

## Faster R-CNN Features for Instance Search

Amaia Salvador, Xavier Giró-i-Nieto, Ferran Marqués  
 Universitat Politècnica de Catalunya (UPC)  
 Barcelona, Spain  
 {amaia.salvador,xavier.giro}@upc.edu

Shin'ichi Satoh  
 National Institute of Informatics  
 Tokyo, Japan  
 satoh@nii.ac.jp

### Abstract

Image representations derived from pre-trained Convolutional Neural Networks (CNNs) have become the new state of the art in computer vision tasks such as instance retrieval. This work explores the suitability for instance retrieval of image- and region-wise representations pooled from an object detection CNN such as Faster R-CNN. We take advantage of the object proposals learned by a Region Proposal Network (RPN) and their associated CNN features to build an instance search pipeline composed of a first filtering stage followed by a spatial reranking. We further investigate the suitability of Faster R-CNN features when the network is fine-tuned for the same objects one wants to retrieve. We assess the performance of our proposed system with the Oxford Buildings 5k, Paris Buildings 6k and a subset of TRECVID Instance Search 2013, achieving competitive results.

### 1. Introduction

Visual media is nowadays the most common type of content in social media channels, thanks to the proliferation of ubiquitous cameras. This explosion of online visual content has motivated researchers to come up with effective yet efficient automatic content based image retrieval systems. This work addresses the problem of instance search, understood as the task of retrieving those images from a database that contain an instance of a query.

Recently, Convolutional Neural Networks (CNNs) have been proven to achieve state of the art performance in many computer vision tasks such as image classification [12, 22], object detection [19] or semantic segmentation [14]. CNNs trained with large amounts of data have been shown to learn feature representations that can be generic enough to be used even to solve tasks for which they had not been trained [18]. Particularly for image retrieval, many works in the literature [3, 25, 11] have adopted solutions based on off-the-shelf features extracted from a CNN pretrained for the task of image classification [12, 22, 24], achieving state of



Figure 1. Examples of the rankings and object locations obtained by our proposed retrieval system for query objects (left, depicted with a blue contour) of three different datasets: TRECVID INS 2013, Paris Buildings and Oxford Buildings.

the art performance in popular retrieval benchmarks.

Instance search systems often combine fast first filtering stages, in which all images in a database are ranked according to their similarity to the query, with more computationally expensive mechanisms that are only applied to the top retrieved items. Geometric verification and spatial analysis [10, 29, 15, 28] are common reranking strategies, which are often followed with query expansion (pseudo-relevance feedback) [1, 5].

Spatial reranking usually involves the usage of sliding windows at different scales and aspect ratios over an image. Each window is then compared to the query instance in order to find the optimal location that contains the query,

which requires the computation of a visual descriptor on each of the considered windows. Such strategy resembles that of an object detection algorithm, which usually evaluates many image locations and determines whether they contain the object or not. Object Detection CNNs [8, 9, 7, 19] have rapidly evolved to a point where the usage of exhaustive search with sliding windows or the computation of object proposals [26, 2] is no longer required. Instead, state of the art detection CNNs [19] are trained in an end-to-end manner to simultaneously learn object locations and labels.

This work explores the suitability of both off-the-shelf and fine-tuned features from an object detection CNN for the task of instance retrieval. We make the following three contributions:

- We propose to use a CNN pre-trained for object detection to extract convolutional features both at global and local scale in a single forward pass of the image through the network.
- We explore simple spatial reranking strategies, which take advantage of the locations learned by a Region Proposal Network (RPN) to provide a rough object localization for the top retrieved images of the ranking.
- We analyze the impact of fine-tuning an object detection CNN for the same instances one wants to query in the future. We find such a strategy to be suitable for learning better image representations.

This way, we put together a simple instance retrieval system that uses both local and global features from an object detection network. Figure 1 shows examples of rankings generated with our retrieval pipeline.

The remainder of the paper is structured as follows. Section 2 introduces the related works, Section 3 presents the methodology of this paper, including feature pooling, reranking and fine-tuning strategies. Section 4 includes the performed experiments on three different image retrieval benchmarks as well as the comparison to other state of the art CNN-based instance search systems. Finally, Section 5 draws the conclusions of this work.

## 2. Related Work

**CNNs for Instance Search.** Features from pre-trained image classification CNNs have been widely used for instance search in the literature. Early works in this direction demonstrated the suitability of features from fully connected layers for image retrieval [4]. Razavian *et al.* [18] later improved the results by combining fully connected layers extracted from different image sub-patches.

A second generation of works explored the usage of other layers in the pretrained CNN and found that convolutional layers significantly outperformed fully connected

ones at image retrieval tasks [21]. Babenko and Lempitsky [3] later proposed a compact descriptor composed of the sum of the activations of each of the filter responses in a convolutional layer. Tolias *et al.* introduced R-MAC [25], a compact descriptor composed of the aggregation of multiple region features. Kalantidis *et al.* [11] found significant improvements when applying non-parametric spatial and channel-wise weighting strategies to the convolutional layers.

This work shares similarities with all the former in the usage of convolutional features of a pretrained CNN. However, we choose to use a state-of-the-art object detection CNN, to extract both image- and region-based convolutional features in a single forward pass.

**Object Detection CNNs.** Many works in the literature have proposed CNN-based object detection pipelines. Girshick *et al.* presented R-CNN [8], a version of Krizhevsky’s AlexNet [12], fine-tuned for the Pascal VOC Detection data [6]. Instead of full images, the regions of an object proposal algorithm [26] were used as inputs to the network. At test time, fully connected layers for all windows were extracted and used to train a bounding box regressor and classifier.

Since then, great improvements to R-CNN have been released, both in terms of accuracy and speed. He *et al.* proposed SPP-net [9], which used a Spatial Pyramid based pooling layer to improve classification and detection performance. Additionally, they significantly decreased computational time by pooling region features from convolutional features instead of forward passing each region crop through all layers in the CNN. This way, the computation of convolutional features is shared for all regions in an image. Girshick later released Fast R-CNN [7], which used the same speed strategy as SPP-net but, more importantly, replaced the post-hoc training of SVM classifiers and box regressors with an end-to-end training solution. Ren *et al.* introduced Faster R-CNN [19], which removed the object proposal dependency of former object detection CNN systems by introducing a Region Proposal Network (RPN). In Faster R-CNN, the RPN shares features with the object detection network in [7] to simultaneously learn prominent object proposals and their associated class probabilities.

In this work, we take advantage of the end-to-end self-contained object detection architecture of Faster R-CNN to extract both image and region features for instance search.

## 3. Methodology

### 3.1. CNN-based Representations

This paper explores the suitability of using features from an object detection CNN for the task of instance search. In our setup, query instances are defined by a bounding box over the query images. We choose the architecture and pretrained models of Faster R-CNN [19] and use it as a feature

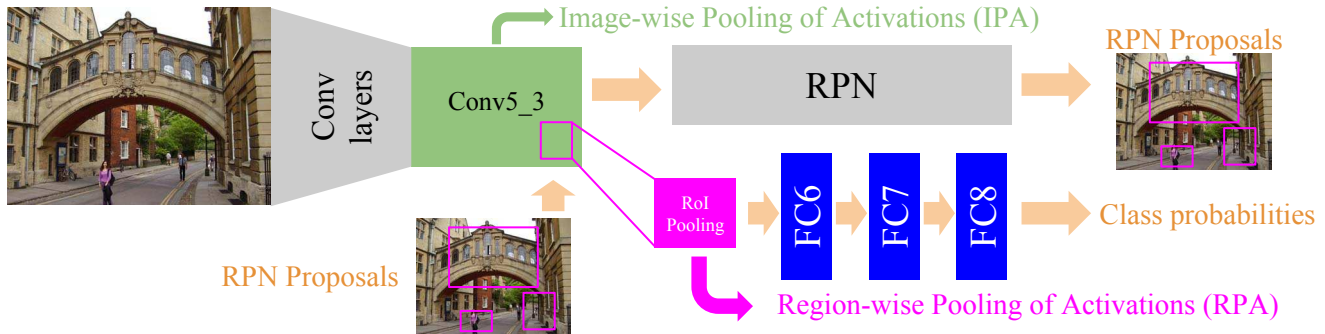


Figure 2. Image- and region-wise descriptor pooling from the Faster R-CNN architecture.

extractor at both global and local scales. Faster R-CNN is composed of two branches that share convolutional layers. The first branch is a Region Proposal Network that learns a set of window locations, and the second one is a classifier that learns to label each window as one of the classes in the training set.

Similarly to other works [3, 25, 11] our goal is to extract a compact image representation built from the activations of a convolutional layer in a CNN. Since Faster R-CNN operates at global and local scales, we propose the following strategies of feature pooling:

**Image-wise pooling of activations (IPA).** In order to construct a global image descriptor from Faster R-CNN layer activations, one can choose to ignore all layers in the network that operate with object proposals and extract features from the last convolutional layer. Given the activations of a convolutional layer extracted for an image, we aggregate the activations of each filter response to construct an image descriptor of the same dimension as the number of filters in the convolutional layer. Both max and sum pooling strategies are considered and compared in Section 4.3 of this paper.

**Region-wise pooling of activations (RPA).** After the last convolutional layer, Faster R-CNN implements a region pooling layer that extracts the convolutional activations for each of the object proposals learned by the RPN. This way, for each one of the window proposals, it is possible to compose a descriptor by aggregating the activations of that window in the RoI pooling layer, giving raise to the region-wise descriptors. For the region descriptor, both max and sum pooling strategies are tested as well.

Figure 2 shows a schematic of the Faster R-CNN architecture and the two types of descriptor pooling described above.

Following several other authors [3, 11], sum-pooled features are  $l_2$ -normalized, followed by whitening and a second round of  $l_2$ -normalization, while max-pooled features are just  $l_2$ -normalized once (no whitening).

### 3.2. Fine-tuning Faster R-CNN

This paper explores the suitability of fine-tuning Faster R-CNN to 1) obtain better feature representations for image retrieval and 2) improve the performance of spatial analysis and reranking. To achieve this, we choose to fine tune Faster R-CNN to detect the query objects to be retrieved by our system. This way, we modify the architecture of Faster R-CNN to output the regressed bounding box coordinates and the class scores for each one of the query instances of the tested datasets.

In our experiments, we explore two modalities of fine-tuning:

- **Fine-tuning Strategy #1:** Only the weights of the fully connected layers in the classification branch are updated (i.e. the convolutional layers and the RPN are left unchanged).
- **Fine-tuning Strategy #2:** Weights of all layers after the first two convolutional layers are updated. This way, convolutional features, RPN proposals and fully connected layers are modified and adapted to the query instances.

The resulting fine-tuned networks are to be used to extract better image and region representations and to perform spatial reranking based on class scores instead of feature similarities.

### 3.3. Image Retrieval

The three stages of the proposed instance retrieval pipeline are described in this section: filtering stage, spatial reranking and query expansion.

**Filtering Stage.** The Image-wise pooling (IPA) strategy is used to build image descriptors for both query and database images. At test time, the descriptor of the query image is compared to all the elements in the database, which are then ranked based on the cosine similarity. At this stage, the whole image is considered as the query.



**Spatial Reranking.** After the Filtering Stage, the top  $N$  elements are locally analyzed and reranked. We explore two reranking strategies:

- *Class-Agnostic Spatial Reranking (CA-SR).* For every image in the top  $N$  ranking, the region-wise descriptors (RPA) for all RPN proposals are compared to the region descriptor of the query bounding box. The region-wise descriptors of RPN proposals are pooled from the RoI pooling layer of Faster R-CNN (see Figure 2). To obtain the region-wise descriptor of the query object, we warp its bounding box to the size of the feature maps in the last convolutional layer and pool the activations within its area. The region with maximum cosine similarity for every image in the top  $N$  ranking gives the object localization, and its score is kept for ranking.
- *Class-Specific Spatial Reranking (CS-SR).* Using a network that has been fine-tuned with the same instances one wishes to retrieve, it is possible to use the direct classification scores for each RPN proposal as the similarity score to the query object. Similarly to CA-SR, the region with maximum score is kept for visualization, and the score is used to rank the image list.

**Query Expansion (QE).** The image descriptors of the top  $M$  elements of the ranking are averaged together with the query descriptor to perform a new search.

## 4. Experiments

### 4.1. Datasets

The methodologies described in Section 3 are assessed with the following datasets:

- **Oxford Buildings [16].** 5,063 images, including 55 query images of 11 different buildings in Oxford (5 images/instance are provided). A bounding box surrounding the target object is provided for query images.
- **Paris Buildings [17].** 6,412 still images of Paris landmarks, including 55 query images of 11 buildings with associated bounding box annotations.
- **INS 2013 [23].** A subset of 23,614 keyframes from TRECVID Instance Search (INS) dataset containing only those keyframes that are relevant for at least one of the queries of INS 2013.

### 4.2. Experimental Setup

We use both the VGG16 [22] and ZF [27] architectures of Faster R-CNN to extract image and region features. In both cases, we use the last convolutional layer (*conv5* and

*conv5\_3* for ZF and VGG16, respectively) to build the image descriptors introduced in Section 3, which are of dimension 256 and 512 for the ZF and VGG16 architectures, respectively. Region-wise features are pooled from the RoI pooling layer of Faster R-CNN. Images are re-scaled such that their shortest side is 600 pixels. All experiments were run in an Nvidia Titan X GPU.

### 4.3. Off-the-shelf Faster R-CNN features

In this section, we assess the performance of using off-the-shelf features from the Faster R-CNN network for instance retrieval.

First, we compare the sum and max pooling strategies of image- and region-wise descriptors. Table 1 summarizes the results. According to our experiments sumpooling is significantly superior to maxpooling for the filtering stage. Such behaviour is consistent with other works in the literature [3, 11]. Sumpooling is, however, consistently outperformed by maxpooling when reranking using region-wise features for all three datasets. Specifically for the Oxford and Paris datasets, we find the spatial reranking with maxpooling to be beneficial after filtering (gain of 0.10 and 0.03 mAP points for Oxford and Paris, respectively). However, the spatial reranking (either with max or sum pooling) has little or no effect for the INS13 dataset. To further interpret these results, we qualitatively evaluate the two pooling strategies. Figure 3 shows examples of top rankings for INS13 queries, spatially reranked with region-wise max and sum pooled descriptors. These examples indicate that, although mAP is similar, the object locations obtained with maxpooling are more accurate. According to this analysis, we set IPA-sum descriptors for the filtering stage and RPA-max descriptors for the spatial reranking in all the upcoming experiments of this paper.

Table 2 shows the performance of different Faster R-CNN architectures (ZF and VGG16) trained on two datasets (Pascal VOC and COCO [13]), including experiments with query expansion with the  $M = 5$  top retrieved images as well. As expected, features pooled from the deeper VGG16 network perform better in most cases, which is consistent with previous works in the literature showing that features from deeper networks reach better performance. Query expansion applied after the spatial reranking achieves significant gains for all tested datasets. Such behaviour was expected in particular with Oxford and Paris datasets, for which the spatial reranking already provided a significant gain. Interestingly, query expansion is also most beneficial after spatial reranking for the INS13 dataset, which suggests that, although in this case the spatial reranking does not provide any gain in mAP, the images that fall on the very top of the ranking are more useful to expand the query than the ones in the top of the first ranking.



Figure 3. Examples of top 4 rankings and object locations obtained for queries 9098: *a P (parking automat) sign* and 9076: *this monochrome bust of Queen Victoria* from the INS 2013 dataset (query images surrounded in blue). Comparison between the rankings generated using RPA-sum (top) and RPA-max (bottom), after the filtering stage with IPA-sum. Regressed bounding box coordinates have been disabled for visualization.

Table 1. Mean Average Precision (mAP) comparison between sum and max pooling strategies for both filtering and reranking stages using *conv5* features from the ZF Faster R-CNN model.

Filtering	Reranking	Oxford 5k	Paris 6k	INS 13
IPA-sum	None	0.505	0.612	<b>0.215</b>
	RPA-sum	0.501	0.621	0.196
	RPA-max	<b>0.602</b>	<b>0.641</b>	0.206
IPA-max	None	0.478	0.540	0.131
	RPA-sum	0.508	0.565	0.135
	RPA-max	0.559	0.561	0.138

Table 2. mAP of pre-trained Faster R-CNN models with ZF and VGG16 architectures. (P) and (C) denote whether the network was trained with Pascal VOC or Microsoft COCO images, respectively. In all cases, IPA-sum descriptors are used for the filtering stage. The CA-SR column specifies whether Class-Agnostic Spatial Reranking with RPA-max is applied to the top  $N = 100$  elements of the ranking. When indicated, QE is applied with  $M = 5$ .

Net	CA-SR	QE	Oxford 5k	Paris 6k	INS 13
ZF (P)	No	No	0.505	0.612	0.215
		Yes	0.515	0.671	0.246
	Yes	No	0.602	0.640	0.206
		Yes	0.622	0.707	<b>0.261</b>
VGG16 (P)	No	No	0.588	0.657	0.172
		Yes	0.614	0.706	0.201
	Yes	No	0.641	0.683	0.171
		Yes	<b>0.679</b>	0.729	0.242
VGG16 (C)	No	No	0.588	0.656	0.216
		Yes	0.600	0.695	0.250
	Yes	No	0.573	0.663	0.192
		Yes	0.647	<b>0.732</b>	0.241

#### 4.4. Fine-tuning Faster R-CNN

In this section, we assess the impact in retrieval performance of fine-tuning a pretrained network with the query objects to be retrieved. We choose to fine-tune the VGG16 Faster R-CNN model, pretrained with the objects of the Mi-

crosoft COCO dataset.

In the case of Oxford and Paris, we modify the output layers in the network to return 12 class probabilities (11 buildings in the dataset, plus an extra class for the background), and their corresponding regressed bounding box coordinates. We use the 5 images provided for each one of the buildings and their bounding box locations as training data. Additionally, we augment the training set by performing a horizontal flip on the training images ( $11 * 5 * 2 = 110$  training images in total). For INS 13, we have 30 different query instances, with 4 images each, giving raise to  $30 * 4 * 2 = 240$  training examples. The number of output classes for INS 13 is 31 (30 queries plus the background class).

The original Faster R-CNN training parameters described in [19] are kept for fine-tuning, except for the number of iterations, which we decreased to 5.000 considering our small number of training samples. We use the approximate joint training strategy introduced in [20], which trains the RPN and classifier branches at the same time, using the multi-task loss defined in [19]. This way, we train a separate network for each one of the tested datasets, using the two different fine-tuning modalities described in Section 3.2. Fine-tuning was performed on a Nvidia Titan X GPU and took around 30 and 45 minutes for finetuning strategies #1 and #2, respectively.

We first take the networks fine-tuned with strategy #1 and run the retrieval pipeline from scratch. Table 3 shows the obtained results (ft#1 columns). Results of the filtering and CA-SR stages are the same as those obtained with the original Faster R-CNN model, which is because the weights for the convolutional layers were not modified during fine-tuning. Results indicate that, although mAP is not always improved after CS-SR (e.g. from 0.588 to 0.543 for Oxford 5k), it is significantly better than CA-SR for Oxford and Paris when followed with query expansion. In case of the INS 13 dataset, we do not find significant improvements when using CS-SR, which suggests that only fine-tuning fully connected layers might not be sufficient to effectively

detect the challenging query objects in this dataset.

The second experiment in this section involves fine-tuning a higher number of layers in the Faster R-CNN architecture (Fine-tuning Strategy #2). Using this modality, the weights in the last convolutional layer are modified. Figure 4 shows the difference in the activations in conv5\_3 after fine-tuning it for the query instances in each dataset. These visualizations indicate that, after fine-tuning, more neurons in the convolutional layer positively react to the visual patterns that are present in the query objects of the dataset.

We then use the fine-tuned networks of the Fine-tuning Strategy #2 for each one of the datasets to extract image- and region-wise descriptors to perform instance search. Table 3 presents the results (ft#2 columns). As expected, fine-tuned features significantly outperform raw Faster R-CNN features for all datasets (mAP is  $\sim 20\%$  higher for Oxford and Paris, and 8% higher for INS 13). Results indicate that, for Oxford and Paris datasets, the gain of CA-SR + QE is higher with raw features (10% and 11% mAP increase for Oxford and Paris, respectively) than with fine-tuned ones (8% and 3% mAP increase, respectively). This suggests that fine-tuned features are already discriminant enough to correctly retrieve the objects in these two datasets. However, results for the INS 13 dataset show that CA-SR + QE is most beneficial when using fine-tuned features (11% and 41% mAP increase for raw and fine-tuned features, respectively). This difference between the performance for Oxford/Paris and INS13 suggests that queries for the latter are more challenging and therefore benefit from fine-tuned features and spatial reranking the most. A similar behaviour is observed for CS-SR which, for Oxford and Paris, is most beneficial when applied to a ranking obtained with raw features. For INS 13, however, the gain is greater when using fine-tuned features. Overall, the performance of reranking + query expansion is higher for CS-SR than CA-SR. Figure 1 shows examples of rankings for queries of the three different datasets after applying CS-SR. For visualization, we disable the regressed bounding box coordinates predicted by Faster R-CNN and choose to display those that are directly returned by the RPN. We find that the locations returned by the regression layer are inaccurate in most cases, which we hypothesize is caused by the lack of training data.

Finally, in Figure 5 we qualitatively evaluate the object detections after CS-SR using the fine-tuned strategies #1 and #2. The comparison reveals that locations obtained with the latter are more accurate and tight to the objects. The Fine-tuning Strategy #2 allows the RPN layers to adapt to the query objects, which causes the network to produce object proposals that are more suitable for the objects in the test datasets.

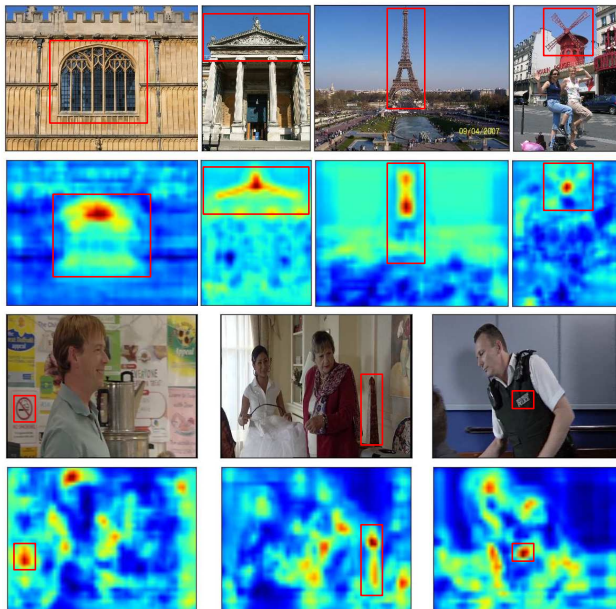


Figure 4. Difference between conv5\_3 features (sum pooled over feature maps) extracted from the original Faster R-CNN model pretrained with MS COCO with conv5\_3 features from the same model fine-tuned for INS13 (bottom), Oxford and Paris (top) queries.

Table 3. Comparison between Fine-tuning strategies #1 (ft#1) and #2 (ft#2) on the three datasets. Spatial Reranking (R) is applied to the  $N = 100$  top elements of the ranking. QE is performed with  $M = 5$ .

R	QE	Oxford 5k		Paris 6k		INS 13	
		ft#1	ft#2	ft#1	ft#2	ft#1	ft#2
No	No	0.588	0.710	0.656	0.798	0.216	0.234
No	Yes	0.600	0.748	0.695	0.813	<b>0.250</b>	0.259
CA-SR	No	0.573	0.739	0.663	0.801	0.192	0.248
CA-SR	Yes	0.647	0.772	0.732	0.824	0.241	0.330
CS-SR	No	0.543	0.751	<b>0.793</b>	0.807	0.181	0.250
CS-SR	Yes	<b>0.678</b>	<b>0.786</b>	0.784	<b>0.842</b>	<b>0.250</b>	<b>0.339</b>

#### 4.5. Comparison with state-of-the-art

In this section, we compare our results with several instance search works in the literature. Table 4 shows the results of this comparison.

Our proposed pipeline using Faster R-CNN features shows competitive results with respect to the state of the art. However, other works [11, 25] achieve a very high performance without any reranking nor query expansion strategies using similar feature pooling strategies. We hypothesize that the difference in the CNN architecture (Faster R-CNN vs. VGG16), training data (Pascal VOC vs ImageNet) and input image size (600px wide vs. full resolution) between these works and ours might be the reasons of the gap in performance. Our proposed reranking strategy CA-SR fol-






Figure 5. Ranking examples after CS-SR with fine-tuned strategies #1 (left) and #2 (right).

Table 4. Comparison with CNN-based state-of-the-art works on instance retrieval.

	Oxford 5k	Paris 6k
Razavian <i>et al.</i> [18]	0.556	0.697
Tolias <i>et al.</i> [25]	0.668	<b>0.830</b>
Kalantidis <i>et al.</i> [11]	0.682	0.796
Babenko and Lempitsky [3]	0.657	-
Ours	0.588	0.656
Ours (ft#2)	<b>0.710</b>	0.798
Tolias <i>et al.</i> (+ R + QE) [25]	0.770	<b>0.877</b>
Kalantidis <i>et al.</i> (+ QE) [11]	0.722	0.855
Ours (+ CA-SR + QE)	0.647	0.732
Ours (ft#1) (+ CS-SR + QE)	0.678	0.784
Ours (ft#2) (+ CS-SR + QE)	<b>0.786</b>	0.842

lowed by query expansion is demonstrated to provide similar mAP gains compared to the one proposed in [25]. While CA-SR + QE gives us a gain in mAP of  $\sim 10\%$  both for Oxford and Paris (using raw Faster R-CNN features), Tolias *et al.* [25] use their reranking strategy to raise their mAP by 5 and 15% for the two datasets, respectively.

As expected, results obtained with fine-tuned features (ft#2) achieve very competitive results compared to those in the state of the art, which suggests that fine-tuning the network for the object queries is an effective solution when time is not a constraint.

## 5. Conclusion

This paper has presented different strategies to make use of CNN features from an object detection CNN. It provides a simple baseline that uses off-the-shelf Faster R-CNN features to describe both images and their sub-parts. We have shown that is possible to greatly improve the performance of an off-the-shelf based system, at the cost of fine tuning the CNN for the query images that include objects that one wants to retrieve.

## Acknowledgements

This work has been developed in the framework of the project BigGraph TEC2013-43935-R, funded by the Spanish Ministerio de Economía y Competitividad and the European Regional Development Fund (ERDF). The Image Processing Group at the UPC is a SGR14 Consolidated Research Group recognized and sponsored by the Catalan Government (Generalitat de Catalunya) through its AGAUR office. Amaia Salvador developed this work thanks to the NII International Internship Program 2015. We gratefully acknowledge the support of NVIDIA Corporation with the donation of the GeForce GTX Titan X used in this work.

## References

- [1] R. Arandjelović and A. Zisserman. Three things everyone should know to improve object retrieval. In *Computer Vision and Pattern Recognition (CVPR)*, pages 2911–2918, 2012. 1
- [2] P. Arbeláez, J. Pont-Tuset, J. Barron, F. Marques, and J. Malik. Multiscale combinatorial grouping. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 328–335, 2014. 2
- [3] A. Babenko and V. Lempitsky. Aggregating local deep features for image retrieval. In *International Conference on Computer Vision (ICCV)*, December 2015. 1, 2, 3, 4, 7
- [4] A. Babenko, A. Slesarev, A. Chigorin, and V. Lempitsky. Neural codes for image retrieval. In *Computer Vision—ECCV 2014*, pages 584–599. 2014. 2
- [5] O. Chum, J. Philbin, J. Sivic, M. Isard, and A. Zisserman. Total recall: Automatic query expansion with a generative feature model for object retrieval. In *International Conference on Computer Vision (ICCV)*, pages 1–8, 2007. 1
- [6] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88(2):303–338, June 2010. 2
- [7] R. Girshick. Fast r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1440–1448, 2015. 2

- [8] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014. [2](#)
- [9] K. He, X. Zhang, S. Ren, and J. Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 37(9):1904–1916, 2015. [2](#)
- [10] H. Jégou, M. Douze, and C. Schmid. Improving bag-of-features for large scale image search. *International Journal of Computer Vision*, 87(3):316–336, 2010. [1](#)
- [11] Y. Kalantidis, C. Mellina, and S. Osindero. Cross-dimensional weighting for aggregated deep convolutional features. *arXiv:1512.04065*, 2015. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#)
- [12] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. [1](#), [2](#)
- [13] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014*, pages 740–755. Springer, 2014. [4](#)
- [14] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3431–3440, 2015. [1](#)
- [15] T. Mei, Y. Rui, S. Li, and Q. Tian. Multimedia search reranking: A literature survey. *ACM Computing Surveys (CSUR)*, 46(3):38, 2014. [1](#)
- [16] J. Philbin, O. Chum, M. Isard, J. Sivic, and A. Zisserman. Object retrieval with large vocabularies and fast spatial matching. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1–8, 2007. [4](#)
- [17] J. Philbin, O. Chum, M. Isard, J. Sivic, and A. Zisserman. Lost in quantization: Improving particular object retrieval in large scale image databases. In *Computer Vision and Pattern Recognition (CVPR)*, 2008. [4](#)
- [18] A. S. Razavian, H. Azizpour, J. Sullivan, and S. Carlsson. CNN features off-the-shelf: an astounding baseline for recognition. In *Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2014. [1](#), [2](#), [7](#)
- [19] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems*, pages 91–99, 2015. [1](#), [2](#), [5](#)
- [20] S. Ren, K. He, R. B. Girshick, and J. Sun. Faster R-CNN: towards real-time object detection with region proposal networks. *CoRR*, abs/1506.01497, 2015. [5](#)
- [21] A. Sharif Razavian, J. Sullivan, A. Maki, and S. Carlsson. A baseline for visual instance retrieval with deep convolutional networks. In *International Conference on Learning Representations*. ICLR, 2015. [2](#)
- [22] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014. [1](#), [4](#)
- [23] A. F. Smeaton, P. Over, and W. Kraaij. Evaluation campaigns and trecvid. In *International Workshop on Multimedia Information Retrieval (MIR)*, 2006. [4](#)
- [24] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–9, 2015. [1](#)
- [25] G. Toliás, R. Sivic, and H. Jégou. Particular object retrieval with integral max-pooling of CNN activations. *ICLR*, 2016. [1](#), [2](#), [3](#), [6](#), [7](#)
- [26] J. R. Uijlings, K. E. van de Sande, T. Gevers, and A. W. Smeulders. Selective search for object recognition. *International journal of computer vision*, 104(2):154–171, 2013. [2](#)
- [27] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *Computer vision–ECCV 2014*, pages 818–833. Springer, 2014. [4](#)
- [28] W. Zhang and C.-W. Ngo. Topological spatial verification for instance search. *IEEE Transactions on Multimedia*, 17(8):1236–1247, 2015. [1](#)
- [29] Y. Zhang, Z. Jia, and T. Chen. Image retrieval with geometry-preserving visual phrases. In *Computer Vision and Pattern Recognition (CVPR)*, pages 809–816, 2011. [1](#)