Challenging deep image descriptors for retrieval in heterogeneous iconographic collections

Dimitri Gominski, Martyna Poreba, Valérie Gouet-Brunet, Liming Chen

► To cite this version:

Dimitri Gominski, Martyna Poreba, Valérie Gouet-Brunet, Liming Chen. Challenging deep image descriptors for retrieval in heterogeneous iconographic collections. Proceedings of the 1st workshop on Structuring and Understanding of Multimedia heritAge Contents (SUMAC'19 Workshop @ ACM Multimedia 2019), Oct 2019, Nice, France. 10.1145/3347317.3357246 . hal-02273236

HAL Id: hal-02273236

<https://hal.science/hal-02273236>

Submitted on 18 Sep 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Challenging Deep Image Descriptors for Retrieval in Heterogeneous Iconographic Collections

Dimitri Gominski

dimitri.gominski@ign.fr

Univ. Paris-Est, IGN-ENSG/LaSTIG

Saint-Mandé, France

École Centrale Lyon - LIRIS

Écully, France

Valérie Gouet-Brunet

Valerie.Gouet@ign.fr

Univ. Paris-Est, IGN-ENSG/LaSTIG

Saint-Mandé, France

Martyna Poreba

Martyna.Poreba@ign.fr

Univ. Paris-Est, IGN-ENSG/LaSTIG

Saint-Mandé, France

Liming Chen

liming.chen@ec-lyon.fr]

École Centrale Lyon - LIRIS

Écully, France

Figure 1: Examples from the Alegoria benchmark. Each row shows different images for the same scene landmark: a. Arc de Triomphe (class 20); b. Palace of Versailles (class 34); c. Gerland Stadium & Port of Lyon (class 17); d. Dock of Saint-Tropez (class 39).

Royalty-free images from gallica.bnf.fr - Bibliothèque nationale de France, and from IGN (France) - Photothèque Nationale.

ABSTRACT

This article proposes to study the behavior of recent and efficient state-of-the-art deep-learning based image descriptors for content-based image retrieval, facing a panel of complex variations appearing in heterogeneous image datasets, in particular in cultural collections that may involve multi-source, multi-date and multi-view

contents. For this purpose, we introduce a novel dataset, namely Alegoria dataset, consisting of 12,952 iconographic contents representing landscapes of the French territory, and encapsulating a large range of intra-class variations of appearance which were finely labelled. Six deep features (DELF, NetVLAD, GeM, MAC, RMAC, SPoC) and a hand-crafted local descriptor (ORB) are evaluated against these variations. Their performance are discussed, with the objective of providing the reader with research directions for improving image description techniques dedicated to complex heterogeneous datasets that are now increasingly present in topical applications targeting heritage valorization.

CCS CONCEPTS

• **Information systems** → **Retrieval effectiveness**; • **Computing methodologies** → **Image representations**.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SUMAC '19, October 21, 2019, Nice, France

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6910-7/19/10...\$15.00

<https://doi.org/10.1145/3347317.3357246>

KEYWORDS

CBIR, deep learning, image descriptors, evaluation

ACM Reference Format:

Dimitri Gominski, Martyna Poreba, Valérie Gouet-Brunet, and Liming Chen. 2019. Challenging Deep Image Descriptors for Retrieval in Heterogeneous Iconographic Collections. In *1st Workshop on Structuring and Understanding of Multimedia heritAge Contents (SUMAC '19)*, October 21, 2019, Nice, France. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3347317.3357246>

1 INTRODUCTION

Digital or digitized cultural collections of images offer the interesting challenge of managing a wide variety of images together, mixing photographs with hand-made content, old black and white images with color pictures, low-quality scans or blurry photographs with digital pictures, etc. These collections are usually hosted within various institutions (e.g. GLAM - Galleries, Libraries, Archives and Museums), then organized in silos, without any interconnection between themselves while parts of them may address the same content (e.g. old paintings and recent photographs of Notre-Dame). The annotations associated may be of variable quality, or application-driven, making the interconnection of the fragmented contents not always easy. However, a better overall organization of these funds would be profitable in many areas where the complementarity of the available resources improves the analysis, ranging from SSH to landscape ecology, including urban planning and multimedia. In addition, there is currently a significant interest in the massive digitization of these heritage collections, with a desire to make them accessible to as many people as possible for multiple uses, associated with relevant structuring and consultation tools.

In this context, content-based image retrieval (CBIR) offers a powerful tool for establishing connections between images, independently of their native organization in the collections. Because of the variety of such contents in terms of acquisition source, view-point and illumination changes, evolution of the content across time, the key-word characterizing these collections is *heterogeneity*. These constraints introduce difficulty in the determination of efficient and robust content-based descriptors, from pixel level (alterations due to the digitization process or the aging of photographic chemicals) to semantic level (is a place still considered the same place if every building has changed?).

In this work, our objective is to establish an extensive evaluation of the most recent content-based descriptors, relying on deep features, for heterogeneous data such as those available in cultural collections, when considering an image retrieval task. This problem is sometimes referred to as "cross-domain" retrieval as in [20] or [4], but we will refrain from using that term since it suggests well-defined domains with their specific characteristics. Instead, we propose to speak of "heterogeneous" content-based image retrieval.

The paper is organized as follows: in section 2, we revisit the characteristics of the main public image collections available and present the one we propose for experiments in this work, which is related to cultural heritage. Section 3 is dedicated to the presentation of recent deep learning based descriptors. They are experimented in section 4, where we discuss the impact of the photometric and geometric transformations available on the heterogeneous contents related to cultural collections, before concluding in the last section.

2 HETEROGENEOUS COLLECTIONS OF IMAGES

This section is dedicated to image datasets involving very heterogeneous contents, with discussions about their main characteristics and the way of exploiting the latter in the qualitative and quantitative evaluation of image analysis and indexing techniques. We focus on cultural collections which gather interesting properties that continue to challenge the most efficient state-of-the-art image descriptors. In section 2.1, we present the Alegoria benchmark, which is a new challenging image dataset characterized by several interesting intra-class types of variations, before explaining in section 2.2 how we model these variations in order to enable a sharp evaluation of state-of-the-art image descriptors. Section 2.3 revisits other classical benchmarks and positions the Alegoria one alongside them.

2.1 ALEGORIA dataset

To address the retrieval problem in heterogeneous collections, we propose a benchmark consisting of **12952** grey and color images of outdoor scenes. Each image (in JPEG format) has a maximum side length of 800 pixels. Street-view, oblique, vertical aerial images, sketches, or old postcards can be found, taken from different view-points, by different type of cameras at different periods and sometimes even under different weather conditions. These geographic iconographic contents describe the French territory at different times since the inter-war period to the present day. They contain multiple objects and cultural artifacts: buildings (also stadiums and train stations), churches and cathedrals, historical sites (e.g. palaces, the most important monuments of Paris), seashores, suburbs of large cities, countrysides, etc. Some example images are shown in Figure 1, each row represents the same geographical site.

To enable quantitative comparisons, a subset of **681** images of this database have been manually selected and labelled. There are **39** classes of at least 10 images, each one is associated with the same topics site, for example Eiffel Tower, Arc de Triomphe, Notre-Dame de Paris, Sacré-Coeur Basilica, Palace of Versailles, Palace of Chantilly, Nanterre, Saint-Tropez, Stadium Lyon Gerland, Perrache train station. This benchmark can be used for CBIR in several applications such as place recognition, image-based localization or semantic segmentation.

2.2 Annotation of the appearance variations

The Alegoria dataset is a good illustration of a highly multi-source, multi-date and multi-view dataset. This heterogeneity allows to highlight significant variations of appearance, such as landscape transformations (site development, seasonal changes in vegetation), perspective (significant change in angle of view) or quality (color, B&W or sepia old photos). In order to evaluate the impact of these variations on current approaches of image analysis and indexing, we annotated the Alegoria dataset by considering a set of the most common and important intra-class variations. A total of 10 variation types were taken into account, including the usual Scale, Illumination and Orientation changes, plus variations that are more specific to cultural heritage : Alterations (chemical degradations or damages on the photographic paper before digitization), Color domains (grayscale, sepia, etc.), Domain representation (picture, drawing,

painting), Time changes (impact of large time spans) ; and general indicators of difficulty like Clutter, Positionning (when the main object of interest is not central to the picture) and Undistinctiveness (when the object of interest is not clear even to the human eye). Only variations presenting a high level of difficulty were counted: obviously there is always a degree of scale variation between two images of the same class, but we counted it only if it clearly adds difficulty when we visually compare two images.

To quantify this *variation predominance*, we use the following annotation process. For each class, we define a reference image, carefully chosen as the image depicting the object in the most common way in the class. For example, the second image of row a. on Figure 1 is a good reference, because most of the images in this class are in grayscale, with orientations between horizontal and vertical, with overall low quality, etc. We then get the *variation predominance* score by comparing all images in the class to the reference image, and measuring the frequency of occurrence: 0 when no image is concerned, 1 when anecdotal (one or less image from the class concerned), 2 when multiple occurrences are present, 3 when the variation is predominant (more than a third of the class). This also allows the computation of an overall difficulty score, that gives hints about the most difficult classes associated with multiple severe types of variations.

2.3 Relations to other benchmarks

The standard benchmarks for evaluation of content-based image retrieval techniques dedicated to landscapes are historically Oxford5k [16], Paris6k [17], and to a lesser extent Holidays [14] and UKBench [10]. Recently, Radenović et al. [18] proposed a revised version of Oxford and Paris datasets ($\mathcal{ROxford}$ and $\mathcal{RParis}$), correcting mistakes in the annotation and proposing three protocols of evaluation with varying levels of difficulty. The main variations in these small datasets (55 queries in 11 classes) arise due to image capture conditions like different viewpoints, occlusions, weather and illumination. Google also proposed its own dataset, namely Google-Landmarks [15], which contains around 5 million images of more than 200000 different landmarks. But this dataset is for now mainly used for training descriptors.

By introducing the Alegoria dataset, we aim at proposing a complementary benchmark, designed for precise evaluation of robustness on a broader panel of appearance variations. These variations bring into play challenging conditions such as long-term acquisitions (multi-date contents) as well as multi-source contents (drawings, engraving, photographs, etc.) that are not widely represented in the other popular datasets and have the additional interest of bridging cultural heritage and geographical information domains. We also generalize the content to a larger panel of geographical landscapes, including urban contents and landmarks as well as more natural landscapes such as mountains and rivers. The cathedral Notre-Dame de Paris is a good example of this complementarity: this landmark can be found in both Alegoria and Paris datasets, the difference being in what is evaluated. On Paris dataset, we assess the absolute performance of the retrieval method, whereas on Alegoria we can assess how the method reacts to different types of variations, including variations due to multi-date and multi-source contents.

3 DEEP FEATURES

Recently, convolutional neural networks (CNNs), aided with GPUs [12], have been proven to be powerful feature extractors. Contrary to the hand-crafted methods where descriptors were carefully designed to maximize invariance and discriminability (e.g. SIFT, ORB, SURF...), deep learning offers a way of letting an optimization algorithm to determine how to get these characteristics. In the seminal work of Babenko et al. [3], features were simply obtained at the output of the fully connected layers in early networks such as Alexnet. Babenko and Lempitsky [2] proved that aggregating features from the last convolutional layers produced better global descriptors. But these techniques lacked the core advantage of deep learning : optimizing the network directly for the retrieval task.

Arandjelović et al. [1] resolved this problem by proposing an end-to-end learning of deep features. Also declined in local [29], or patch-based [28] versions, these works completed the toolbox of deep features that was now ready to replace hand-crafted methods.

The backbone of recent image retrieval methods using deep features relies on a CNN, applying a function f_i to the input image I (depending on what layer i is considered), and producing a tensor of activations $T_i : T_i = f_i(I)$. Through the training process, we aim at optimizing f_i so that this 3D tensor, with width W and height H depending on the dimensions of the input image, and depth D depending on i , contain discriminative information. However T_i is too big to be indexed and compared during the retrieval process, it is thus mandatory to design a method reducing the memory cost of describing the images. The following methods will be presented along this guideline of how T_i is handled, giving either a local or a global descriptor.

3.1 Local methods

Local methods use a set of carefully selected points on an image. This involves identifying points that maximize invariance, and describing small patches around these points to extract information. Early deep methods replaced parts of the historical hand-crafted pipeline by trainable pieces. Verdie et al. [27] designed a learnable keypoint detector maximizing repeatability under drastic illumination changes. Yi et al. [29] extended the architecture for full point detecting and describing. However these two methods aim at building a robust detector/descriptor, whereas we want to *fit* our features to the data.

Noh et al. [15] solved the issue with a pipeline producing a set of local descriptors in a single forward pass, and that can be trained directly on any dataset with image-level labels. In their method, T_i is seen as a dense grid of local descriptors, where each position in the activation map is a D -dimensional local feature, whose receptive field in the input image is known. Additionally, an on-top function assigns a relevance score to each feature, and a threshold is set to only select most meaningful features. The output is a set of N **DEL**F descriptors per image. See Figure 3a for a visual interpretation. This departs from the traditional detect-then-describe process by selecting points *after* describing them, but it is simple to train and has shown good results on standard benchmarks [18]. Note that D typically ranges from 512 to 2048 in the last layers of the CNN, hence the PCA reduction to $D = 40$

Figure 2: Deep features extraction

proposed by the authors. Optimization on relevant data is done uniquely with a classification loss.

Dusmanu et al. [6] expanded this work with a detect-and-describe approach where they enforce keypoint repeatability and descriptor robustness using a Structure-from-Motion dataset with corresponding points on different images.

To perform the database similarity measure, local features must be aggregated. This is usually done with the learning of a dictionary (as in the well-known bag-of-words [23]), image are then described with a sparse vector suited for the inverted index structure.

Some local methods aggregate local descriptors into a compact representation, like the VLAD descriptor [9]. Arandjelović et al. [1] mimic VLAD with a learnable pooling layer, giving **NetVLAD**. By replacing the hard assignment step with soft assignment to multiple clusters, they can train this layer. In this work, T_i is considered, as in **DELF**, as a block containing $W \times H$ D-dimensional descriptors. The recent work of Teichmann et al. [24] builds upon the same idea, they describe selected candidate regions in the query with VLAD, and then propose a regional version of the ASMK [25] (another aggregation method) to aggregate these descriptors into a global descriptor.

Finally, the also recent work of Siméoni et al. [22] proposes a new view, following the observation that shapes of objects of interest in the input image can be found in some channels of T_i .

They perform detection and description of interest points in T_i using a hand-crafted detector (MSER [13]), and then match images based on spatial verification. However, this method is not suitable for large-scale retrieval since it works on pairs of images.

3.2 Global methods

Global methods describe an image as a whole, embedding all important information in a single vector. This is conceptually closer to the classification task for which common architectures like Resnet [7] or VGG [21] were designed. Babenko et al. [3] indeed showed that simply taking intermediate features from a classification network and using them for retrieval yields good performance.

To handle varying sizes in images and allegedly get more invariance, a standard seems to have emerged in deep global descriptors : extracting information at one of the last layers with a pooling process giving one value per channel. Following our notation, here the tensor T_i is seen as a set of D activation maps that contain each a different type of highly semantic information. To get a global descriptor, a straightforward approach is thus to get the most meaningful value per channel. We can then compare two images simply with dot-product similarity of their descriptors. See figure 3b for an example with the MAC descriptor detailed below. Babenko and Lempitsky [2] propose to sum the activations per channel, establishing the **SPoC** descriptor. This is equivalent to an average pooling operation. Differently, Kalantidis et al. [11] tweak the SPoC descriptor with a spatial and channel weighting, while Tolias et al. [26] get better results by using the maximum value per channel (**MAC** descriptor). They also propose the regional version **RMAC**, by sampling windows at different scales and describing them separately. Radenović et al. [19] generalize the preceding approaches with a generalized mean pooling (**GeM**) including a learnable parameter.

Global methods allow efficient fine-tuning on relevant data, either with the triplet loss [8] or the contrastive loss [5].

4 PERFORMANCE EVALUATION

This section is dedicated to the evaluation of the most efficient and recent approaches of image description for image retrieval, revisited in section 3, on the benchmark Alegoria presented in section 2, with the ambition of highlighting their behavior according to several types of appearance variations.

We evaluate the performance of **DELF** and **NetVLAD** for deep local-based descriptors, and **GeM**, **MAC**, **RMAC** and **SPoC** for global-based descriptors. We also include the hand-crafted descriptor **ORB** for reference.

Since there is no dataset for fully training a deep feature network with heterogeneous data involving all the types of variations we consider, we evaluate methods in an off-the-shelf manner, with no fine-tuning. However, all these methods were trained on contents close to the Alegoria contents: DELF was trained on the Landmarks dataset, which is a large-scale noisy dataset of landmarks with some typical variations such as Scale, Orientation and Occlusion. NetVLAD, GeM, MAC, RMAC and SPoC were all trained using the code provided by Radenović et al. [18], on the retrieval SfM 120k dataset. This one also focused on specific objects, mostly landmarks, with also interior pictures and standard variations (mostly Scale,

Table 1: Correlation between variations and performance

	DELf	ORB	NetVLAD	GeM	MAC	RMAC	SPoC
Scale	-0.41	-0.36	-0.47	-0.45	-0.47	-0.50	-0.43
Illumination	-0.32	-0.39	-0.53	-0.48	-0.45	-0.49	-0.45
Orientation	-0.42	-0.37	-0.50	-0.46	-0.44	-0.49	-0.45
Color	-0.28	-0.14	-0.09	-0.58	-0.58	-0.52	-0.54
Representation domain	0.07	-0.23	-0.10	-0.22	-0.18	-0.15	-0.23
Occlusion	0.05	-0.26	-0.38	-0.33	-0.37	-0.34	-0.32
Positioning	-0.17	0.07	-0.13	-0.38	-0.41	-0.41	-0.33
Clutter	-0.13	0.02	0.11	-0.10	-0.06	0.00	-0.10
Undistinctiveness	-0.02	0.09	0.09	0.22	0.16	0.24	0.23
Alterations	-0.31	-0.07	-0.15	-0.23	-0.22	-0.19	-0.23
Time changes	0.05	0.23	0.35	0.19	0.13	0.13	0.22
Overall difficulty	-0.42	-0.31	-0.41	-0.62	-0.64	-0.62	-0.58

Table 2: Color experiment: influence on the intra-class color variation on several descriptors.

Global mAP	DELf	GeM	RMAC
Mixed color domains	0.402	0.277	0.294
Grayscale only	0.421	0.282	0.299

Orientation, Illumination). We use Resnet101 as the backbone architecture, giving 2048-dimensionnal descriptors, except for NetVLAD, for which we use the standard parameters of the original paper (64 clusters * 512 channels in T_i).

For fair comparison, we discard any post-processing step. Global methods are compared with simple dot-product. DELf (dimension 1024 as in the original paper without PCA) and ORB descriptors are matched one-to-one giving a number-of-inliers score assessing similarity, using a product-quantized index for efficient memory management.

To compare the image descriptors evaluated, we use the classical mAP score, computed per class. For each query q , the average precision (AP) is computed on the sublist of results from 1 to k , k being the index of the last positive result.

The mAP per class is obtained by averaging the AP over all query images from a single class.

4.1 Results

The reader can refer to table 3 for the full lists of mAPs computed per class and the associated evaluation of variation predominance. To give an indicator of the overall difficulty of each class, we summed the predominance score of all the variations in the last column. We also computed the correlation matrix (see table 1) between the results of each methods and the predominance scores, considering each column as a series of 39 observations. Lower values indicate negative correlation: this variation is highly correlated with a decrease in performance. This does not imply causality but gives insights on the correspondences between variations and the performance of the descriptor. Note that there are also positive values, notably for Time changes and Undistinctiveness, indicating that these variations are correlated with other factors that on the

contrary improve performance. This might be a result of a bias in the selection of the pictures: when selecting pictures on a long time range for example, we tend to reduce the actual difficulty.

The Overall difficulty correlation score gives us a sanity check : overall difficulty indeed has a consistent negative correlation with the performance of all methods.

4.2 Discussion

In this section, we discuss the results obtained according to several criteria:

4.2.1 Local vs global description. The local DELf descriptor yields the best results with a consistent and significant margin. Table 1 shows that it is particularly stronger than global methods against Occlusion and Representation domain changes. This highlights the well-known advantage of local methods: by focusing on a set of local areas, they avoid the semantic noise captured by global methods. They also avoid the usual centering bias of global methods, as shown with the better Positioning score. However, on classes consisting of aerial images (e.g.class 7) RMAC gets better results. We believe this is due to the training dataset of DELf: it does not contain much aerial images, and DELf enforces a selection of important keypoints with its learned attention mechanism. DELf thus fails to find enough discriminative points on this type of data, whereas RMAC captures information on a large part of the image, allowing better results. Figure 4a gives an example where DELf fails to find true correspondences between two images from class 8.

4.2.2 Pooling. NetVLAD, GeM, MAC and SPoC only differ in the way they pool the tensor T to get a single global descriptor. We note that GeM gets overall better results, this can be explained by the fact that it generalizes MAC and SPoC with a tunable parameter, getting the best of both methods. NetVLAD performs consistently worse than others.

4.2.3 Attention mechanisms. RMAC has overall better performance than other global descriptors. We confirm the observation from the original paper [26] that the Region Proposal mechanism (which is basically an attention mechanism) of RMAC allows it to outperform

its simpler version MAC, and we note that this is also true against other pooling methods.

See figure 4b for an example where RMAC gets better results because it focuses on the parts of the image considered to contain the object of interest, whereas GeM returns negative but visually similar images. As showed in the ablation studies in the original paper of DELF, its attention mechanism is also responsible for a performance boost, but we lack other deep local methods to highlight this fact on our data.

4.2.4 Types of variations. Table 1 shows that Scale, Illumination, Orientation and Color are consistently associated with worse results. This shows that the main problem of image retrieval, even with modern deep learning methods, is still about getting invariance against basic variations. To support this analysis, we propose to do a simple experiment: mapping all the images in the same color domain (grayscale) before performing the description and the matching. We compute the global mAP for DELF, GeM and RMAC for the original dataset and for the grayscale dataset; see table 2 for the results. This normalization step reduces intra-class variance, but also the discriminative power of the descriptors. The mild but noticeable improvement shows that the former prevails, and we argue that a careful fine-tuning [19] can maintain this discriminative power with reduced variance.

(a) DELF matching failure case on aerial images

(b) RMAC vs. GeM descriptors on difficult cases. The top row shows the first 5 results for RMAC by decreasing order of similarity, the bottom row for GeM. The queries are in a white box, correct retrieved images in a green box and incorrect in a red box.

Figure 4: Retrieval examples

5 CONCLUSION

We proposed a new benchmark for the evaluation of deep features on heterogeneous data, and discussed how most recent and efficient features react to the panel of variations encountered. Our

results show that there is still many difficult cases to be handled by image retrieval methods, we thus presented insights on how to gain robustness with attention mechanisms and intra-class variance reduction. We believe this evaluation is necessary to allow image retrieval knowledge to be applied to real-world situations, and encourage future research to include detailed robustness studies, and to carefully design deep learning architectures for robust feature extraction regarding these variations.

ACKNOWLEDGMENTS

This work is supported by ANR (French National Research Agency) and DGA (French Directorate General of Armaments) within the ALEGORIA project, under respective Grant ANR-17-CE38-0014-01 and DGA Agreement 2017-60-0048.

REFERENCES

- [1] Relja Arandjelović, Petr Gronat, Akihiko Torii, Tomas Pajdla, and Josef Sivic. 2017. NetVLAD: CNN architecture for weakly supervised place recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* XX (June 2017). <https://hal.inria.fr/hal-01557234>
- [2] Artem Babenko and Victor S. Lempitsky. 2015. Aggregating Local Deep Features for Image Retrieval. In *2015 IEEE International Conference on Computer Vision (ICCV)*. 1269–1277. DOI: <http://dx.doi.org/10.1109/ICCV.2015.150>
- [3] Artem Babenko, Anton Slesarev, Alexandr Chigorin, and Victor Lempitsky. 2014. Neural Codes for Image Retrieval. In *LNCS*, Vol. 8689. DOI: http://dx.doi.org/10.1007/978-3-319-10590-1_38
- [4] N. Bhowmik, Li Weng, V. Gouet-Brunet, and B. Sohelian. 2017. Cross-domain image localization by adaptive feature fusion. In *2017 Joint Urban Remote Sensing Event (JURSE)*. 1–4. DOI: <http://dx.doi.org/10.1109/JURSE.2017.7924572>
- [5] S. Chopra, R. Hadsell, and Y. LeCun. 2005. Learning a similarity metric discriminatively, with application to face verification. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, Vol. 1. 539–546 vol. 1. DOI: <http://dx.doi.org/10.1109/CVPR.2005.202>
- [6] M. Dusmanu, I. Rocco, T. Pajdla, M. Pollefeys, J. Sivic, A. Torii, and T. Sattler. 2019. D2-Net: A Trainable CNN for Joint Detection and Description of Local Features. In *IEEE Conference on Computer Vision and Pattern Recognition*. <https://dsnn.ml/publications/d2-net.html>
- [7] K. He, X. Zhang, S. Ren, and J. Sun. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 770–778. DOI: <http://dx.doi.org/10.1109/CVPR.2016.90>
- [8] Elad Hoffer and Nir Ailon. 2015. Deep Metric Learning Using Triplet Network. In *Similarity-Based Pattern Recognition (Lecture Notes in Computer Science)*, Aasa Feragen, Marcello Pelillo, and Marco Loog (Eds.). Springer International Publishing, 84–92.
- [9] Herve Jegou, Matthijs Douze, Cordelia Schmid, and Patrick Perez. 2010. Aggregating local descriptors into a compact image representation. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. IEEE, San Francisco, CA, USA, 3304–3311. DOI: <http://dx.doi.org/10.1109/CVPR.2010.5540039>
- [10] Hervé Jégou, Matthijs Douze, and Cordelia Schmid. 2008. *Hamming Embedding and Weak Geometry Consistency for Large Scale Image Search - extended version*. report. <https://hal.inria.fr/inria-00548651>
- [11] Yannis Kalantidis, Clayton Mellina, and Simon Osindero. 2016. Cross-Dimensional Weighting for Aggregated Deep Convolutional Features. In *Computer Vision – ECCV 2016 Workshops (Lecture Notes in Computer Science)*, Gang Hua and Hervé Jégou (Eds.). Springer International Publishing, 685–701.
- [12] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems 25*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (Eds.). Curran Associates, Inc., 1097–1105. <http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf>
- [13] J. Matas, O. Chum, M. Urban, and T. Pajdla. 2004. Robust wide-baseline stereo from maximally stable extremal regions. *Image and Vision Computing* 22, 10 (Sept. 2004), 761–767. DOI: <http://dx.doi.org/10.1016/j.imavis.2004.02.006>
- [14] D. Nister and H. Stewenius. 2006. Scalable Recognition with a Vocabulary Tree. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition - Volume 2 (CVPR'06)*, Vol. 2. IEEE, New York, NY, USA, 2161–2168. DOI: <http://dx.doi.org/10.1109/CVPR.2006.264>
- [15] H. Noh, A. Araujo, J. Sim, T. Weyand, and B. Han. 2017. Large-Scale Image Retrieval with Attentive Deep Local Features. In *2017 IEEE International Conference on Computer Vision (ICCV)*. 3476–3485. DOI: <http://dx.doi.org/10.1109/ICCV.2017.374>

Table 3: Results: retrieval accuracy against class variations

Class	#	Descriptor								Variations											
		DELf	ORB	NetVLAD	GeM	MAC	RMAC	SPoC	Avg	Sc.	Il.	Or.	Co.	Do.	Oc.	Po.	Cl.	Un.	Al.	Ti.	Overall
1	48	0.249	0.044	0.050	0.513	0.401	0.464	0.485	0.210	3	2	3	2	1	0	2	2	0	2	1	18
2	13	0.807	0.091	0.095	0.370	0.306	0.357	0.296	0.332	0	1	0	2	1	0	3	2	0	2	3	14
3	16	0.604	0.171	0.071	0.275	0.218	0.303	0.263	0.272	3	3	3	3	0	2	3	2	0	3	0	22
4	22	0.389	0.054	0.050	0.231	0.186	0.211	0.220	0.192	2	3	3	3	0	2	3	3	2	3	1	25
5	18	0.270	0.059	0.060	0.236	0.215	0.209	0.211	0.180	2	2	2	0	0	1	0	1	0	3	2	13
6	11	0.623	0.269	0.210	0.562	0.475	0.588	0.538	0.467	0	2	0	0	0	0	3	2	0	2	2	11
7	11	0.729	0.123	0.099	0.651	0.514	0.822	0.626	0.509	0	0	0	0	0	0	0	3	3	3	0	9
8	22	0.514	0.082	0.075	0.661	0.607	0.628	0.554	0.446	0	2	0	0	0	0	0	3	0	3	3	11
9	15	0.628	0.075	0.072	0.333	0.235	0.332	0.338	0.288	3	3	2	2	0	2	3	1	0	1	0	17
10	11	0.211	0.093	0.093	0.233	0.212	0.246	0.210	0.185	3	2	3	3	0	0	3	3	3	3	3	26
11	19	0.238	0.071	0.057	0.153	0.138	0.168	0.147	0.139	2	3	0	3	0	0	3	3	3	3	0	20
12	12	0.152	0.099	0.087	0.156	0.111	0.150	0.151	0.129	2	2	1	3	0	0	3	3	1	3	3	21
13	13	0.296	0.136	0.099	0.273	0.225	0.338	0.292	0.237	0	3	0	3	0	0	3	3	0	3	1	16
14	10	0.519	0.142	0.128	0.353	0.233	0.403	0.348	0.304	0	2	0	3	3	2	0	3	0	2	3	18
15	14	0.149	0.079	0.076	0.156	0.158	0.144	0.133	0.128	3	2	3	3	0	0	3	2	0	3	3	22
16	26	0.378	0.105	0.127	0.394	0.318	0.412	0.327	0.294	2	0	0	3	0	0	3	3	3	3	3	20
17	11	0.157	0.092	0.097	0.115	0.127	0.106	0.124	0.117	0	2	3	3	1	0	3	3	0	3	3	21
18	19	0.467	0.209	0.121	0.335	0.278	0.353	0.346	0.301	1	0	0	3	0	0	3	3	2	3	3	18
19	27	0.167	0.044	0.042	0.103	0.105	0.105	0.085	0.093	3	3	3	3	2	0	2	3	0	3	0	22
20	15	0.658	0.076	0.074	0.203	0.160	0.268	0.235	0.239	1	3	0	3	2	2	2	3	0	2	3	21
21	14	0.198	0.073	0.075	0.095	0.100	0.103	0.089	0.105	2	3	3	3	2	2	3	3	0	3	3	27
22	23	0.347	0.048	0.061	0.110	0.104	0.113	0.094	0.125	3	3	3	3	2	2	2	3	0	3	0	24
23	17	0.169	0.081	0.068	0.134	0.139	0.141	0.109	0.120	3	2	1	3	2	2	2	3	0	3	0	21
24	11	0.186	0.093	0.095	0.174	0.181	0.181	0.146	0.151	2	2	0	3	1	0	1	3	1	3	0	16
25	12	0.560	0.086	0.126	0.449	0.452	0.541	0.382	0.371	0	2	0	3	3	0	0	3	0	1	0	12
26	26	0.322	0.047	0.046	0.111	0.129	0.114	0.067	0.119	2	3	2	3	1	3	2	3	0	3	0	22
27	36	0.268	0.037	0.041	0.073	0.065	0.089	0.060	0.090	1	2	0	3	2	2	3	3	0	3	2	21
28	14	0.507	0.105	0.099	0.239	0.247	0.259	0.200	0.237	0	1	0	3	0	1	1	3	0	3	0	12
29	21	0.297	0.058	0.066	0.254	0.192	0.195	0.241	0.186	3	2	1	1	1	3	2	0	3	1	2	19
30	11	0.417	0.131	0.160	0.472	0.329	0.566	0.450	0.361	0	0	0	2	0	0	0	3	2	3	2	12
31	16	0.353	0.064	0.074	0.181	0.174	0.184	0.153	0.169	1	2	2	3	0	1	1	3	0	3	0	16
32	30	0.197	0.105	0.107	0.176	0.160	0.193	0.135	0.111	0	3	0	1	0	0	1	3	0	3	0	11
33	17	0.715	0.111	0.117	0.305	0.264	0.362	0.292	0.309	1	0	0	3	2	0	2	3	0	3	3	17
34	10	0.284	0.121	0.104	0.240	0.216	0.231	0.213	0.201	1	1	1	3	0	0	1	3	0	1	3	14
35	10	0.317	0.102	0.103	0.193	0.176	0.186	0.170	0.178	1	1	0	3	0	1	1	3	0	3	0	13
36	18	0.309	0.070	0.065	0.197	0.158	0.194	0.186	0.169	2	2	2	2	0	0	3	3	0	3	0	17
37	18	0.226	0.063	0.078	0.237	0.171	0.225	0.233	0.176	1	2	2	3	0	0	2	3	0	3	3	19
38	13	0.446	0.110	0.155	0.439	0.309	0.439	0.463	0.337	2	2	0	3	0	0	2	3	0	3	3	18
39	11	0.294	0.095	0.097	0.209	0.201	0.216	0.200	0.187	2	2	1	3	1	0	3	3	0	2	1	18

Best seen in color

Sc.: Scale, Il.: Illumination, Or.: Orientation, Co.: Color, Do.: Domain representation, Oc.: Occlusion, Po.: Positioning, Cl.: Clutter, Un.: Undistinctiveness, Al.: Alterations, Ti.: Time changes, Overall: Overall difficulty. For each class, the 3 top mAPs are in green, ranging from dark (best result) to light green (3rd result). Column Avg. indicates the average performance of all methods. Column Overall gives an overall score of difficulty (sum of the variation scores).

- [16] J. Philbin, O. Chum, M. Isard, J. Sivic, and A. Zisserman. 2007. Object retrieval with large vocabularies and fast spatial matching. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*. 1–8. DOI : <http://dx.doi.org/10.1109/CVPR.2007.383172>
- [17] J. Philbin, O. Chum, M. Isard, J. Sivic, and A. Zisserman. 2008. Lost in quantization: Improving particular object retrieval in large scale image databases. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*. 1–8. DOI : <http://dx.doi.org/10.1109/CVPR.2008.4587635>

- [18] Filip Radenović, Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondřej Chum. 2018. Revisiting Oxford and Paris: Large-Scale Image Retrieval Benchmarking. DOI : <http://dx.doi.org/10.1109/CVPR.2018.00598>
- [19] F. Radenović, G. Tolias, and O. Chum. 2019. Fine-Tuning CNN Image Retrieval with No Human Annotation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41, 7 (July 2019), 1655–1668. DOI : <http://dx.doi.org/10.1109/TPAMI.2018.2846566>
- [20] Abhinav Shrivastava, Tomasz Malisiewicz, Abhinav Gupta, and Alexei A. Efros. 2011. Data-driven Visual Similarity for Cross-domain Image Matching. *ACM*

- Transaction of Graphics (TOG) (Proceedings of ACM SIGGRAPH ASIA)* 30, 6 (2011).
- [21] Karen Simonyan and Andrew Zisserman. 2014. Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv:1409.1556 [cs]* (Sept. 2014). <http://arxiv.org/abs/1409.1556> arXiv: 1409.1556.
 - [22] Oriane Siméoni, Yannis Avrithis, and Ondrej Chum. 2019. Local Features and Visual Words Emerge in Activations. *arXiv:1905.06358 [cs]* (May 2019). <http://arxiv.org/abs/1905.06358> arXiv: 1905.06358.
 - [23] J. Sivic and A. Zisserman. 2003. Video Google: A Text Retrieval Approach to Object Matching in Videos. In *Proceedings of the International Conference on Computer Vision*, Vol. 2. 1470–1477. <http://www.robots.ox.ac.uk/vgg>
 - [24] Marvin Teichmann, Andre Araujo, Menglong Zhu, and Jack Sim. 2018. Detect-to-Retrieve: Efficient Regional Aggregation for Image Search. *arXiv:1812.01584 [cs]* (Dec. 2018). <http://arxiv.org/abs/1812.01584> arXiv: 1812.01584.
 - [25] Giorgos Tolias, Yannis Avrithis, and Hervé Jégou. 2016. Image Search with Selective Match Kernels: Aggregation Across Single and Multiple Images. *International Journal of Computer Vision* 116, 3 (Feb. 2016), 247–261. DOI : <http://dx.doi.org/10.1007/s11263-015-0810-4>
 - [26] Giorgos Tolias, Ronan Sicre, and Hervé Jégou. 2016. Particular Object Retrieval With Integral Max-Pooling of CNN Activations. In *ICL 2016 - International Conference on Learning Representations (International Conference on Learning Representations)*. San Juan, Puerto Rico, 1–12. <https://hal.inria.fr/hal-01842218>
 - [27] Y. Verdie, Kwang Moo Yi, P. Fua, and V. Lepetit. 2015. TILDE: A Temporally Invariant Learned DETector. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 5279–5288. DOI : <http://dx.doi.org/10.1109/CVPR.2015.7299165>
 - [28] Xufeng Han, T. Leung, Y. Jia, R. Sukthankar, and A. C. Berg. 2015. MatchNet: Unifying feature and metric learning for patch-based matching. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 3279–3286. DOI : <http://dx.doi.org/10.1109/CVPR.2015.7298948>
 - [29] Kwang Moo Yi, Eduard Trulls, Vincent Lepetit, and Pascal Fua. 2016. LIFT: Learned Invariant Feature Transform. In *Computer Vision – ECCV 2016 (Lecture Notes in Computer Science)*, Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling (Eds.). Springer International Publishing, 467–483.