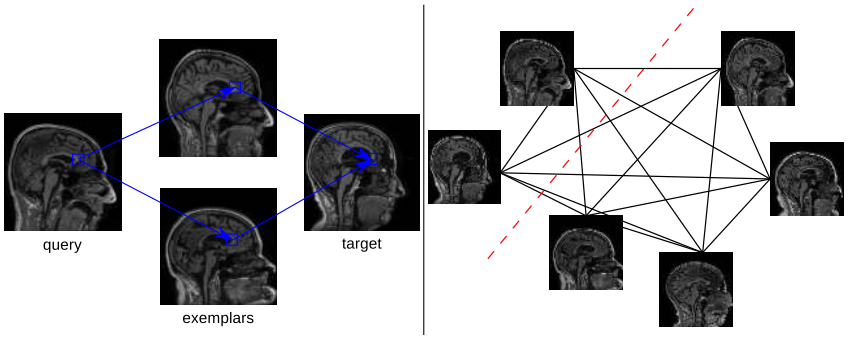


**Fig. 3.** *Left*: Exemplar-based registration. *Right*: Repository graph. The red dashed line illustrates the graph cut into exemplars and other images.

### 3.1 Exemplar Images Selection

The objective of exemplar selection is to pick a fixed number ($K$) of repository images, such that they can be accurately registered with the remaining ones. In this section we describe several ways of deterministic (non-random) exemplar selection from a set of images based on their pairwise registration errors. For instance, the exemplars selection can be carried out on the full repository or its subset. We stress that deterministic exemplar selection (including pairwise registrations) is performed off-line and has no impact on the query times; random exemplar selection does not require any additional processing at all.

It is natural to represent an image repository as a fully-connected graph with vertices corresponding to images and edges weighted by the registration errors, as shown in Fig. 3 (right). We employ the repository graph formalism to describe the objective functions for deterministic exemplar selection.

**Min-Sum Selection.** The task of exemplar images selection can be formulated as a min-cut problem on a repository graph. Indeed, we aim at splitting the set of vertices (images) into two partitions such that the sum of edge weights (registration errors) between vertices lying in these partitions is minimal. The resulting optimisation problem is as follows:

$$\alpha = \arg\min_{\alpha} \sum_{i,j} \alpha_i (1 - \alpha_j) d_{ij}, \quad \text{s.t.} \sum_i \alpha_i = K \tag{2}$$

where $\alpha \in \{0,1\}^N$ is a binary vector such that $a_k = 1$ iff $k$-th repository image is selected as an exemplar. The optimisation of (2) is NP-hard. While efficient approximate solutions exist [7], we leave their evaluation for future work. In this paper, we evaluate exemplar selection based on the following simplified objective:

$$\alpha = \arg\min_{\alpha} \sum_{i,j} \alpha_i d_{ij}, \quad \text{s.t.} \sum_i \alpha_i = K \tag{3}$$

which corresponds to selecting $K$ exemplars such that the sum of edge weights between them and *all* other images (including other exemplars) is minimal. The solution of (3) can be easily obtained by ranking the images in the ascending order of $q_i = \sum_j d_{ij}$ and then selecting the top-$K$ images as exemplars.

**Spectral Clustering Selection.** Another approach to deterministic exemplar selection is based on clustering the repository images into $K$ clusters followed by the selection of a single exemplar in each of these clusters. Given the pairwise similarity matrix $s_{ij} = 1 - d_{ij}$, we use the normalised cuts technique [13] to split the vertices (images) into a set of classes such that the similarity between images in different clusters is small, and the similarity between images in the same cluster is large. This corresponds to computing a normalised cut in the repository graph. Once the images are divided into clusters, a single exemplar is selected in each of the clusters as the image with minimal sum of registration errors to the others (3).

### 3.2 Shortest Path Aggregation

Once the exemplars are selected and fixed, the way of aggregating several registrations into one should be defined (function $f$ in (1)). In general, taking the mean or median does not account for the exemplars registration error, which can be large for certain pairs of query and target images. One of the possible ways to account for these errors is to pick a single registration which corresponds to the shortest path in the graph from the query to the target vertices and goes through exactly one exemplar (Fig. 3, left). In other words, for a given (query, target) pair of images, only one exemplar is selected, which has the lowest sum of registration errors with these images:

$$\begin{aligned} &f(q,t)(\mathbf{x}) = \left(T_{s,t} \circ T_{q,s}\right)(\mathbf{x}), \\ &s = \arg\min_k d_{qk} + d_{kt} \end{aligned} \tag{4}$$

### 3.3 Evaluation

In this section, we compare the registration accuracy of different combinations of exemplar selection and transform aggregation techniques. For exemplar selection, we consider random selection ("rand"), "min-sum" selection, and spectral clustering selection (Sect. 3.1). For transform aggregation, "median", "mean", and the shortest path exemplar ("single", Sect. 3.2) are compared. The evaluation was performed on the dataset described in Sect. 2.1, which was randomly split into 45 training and 45 testing images. Exemplar selection was performed on the training set, registration evaluation – on the test set. The experiment was repeated three times. For each pair of test images, the accuracy of registration was assessed using two criteria. First, we measured the mean distance (in mm) between points projected using pairwise (between query and target) and

exemplar-based transformations. The measure describes how different exemplar-based registration is from the pairwise registration. The points were selected to be the centers of mass of the 83 parcellated anatomical structures. The second measure is the mean overlap ratio (Jaccard coefficient) of 83 anatomical structure bounding boxes, projected from the query image to the target image, with the bounding boxes in the target image. We used the bounding boxes of the parcellated anatomical structure volumes instead of the volumes themselves because it more closely follows the search engine use case scenario, where we operate on the level of bounding boxes. We note that this measure is noisy due to the possible inaccuracies of the "gold standard" parcellation. In Table 1 we report the mean and standard deviation of the two measures across all test image pairs for different number $K$ of exemplar images.

Based on the presented results, we can conclude that all three exemplar selection methods (including the random choice) exhibit similar levels of performance when coupled with robust median aggregation. Aggregation based on shortest path selection performs worse, and the mean aggregation is the worst. The reason for such a behaviour could be that the global registration error, which we used for exemplar selection, does not account for the local inaccuracies. Another reason for similar performance can be the lack of strong image variation in our dataset. At the same time, using a single exemplar ($K = 1$) results in worse accuracy compared to several exemplar images. The accuracy of exemplar-based registration with median aggregation is at the same level as that of pairwise registration without exemplars. The average distance between the points projected using the two registrations is less than 1.4 mm. Considering its low computational complexity, in our practical implementation we used the randomised selection of $K = 5$ exemplars and the median aggregation of the composed transforms.

**Table 1.** Exemplar-based registration accuracy. The overlap ratio of pairwise registration (without exemplars) is $0.568 \pm 0.076$. For the overlap ratio higher is better, and for the distance smaller means closer to the direct registration without exemplars.

| exemplar | aggregation | overlap ratio | | | distance (mm) | | |
|---|---|---|---|---|---|---|---|
| | | $K = 1$ | $K = 5$ | $K = 7$ | $K = 1$ | $K = 5$ | $K = 7$ |
| rand | mean | 0.555 ±0.072 | $0.532 \pm 0.073$ | $0.53 \pm 0.073$ | 2.04 ±0.28 | $1.44 \pm 0.22$ | $1.38 \pm 0.21$ |
| | median | | $0.569 \pm 0.076$ | $\mathbf{0.571 \pm 0.076}$ | | $1.45 \pm 0.23$ | $1.37 \pm 0.23$ |
| | single | | $0.557 \pm 0.073$ | $0.559 \pm 0.073$ | | $1.99 \pm 0.26$ | $1.98 \pm 0.25$ |
| min-sum | mean | 0.557 ±0.072 | $0.531 \pm 0.072$ | $0.529 \pm 0.072$ | 1.94 ±0.26 | $1.42 \pm 0.22$ | $1.37 \pm 0.22$ |
| | median | | $0.569 \pm 0.076$ | $0.57 \pm 0.076$ | | $1.43 \pm 0.23$ | $1.36 \pm 0.23$ |
| | single | | $0.558 \pm 0.072$ | $0.556 \pm 0.072$ | | $1.94 \pm 0.26$ | $2.00 \pm 0.32$ |
| cluster | mean | | $0.531 \pm 0.072$ | $0.529 \pm 0.072$ | | $1.44 \pm 0.22$ | $1.39 \pm 0.22$ |
| | median | | $0.569 \pm 0.076$ | $0.57 \pm 0.076$ | | $1.45 \pm 0.23$ | $1.38 \pm 0.23$ |
| | single | | $0.556 \pm 0.072$ | $0.556 \pm 0.072$ | | $2.03 \pm 0.32$ | $2.03 \pm 0.31$ |

## 4 Learning to Rank Retrieved ROI

In this section we propose a framework for discriminative learning of ROI ranking. In general, we aim at automatically reproducing the ROI ranking provided by a clinician in the form of preference constraints. That is, for a particular query image and ROI the expert selects *pairs* of retrieved images (with corresponding ROI) such that the first image of the pair should be ranked higher than the second one. More formally, the annotation is represented as a set of triplets: $\{(I_q, R_q), (I_h, R_h), (I_l, R_l)\}_i$ where $(I_q, R_q)$ are the query image and ROI, and the (image, ROI) pair $(I_h, R_h)$ should be ranked higher than $(I_l, R_l)$.

We propose to learn a distance in the space of ROI such that ranking of the retrieved ROI based on their distance to the query ROI satisfies the ground-truth preference constraints with a margin [6], i.e.

$$d(R_q^i, R_h^i) + 1 < d(R_q^i, R_l^i) \tag{5}$$

where $d$ is the distance between ROIs in the feature space (we omit image notation for brevity). We constrain the distance to be a generalised squared Mahalanobis distance of the form:

$$d_A(R_u, R_v) = (\phi_u - \phi_v)^T A(\phi_u - \phi_v), \tag{6}$$

where $A \succeq 0$ is a positive semi-definite matrix to be learnt, and $\phi_u$ is the feature vector of the ROI $R_u$ (discussed later). It can be shown that the distance (6) equals the squared Euclidean distance in the projected space defined by a projection matrix $W$ such that $W^T W = A$. Taking into account the preference constraints (5), the large-margin learning objective takes the form:

$$A = \arg\min_{A \succeq 0} \sum_i \max\left(d_A(R_q^i, R_h^i) - d_A(R_q^i, R_l^i) + 1, 0\right) + \frac{\lambda}{2}\|A\|_F^2 \tag{7}$$

where the first term is a sum of ranking hinge losses, and the second term is a Frobenius (Euclidean) norm of the matrix $A$. The parameter $\lambda > 0$, balancing the two parts, is selected on a validation set. The max-margin formulation is closely related to the LMNN formulation of [17]. The cost function (7) is strongly convex and can be efficiently optimised by projected stochastic sub-gradient method.

**ROI Feature Vector.** The proposed ROI distance learning framework is generic and can be applied to different ROI representations. Here we consider the "bag of words" representation [15] which consists in computing visual descriptors inside the ROI, assigning them to a nearest cluster ("visual word" from a vocabulary), and then accumulating the assignments inside an ROI into a histogram of words. The visual words vocabulary is computed using k-means clustering. The visual descriptors, the corresponding vocabulary, and visual word assignments can be pre-computed off-line. At query time, only the histogram over ROI should be computed, which can be done quickly using integral volumes.

**Evaluation.** Here we describe a preliminary experiment which was carried out on the dataset of Sect. 2.1. Considering that expert-annotated preference constraints are not currently available for our data (it is a subject of future research), we used the clinical diagnosis class labels ("AD"(Alzheimer's) and "Control") to generate the constraints of the form (5). Namely, we constrain the distance between same-class ROIs to be smaller than the distance between different-class ROI. As the ROI, we used a bounding box in the hippocampal area, which is known to be relevant to the Alzheimer's disease. Bag of words was computed using single-scale dense textons [16] of size $3 \times 3 \times 3$mm which were quantised into 512 visual words, leading to a 512-D ROI feature vector. We randomly selected 30 images for training, 10 for validation, and 20 for testing. The experiment was repeated three times. Mean average precision of retrieval was measured to be 58.8% using Euclidean distance between feature vectors (i.e. $A = I$), and 63.8% using the learnt distance. This shows that metric learning can indeed improve the retrieval performance. We believe that with appropriate preference constraints annotation and more sophisticated visual features, the results of the proposed learning-to-rank framework can be further improved.

## 5 Summary

In this paper we presented a practical structured image search framework, capable of instant retrieval of brain MRI volumes and corresponding regions of interest from large datasets. Fast ROI alignment in repository images was made possible using scalable exemplar-based registration technique.

The evaluation of different exemplar-based registration methods has shown that random exemplar image selection coupled with robust median transform aggregation achieves registration accuracy on par with optimised exemplar selection and pairwise registration without exemplars. We note that in the case of diverse non-uniform image datasets, deviant images can be unrepresented in the exemplar set. In that case, our conclusions might not be immediately applicable.

Finally, we have presented a discriminative distance learning framework for ranking retrieved ROIs. It was demonstrated that it can indeed improve the ranking performance. It should be noted that while the proposed ranking function has been learnt on a specific anatomical area (hippocampal area), the same approach can be used to learn more generic ranking functions. A web-based demo of 3-D ROI retrieval framework is available at `http://www.robots.ox.ac.uk/~vgg/research/med_search/`

**Acknowledgements.** This work was supported by Microsoft Research PhD Scholarship Program, ERC grant VisRec no. 228180, and CBRC grant no. 168. Data collection and sharing for this project was funded by the Alzheimer's Disease Neuroimaging Initiative (ADNI) (NIH Grant U01 AG024904).

## References

1. Avni, U., Greenspan, H., Konen, E., Sharon, M., Goldberger, J.: X-ray categorization and retrieval on the organ and pathology level, using patch-based visual words. IEEE Trans. Med. Imag. 30(3), 733–746 (2011)

2. Burner, A., Donner, R., Mayerhoefer, M., Holzer, M., Kainberger, F., Langs, G.: Texture bags: Anomaly retrieval in medical images based on local 3d-texture similarity. In: Müller, H., Greenspan, H., Syeda-Mahmood, T.F. (eds.) MCBR-CDS. LNCS, vol. 7075, pp. 116–127. Springer, Heidelberg (2011)
3. Cardoso, M.J., Clarkson, M.J., Ridgway, G.R., Modat, M., Fox, N.C., Ourselin, S.: LoAd: A locally adaptive cortical segmentation algorithm. NeuroImage 56(3), 1386–1397 (2011)
4. Cardoso, M., Modat, M., Ourselin, S., Keihaninejad, S., Cash, D.: Multi-STEPS: Multi-label similarity and truth estimation for propagated segmentations. In: Math. Meth. in Biomed. Im. Anal., IEEE Workshop on, pp. 153–158 (2012)
5. Jack, C.R., Shiung, M.M., Gunter, J.L., O'Brien, P.C., Weigand, S.D., Knopman, D.S., Boeve, B.F., Ivnik, R.J., Smith, G.E., Cha, R.H., Tangalos, E.G., Petersen, R.C.: Comparison of different MRI brain atrophy rate measures with clinical disease progression in AD. Neurology 62(4), 591–600 (2004)
6. Joachims, T.: Optimizing search engines using clickthrough data. In: ACM SIGKDD Int. Conf. on Knowl. Disc. and Data Mining, pp. 133–142. ACM Press, New York (2002)
7. Lim, Y., Jung, K., Kohli, P.: Energy minimization under constraints on label counts. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) ECCV (2). LNCS, vol. 6312, pp. 535–551. Springer, Heidelberg (2010)
8. Modat, M., Taylor, Z., Barnes, J., Hawkes, D., Fox, N., Ourselin, S.: Fast free-form deformation using graphics processing units. Comp. Meth. and Prog. in Biomed. 98(3), 278–284 (2010)
9. Mueller, S., Weiner, M., Thal, L., Petersen, R., Jack, C., Jagust, W., Trojanowski, J., Toga, A., Beckett, L.: Ways toward an early diagnosis in Alzheimer's disease: The Alzheimer's Disease Neuroimaging Initiative (ADNI). Alzheimer's and Dementia 1(1), 55–66 (2005)
10. 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 Med. Inform. 73(1), 1–23 (2004)
11. Napel, S.A., Beaulieu, C.F., Rodriguez, C., Cui, J., Xu, J., Gupta, A., Korenblum, D., Greenspan, H., Ma, Y., Rubin, D.L.: Automated retrieval of CT images of liver lesions on the basis of image similarity: method and preliminary results. Radiology 256(1), 243–252 (2010)
12. Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L.G., Leach, M.O., Hawkes, D.J.: Nonrigid registration using free-form deformations: application to breast MR images. IEEE Trans. Med. Imag. 18(8), 712–721 (1999)
13. Shi, J., Malik, J.: Normalized cuts and image segmentation. IEEE Trans. on Patt. Anal. and Mach. Intell. 22(8), 888–905 (2000)
14. Simonyan, K., Zisserman, A., Criminisi, A.: Immediate structured visual search for medical images. In: Fichtinger, G., Martel, A.L., Peters, T.M. (eds.) MICCAI (3). LNCS, vol. 6893, pp. 288–296. Springer, Heidelberg (2011)
15. Sivic, J., Zisserman, A.: Video Google: A text retrieval approach to object matching in videos. In: IEEE Int. Conf. on Comp. Vis., pp. 1470–1477. IEEE Press, New York (2003)
16. Varma, M., Zisserman, A.: Texture classification: Are filter banks necessary? In: IEEE Int. Conf. on Comp. Vis. and Pat. Rec. vol. 2, pp. 691–698. IEEE Press, New York (2003)
17. Weinberger, K., Saul, L.: Distance metric learning for large margin nearest neighbor classification. J. Mach. Learn. Res. 10, 207–244 (2009)