

Efficient On-the-fly Category Retrieval Using ConvNets and GPUs

Ken Chatfield^(✉), Karen Simonyan, and Andrew Zisserman

Department of Engineering Science, University of Oxford, Oxford, UK
{ken,karen,az}@robots.ox.ac.uk

Abstract. We investigate the gains in precision and speed, that can be obtained by using Convolutional Networks (ConvNets) for on-the-fly retrieval – where classifiers are learnt at run time for a textual query from downloaded images, and used to rank large image or video datasets.

We make three contributions: (i) we present an evaluation of state-of-the-art image representations for object category retrieval over standard benchmark datasets containing 1M+ images; (ii) we show that ConvNets can be used to obtain features which are incredibly performant, and yet much lower dimensional than previous state-of-the-art image representations, and that their dimensionality can be reduced further without loss in performance by compression using product quantization or binarization. Consequently, features with the state-of-the-art performance on large-scale datasets of millions of images can fit in the memory of even a commodity GPU card; (iii) we show that an SVM classifier can be learnt within a ConvNet framework on a GPU *in parallel* with downloading the new training images, allowing for a continuous refinement of the model as more images become available, and simultaneous training and ranking. The outcome is an on-the-fly system that significantly outperforms its predecessors in terms of: precision of retrieval, memory requirements, and speed, facilitating accurate on-the-fly learning and ranking in under a second on a single GPU.

1 Introduction

On-the-fly learning offers a way to overcome the ‘closed world’ problem in computer vision, where object category recognition systems are restricted to only those pre-defined classes that occur in the carefully curated datasets available for training – for example ImageNet [1] for object categories or UCF-101 [2] for human actions in videos. What is more, it offers the tantalising prospect of developing large-scale general purpose object category retrieval systems which can operate over millions of images in a few seconds, as is possible in the specific instance retrieval systems [3–7] which have reached the point of commercialisation in products such as Google Goggles, Kooaba and Amazon’s SnapTell.

Current on-the-fly systems typically proceed in three stages [8–11]: first, training data for the user query are compiled, commonly by bootstrapping the process via text-to-image search using *e.g.* Google Image Search as a source of

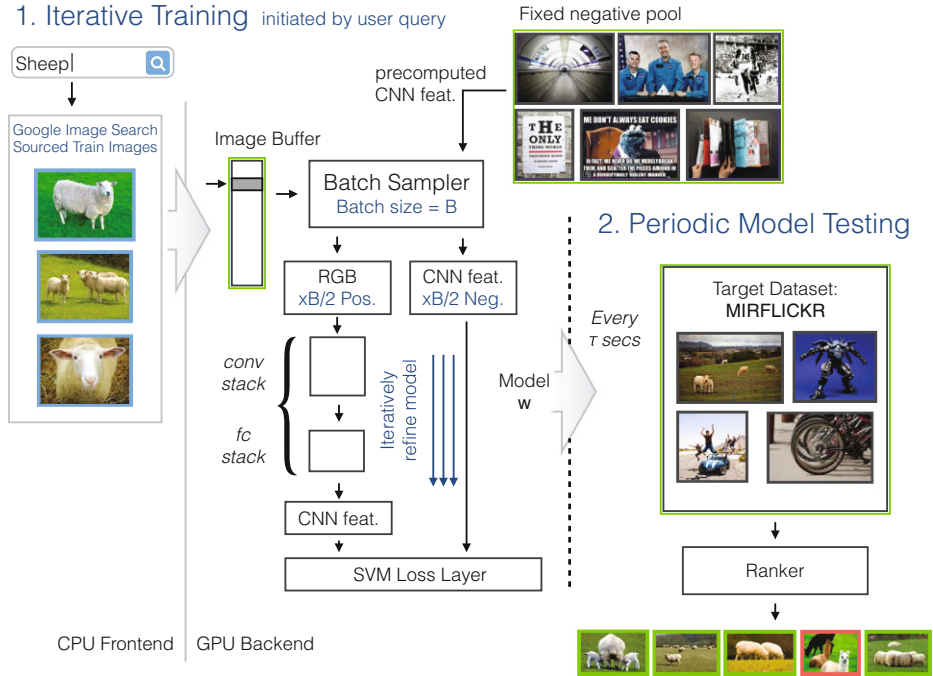


Fig. 1. Architecture of our on-the-fly object category retrieval system. The entire framework aside from the image downloader is resident on the GPU, with data stored in GPU memory outlined in green. Its operation is split into: (i) iterative training, as initiated by a user text query and (ii) periodic model testing to obtain a ranking over the target dataset (refer to text for further details) (Color figure online).

training images; second, a classifier or ranker is learnt for that category; third, all images/videos in a dataset are ranked in order to retrieve those containing the category. The aim is for these stages to happen on-line in a matter of seconds, rather than hours.

Previous methods for on-the-fly learning have been limited by the retrieval-performance/memory/speed trade off. In particular, very high-dimensional feature vectors were required for state-of-the-art classification performance [12–14], but this incurred both a severe memory penalty (as features for the dataset need to be in memory for fast retrieval) and also a severe speed penalty (as computing a scalar product for high-dimensional features is costly) both in training and ranking. Despite the excellent progress in compression methods for nearest neighbour search by using product quantization [15] or binary encoding [16, 17], compromises still had to be made.

In this paper we show that in the context of on-the-fly category retrieval, Convolutional Networks (ConvNets) [18] with GPU training [19] can significantly improve on all three of: retrieval precision, memory requirements, and ranking speed. The whole pipeline, from computing the training image features

and learning the model to scoring and ranking the dataset images is implemented on the GPU and runs in a highly-parallel, online manner. We thus demonstrate a system that is able to go from a cold-query to results in a matter of second(s) on a dataset of million(s) of images. The architecture of our proposed system, from input of text query to display of ranked results, is summarized in Fig. 1 (refer to Sect. 4 for details).

In terms of retrieval performance, we build on the recent research that shows that deep ConvNet features significantly outperform shallow features, such as Fisher Vectors [12,13], on the image classification task [19–21]. However, our contributions go further than simply using ConvNet features in an on-the-fly architecture: we take the full advantage of the GPU computation for all retrieval stages, in parallel with downloading the new training images on the CPU. This novel GPU-based architecture allows a time budget to be set, so that an SVM, trained on the available images within the time limit, can be used to (re-)rank the dataset images at any stage of the process (for instance, every 0.5s). This architecture is in strong contrast to the standard on-the-fly architectures [8], where SVM training only begins once all training images have been downloaded and processed, and ranking follows after that.

We first perform a comprehensive evaluation of the performance of ConvNet-based image features for category-based image retrieval (Sect. 2). We start with a standard object category recognition benchmark (PASCAL VOC 2007 [22]), and then add a large number of distractor images to take the dataset size to 1M+ images (the datasets are described in Sect. 2). We assess retrieval performance over these two datasets (VOC and VOC + distractors) under variation in the training data – either using VOC training images (i.e. a curated dataset) or using images from Google Image search (i.e. the type of images, possibly with label noise, that will be available in the real-world on-the-fly system).

With our goal being ranking of millions of images on a conventional GPU-equipped PC, we then investigate, in Sect. 3, how retrieval performance is affected by using low-dimensional features (still originating from a ConvNet). Low-dimensional features (*e.g.* hundreds of components rather than thousands) have two advantages: they use less memory, and scalar products are faster, both in training and ranking. We cover a spectrum of methods for achieving a low-dimensional descriptor, namely: (i) reducing the dimensionality of the last ConvNet layer; (ii) product quantization of the ConvNet features and (iii) binarization of the ConvNet features. It is shown that a combination of a low-dimensional final ConvNet feature layer with product quantization produces features that are both highly-compact and incredibly performant.

Finally, based on these investigations, we propose a GPU architecture for on-the-fly object category retrieval in Sect. 4, highly scalable, capable of adapting to varying query complexity and all running on a single commodity GPU. An extended version of this paper is available on arXiv¹.

¹ <http://arxiv.org/abs/1407.4764/>.

2 Evaluating Large-Scale Object Category Retrieval

This section describes the evaluation protocol used to assess the performance of the image representations $\phi(I)$ described in Sect. 3 and of the on-the-fly training architecture introduced in Sect. 4. We begin by describing the datasets used for evaluation, and then describe the three different scenarios in which these datasets are used, with each subsequent scenario moving closer to modelling the conditions experienced by a real-world large-scale object category retrieval system.

One difficulty of evaluating a large-scale object category retrieval system is the lack of large-scale datasets with sufficient annotation to assess retrieval performance fully, in particular to measure recall. The PASCAL VOC dataset [22] provides full annotation for a set of twenty common object classes, facilitating evaluation using common ranking performance measures such as mean average precision (mAP), but is much too small ($\sim 10k$ images) to evaluate the performance of a real-world system. Conversely, the ILSVRC dataset [1], while being much larger ($\sim 1M+$ images), does not have complete annotation of *all* object categories in each image. Therefore, ranking performance (*e.g.* recall or mAP) cannot be measured without further annotation, and only object category *classification* metrics (such as top-N classification error per image), which do not accurately reflect the performance of an object category *retrieval* scenario, can be used. Additionally, in this work we use the ImageNet ILSVRC-2012 dataset to pre-train the ConvNet, so can not also use that for assessing performance.

As a result, for evaluation in this paper, we use a custom combination of datasets, carefully tailored to be representative of the data that could be expected in a typical collection of web-based consumer photographs:

PASCAL VOC 2007 [22] is used as our base dataset, with assessment over seventeen of its twenty classes (‘people’, ‘cats’ and ‘birds’ are excluded for reasons explained below). We use the provided train, validation and test splits.

MIRFLICKR-1M [23,24] is used to augment the data from the PASCAL VOC 2007 test set in our later experiments, and comprises 1M unannotated images (aside from quite noisy image tags). The dataset represents a snapshot of images taken by popularity from the image sharing site Flickr, and thus is more representative of typical web-based consumer photography than ImageNet, which although also sourced from Flickr was collected through queries for often very specific terms from WordNet. In addition, MIRFLICKR-1M has been confirmed to contain many images of the twenty PASCAL VOC classes.

2.1 Evaluation Protocol

A linear SVM is trained for all classes, and used to rank all images in the target dataset. For the object category retrieval setting the ‘goodness’ of the first few pages of retrieved results is critical, as the larger the proportion of true positives

for a given object category at the top of a ranked list, the better the perceived performance. We therefore evaluate using precision @ K , where $K = 100$.

Adopting such an evaluation protocol also has the advantage that we are able to use the 1M images from the MIRFLICKR-1M dataset despite the fact that full annotations are not provided. Since we only need to consider the top K of the ranked list for each class during evaluation, we can take a ‘lazy’ approach to annotating the MIRFLICKR-1M dataset, annotating class instances only as far down the ranked list as necessary to generate a complete annotation for the top- K results (for more details of this procedure, refer to scenario 2 below). This avoids having to generate a full set of annotation for all 1M images.

2.2 Experimental Scenarios

Scenario 1: PASCAL VOC. We train models for seventeen of the twenty VOC object classes (excluding ‘people’, ‘cats’ and ‘birds’) using both the training and validation sets. Following this, a ranked list for each class is generated using images from the test set and precision @ K evaluated.

Scenario 2: Large-Scale Retrieval. Training is undertaken in the same manner as scenario 1, but during testing images are added from the MIRFLICKR-1M dataset. There are two sub-scenarios:

Scenario 2a – we test using images from the PASCAL VOC test set (as in scenario 1) with the addition of the entirety of the MIRFLICKR-1M dataset. For each class, we remove all (lazily annotated) positive class occurrences in the ranked list which are retrieved from MIRFLICKR-1M, as the purpose of this scenario is to test how our features perform when attempting to retrieve a small, known number of class occurrences from a very large number of non-class ‘distractor’ images.²

Scenario 2b – this time we exclude all images from the PASCAL VOC dataset, and instead evaluate precision @ K solely over the MIRFLICKR-1M dataset. The purpose of this scenario is to test how our features perform over a real-world dataset with unknown statistics. In practice, it is an easier scenario than scenario 2a, since the MIRFLICKR-1M dataset contains many instances of all of the PASCAL VOC classes.

Scenario 3: Google Training. Testing is the same as in scenario 2b, but instead of using PASCAL data for training, a query is issued to Google Image search for each of the PASCAL VOC classes, and the top $N \sim 250$ images are used in each case as training data. This scenario assesses the tolerance to training on images that differ from the VOC and MIRFLICKR-1M test images:

² The prevalence of the PASCAL VOC classes ‘people’, ‘cats’ and ‘birds’ in the MIRFLICKR-1M data explains why we exclude them, as restricting the annotation of these classes to reasonable levels proved to be impossible.

the Google images may be noisy and typically contain the object in the centre. It also mirrors most closely a real-world on-the-fly object category retrieval setting, as the queries in practice do not need to be limited to the PASCAL VOC classes. There are again two sub-scenarios, with different data used for the negative training samples in each case:

Scenario 3a – the images downloaded from Google Image Search for all other classes, except for the current class, are used as negative training data (this mirrors the PASCAL VOC setup).

Scenario 3b – a fixed pool of $\sim 16,000$ negative training images is used. These training images are sourced from the web by issuing queries for a set of fixed ‘negative’ query terms³ to both Google and Bing image search, and attempting to download the first $\sim 1,000$ results in each case. This same pool of negative training data is also used in Sect. 4.

3 Retrieval Performance over Image Representations

In this section, we perform an evaluation of recent state-of-the-art image representations for the object category retrieval scenarios described in Sect. 2.2.

ConvNet-based features, which form the basis of our on-the-fly system described in Sect. 4, have been shown to perform excellently on standard image classification benchmarks such as PASCAL VOC and ImageNet ILSVRC [20, 21, 25, 26]. We therefore focus our evaluation on these features, employing 2048-dimensional ‘CNN M 2048’ image features of [21] as the baseline. We compare them to a more traditional shallow feature encoding in the form of the Improved Fisher Vector (IFV) [13]. Implementation details for ConvNets and IFV are given in Sect. 3.2. We explore the effects of reducing the dimensionality of our features on their retrieval performance using the following methods:

Lower-dimensional ConvNet output layer – One way of reducing the dimensionality of ConvNet features consists in retraining the network so that the last fully-connected (feature) layer has a lower dimensionality. Following [21], we consider the ‘CNN M 128’ network configuration with a 128-dimensional feature layer. Using such network in place of the baseline ‘CNN M 2048’ can be seen as discriminative dimensionality reduction by a factor of 16.

Product quantization (PQ) has been widely used as a compression method for image features [15, 27], and works by splitting the original feature into Q -dimensional sub-blocks, each of which is encoded using a separate vocabulary of cluster centres pre-learned from a training set. Here we explore compression using $Q = 4, 8$ -dimensional sub-blocks.

³ Miscellanea, random selection, photo random selection, random objects, random things, nothing in particular, photos of stuff, random photos, random stuff, things.

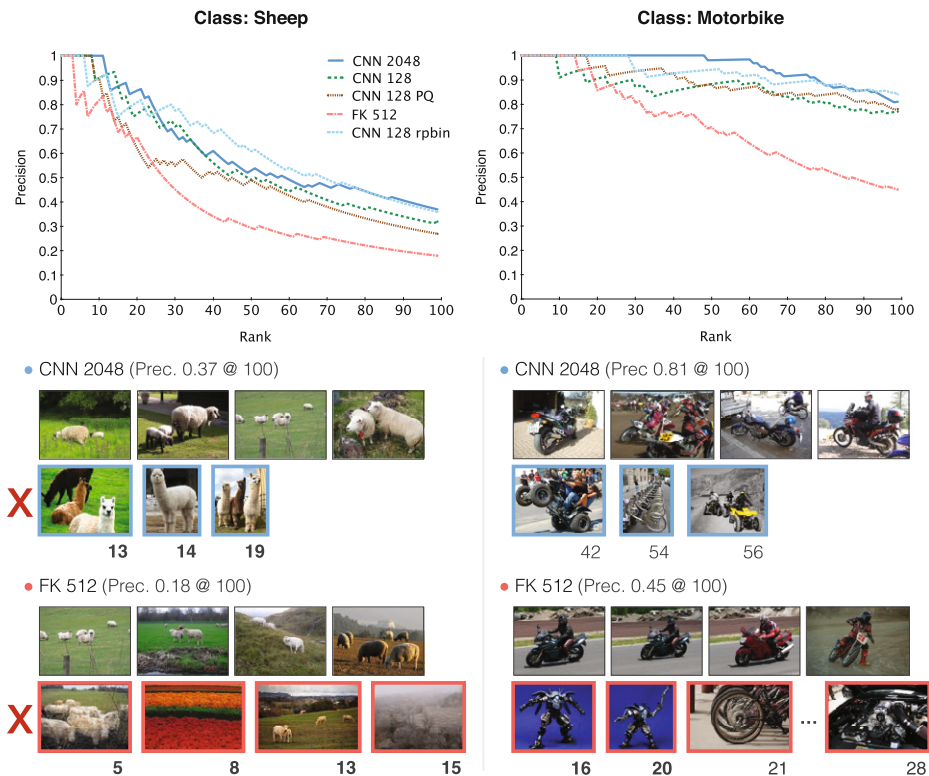


Fig. 2. Sample precision-rank curves and retrieved results for two queries over the combined VOC+MIRFLICKR data (Scenario 2a). In the bottom half of the figure, the top row in each case shows the first few results returned for each method and the second shows the top retrieved false positives with their rank.

Binarization is performed using the tight frame expansion method of [28], which has been recently successfully applied to local patch and face descriptors [29,30]. The binarization of zero-centred descriptors $\phi \in \mathbb{R}^m$ to binary codes $\beta \in \{0, 1\}^n$, $n > m$ is performed as $\beta = \text{sgn}(U\phi)$ where sgn is the sign function: $\text{sgn}(a) = 1$ iff $a > 0$ and 0 otherwise, and the Parseval tight frame U is computed by keeping the first m columns of an orthogonal matrix, obtained from a QR-decomposition of a random $n \times n$ matrix.

3.1 Results and Analysis

Scenario 1 (VOC Train/VOC Test). The PASCAL VOC dataset does not pose any major challenges for any of our features, which is not surprising given the close to decade of research on representations which perform well on this dataset. Even for the most challenging classes (*e.g.* ‘potted plant’) IFV produces fairly good results, with the top 12 images being true positives (Prec

Table 1. Retrieval results (Mean Prec @ 100) for the evaluation scenarios described in Sect. 2.2.

Scenario		VOC Only	Large-scale Retr.		Google Training	
		[1]	[2a]	[2b]	[3a]	[3b]
(a) FK	512	82.3	29.3	80.5		
(b) CNN	2K	92.1	55.4	95.4	88.5	90.9
(c) CNN	2K PQ	90.7	55.1	96.4	88.2	91.9
(d) CNN	128	92.1	51.0	95.1	88.1	92.3
(e) CNN	128 noaug	88.8	45.4	93.1	87.1	91.1
(f) CNN	128 BIN 2K	91.5	52.3	94.0	89.6	
(g) CNN	128 BIN 1K	90.0	50.1	94.0	89.5	
(h) CNN	128 PQ	90.1	50.5	94.6	88.2	92.1
(i) CNN	128 PQ-8	88.8	47.4	93.1	87.7	91.1

@ 100 = 0.58), and the top 92 images being true positives in the case of our 2048-dimensional ConvNet features (Prec @ 100 = 0.83).

Scenario 2a (VOC Train/VOC+distractors Test). Adding 1M distractor images from the MIRFLICKR-1M dataset has a significant impact on the results, with the task now being to retrieve true positives that constitute less than $\sim 0.02\%$ of the dataset. This is a more challenging scenario, and under this setting the superior performance of the ConvNet-based features, when compared to the state-of-the-art shallow representation (IFV), is much clearer to see. Some sample precision-rank curves for two queries, one particularly challenging (‘sheep’) and another less so (‘motorbike’) are shown in Fig. 2. We can make the following observations:

IFV Performance – It can be seen that IFV ([a] in Tables 1 and 2) performs the worst of all methods, despite being much higher dimensional ($\sim 1000\times$) and taking much longer to compute ($\sim 200\times$) compared to our CNN-128 method ([d]). Nonetheless, even for challenging classes such as ‘sheep’ IFV manages to pull out a few true positives at the top of the ranked list. However, the relative performance drop with rank is much sharper than with the ConvNet-based methods.

Bursty Images – Comparing the top-ranked negatives of the FK-512 method ([a] in Table 1) for ‘sheep’ to those of the CNN-2048 method ([b]), it can be seen that IFV appears to mistakenly rank highly ‘bursty’ images comprising repeating patterns or textures. This phenomenon is particularly evident for natural, outdoor scenes which explains why the performance drop of IFV is particularly severe in the ‘sheep’, ‘cow’ and ‘horses’ classes, as it appears that the ConvNet-based features are much more robust to such textured images.

Diversity – The diversity of the retrieved results is also much greater for ConvNet-based representations than for IFV, indicating that the classifier is able to make better generalisations using these features. For example, as seen in Fig. 2, whereas the top four retrieved results for the query ‘motorbike’ for the FK-512 method ([a]) all show a rider in a similar pose, on a racing bike on a race track, the top four retrieved results for the CNN-2048 method ([b]) depict a variety of different motorcycles (road, racing, off-road) from several different angles.

For the most part, compression of the ConvNet features does not appear to reduce their diversity appreciably, with the top-ranked results for all ConvNet methods, whether compressed or not, appearing to exhibit a similar diversity of results.

Compression – As mentioned above, the drop in performance in moving from ConvNet-based features to IFV is much greater than that incurred by any of the compression methods, and this seems to be strongly connected with the robustness of the ConvNet-based features, whether compressed or not, to the kind of ‘bursty’ textured images which IFV is susceptible to. This is remarkable given that comparing the size of the largest uncompressed ConvNet representation CNN-2048 ([b] in Table 2) to the smallest compressed one, CNN-128-PQ-8 ([i]), there is a $\sim 512\times$ size difference. In the case of the CNN-128-BIN-2K method ([f]), the mPrec @ 100 actually increases marginally when compared to the non-compressed codes ([d]) which, when visually inspecting the rankings, again can be explained by the additional robustness brought by compression.

Table 2. Dimensionality, storage requirements and computation time. The rows in this table correspond to those in Table 1. Timings for compression methods are specified as additional time added to the total feature encoding time, and those in parenthesis indicate GPU timings where applicable.

	Dim	Compression (bytes)		New Dim	Storage/1M ims	Comp. Time/im (s)
(a)	83,968	–			312.8 GB	10.32
(b)	2048	–			7.63 GB	0.35 (0.061)
(c)	2048	PQ	4 dims/sq (16 \times)	512	488 MB	+0.061
(d)	128	–			488 MB	0.34 (0.061)
(e)	128	noaug			488 MB	0.083 (0.024)
(f)	128	BIN	2048 bytes (2 \times)	4096	244 MB	+0.38 ms
(g)	128	BIN	1024 bytes (4 \times)	2048	122 MB	+0.22 ms
(h)	128	PQ	4 dims/sq (16 \times)	32	30.5 MB	+3.9 ms
(i)	128	PQ	8 dims/sq (32 \times)	16	15.3 MB	+2.0 ms

VOC Training

- ‘Chair’ – CNN 128 (Prec. 0.92 @ 100)



- ‘Train’ – CNN 128 (Prec. 1.0 @ 100)



Google Training

(Prec. 0.86 @ 100)



(Prec. 1.0 @ 100)



Fig. 3. Difference between retrieved results when training using VOC data and Google training data. Results are shown over the MIRFLICKR-1M dataset (Scenarios 2b and 3b).

The binary representations ([f] & [g]), combined with a linear SVM, also exhibit competitive performance despite the reduced memory footprint. The ranking of such features can be significantly sped-up using hardware-accelerated Hamming distance computation, which, however, requires a different ranking model, which is left for future work. For its superior compression ratios and negligible impact on performance, product quantization remains an obvious choice for the compression of ConvNet features. The fact that the ConvNet features are very sparse, with the CNN-128 representation typically being over 60% zeros, is one reason why they are so amenable to compression, and it is possible that with compression methods geared specifically to capitalise on this sparsity even higher compression ratios could be achieved.

Scenario 2b (VOC Train/MIRFLICKR Test). Given that the MIRFLICKR-1M dataset contains many instances of all of the PASCAL VOC classes, moving to testing solely on MIRFLICKR leads to a jump in performance of the results across all methods. Nonetheless, this scenario provides a closer representation of the performance of a real-world on-the-fly object category retrieval system, given that the image statistics of the MIRFLICKR-1M dataset are not known in advance.

Scenario 3a (Google Train/MIRFLICKR Test). Switching to noisy training images from Google rather than the pre-curated PASCAL VOC training images as expected results in a small drop ($\sim 6\%$) across the board for all methods. However, the precision at the top of the ranking remains subjectively very good. Nonetheless, as shown in Fig. 3, the actual images returned from the dataset are very different, which reflects the differences in the training data sourced from Google Image search versus that from the curated dataset. For example, a query for ‘chair’ returns predominantly indoor scenes with regular dining-table chairs when using VOC training data, and more avant-garde, modern designs, generally centred in the frame when using Google training data.

Scenario 3b (Google Train + negative pool/MIRFLICKR Test). In this scenario, we switch to using a fixed pool of negative data sourced from a set of ‘negative’ queries, and it can be seen how this improves the results by up to $\sim 5\%$. This may be a result of the larger negative training pool size ($\sim 16,000$ images *vs.* $\sim 4,000$ images when using queries for all other VOC classes to provide the negative data as we do in Scenario 3a). Given the assumed lack of coverage in the fixed negative image pool (as it is sourced by issuing queries for deliberately non-specific terms to facilitate its application to as broad a range of queries as possible), this suggests that to a certain extent lack of diversity can be made up for by using a larger number of negative training images.

3.2 Implementation Details

Our implementation of IFV and ConvNet image representations follows that of [21]. In more detail, for IFV computation we use their setting ‘FK IN 512 (x,y)’, which corresponds to: (i) dense rootSIFT [31] local features with spatial extension [32], extracted with 3 pixel step over 7 scales ($\sqrt{2}$ scaling factor); (ii) Improved Fisher vector encoding [13] using a GMM codebook with 512 Gaussians; (iii) intra normalisation [33] of the Fisher vector.

Our ConvNet training and computation framework is based on the publicly available Caffe toolbox [34]. The two ConvNet configurations, considered in this paper (‘CNN M 2048’ and ‘CNN M 128’) are pre-trained on the ImageNet ILSVRC-2012 dataset using the configurations described in [21]⁴. Namely, they contain 5 convolutional and 2 fully-connected layers, interleaved with rectification non-linearities and max-pooling. The stack of layers is followed by a 1000-way soft-max classifier, which is removed after pre-training is finished (turning a ConvNet from an ImageNet classifier to a generic image descriptor). The only difference between the two ConvNets is the dimensionality of the second fully-connected layer, which is 2048 for ‘CNN M 2048’ and 128 for ‘CNN M 128’.

In order to provide a similar setup to our on-the-fly architecture in Sect. 4, which uses a linear predictor $\langle \mathbf{w}, \phi(I) \rangle$ learnt using SVM hinge loss and a quadratic regulariser, as our learning stage we use a standard linear support vector machine implementation. The C parameter is determined using the VOC validation set for scenario 1, and fixed at 0.25 for all other experiments.

4 On-the-fly Architecture

Having evaluated various image representations in Sect. 3, we now describe the architecture of the object category retrieval system, which fully exploits the advantages of ConvNet image representations. From the user experience point of view, the main requirement to our system is instant response: the first ranking of the repository images should be obtained immediately (in under a second), with a potential improvement over time. This dictates the following design choice:

⁴ http://www.robots.ox.ac.uk/~vgg/software/deep_eval/.

downloading the training images from the Internet should be carried out in parallel with training a model on the already downloaded images in the on-line fashion. As a result, at any point of time, the current model can be used to perform ranking of the dataset images.

For this approach to work, however, the chosen image representation should satisfy the following requirements: (i) highly discriminative, so that even a handful of training samples are sufficient to learn a linear ranking model; (ii) fast-to-compute, to maximise the amount of training data processed within the allocated time budget; (iii) low memory footprint, to allow for storing large-scale datasets in the main memory, and ranking them efficiently. As has been demonstrated in Sect. 3, a ConvNet image representation is a perfect match for these requirements. Indeed, pre-training on a large image collection (ImageNet) leads to highly discriminative representation, and even a few training samples are sufficient for training an accurate linear model; ConvNet features can be computed very quickly on the highly-parallel GPU hardware; they have low dimensionality (even without PQ compression) and can be instantly scored using a linear model on the GPU.

Our on-the-fly architecture is illustrated in Fig. 1. It is divided into the CPU-based front-end, which controls the graphical user interface and downloads the training images from the Internet, and the GPU-based back-end, which continually trains the ranking model on the downloaded images and periodically applies it to the repository. The category retrieval is carried out as follows:

Off-line (pre-processing). To allow for fast processing, the ConvNet features for the target dataset images are pre-computed off-line, using the CNN-128 architecture. We also prepare the fixed negative image pool for *all queries* by issuing our negative pool queries (see Sect. 2.2) to both Bing and Google image search, and downloading the returned URLs. The negative image feature features are also pre-computed. The memory requirements for storing the pre-computed features are as follows: 488 MB for the MIRFLICKR-1M dataset and 78 MB for the pool of 16K negative features. It is thus feasible to permanently store the features of both negative and dataset images in the high-speed GPU memory even without compression of any kind (a consumer-grade NVIDIA GTX Titan GPU, used in our experiments, is equipped with 6 GB RAM). As noted in Sect. 2, the ConvNet features can be compressed further by up to 16 \times using product quantization without significant degradation in performance, making datasets of up to 160M images storable in GPU memory, setting 1 GB aside for storage of the model (compared to 10M images without compression), and more if multiple GPUs are used. Many recent laptops are fitted with a GPU containing similar amounts of memory, making our system theoretically runnable on a single laptop. Furthermore, whilst storing the target repository on the GPU is preferable in terms of the ranking time, in the case of datasets of 1B+ images it can be placed in the CPU memory, which typically has a larger capacity.

On-line (CPU front-end). Given a textual query, provided by a user (*e.g.* in a browser window), the front-end starts by downloading relevant images, which will be used as positive samples for the queried category and fed to the GPU

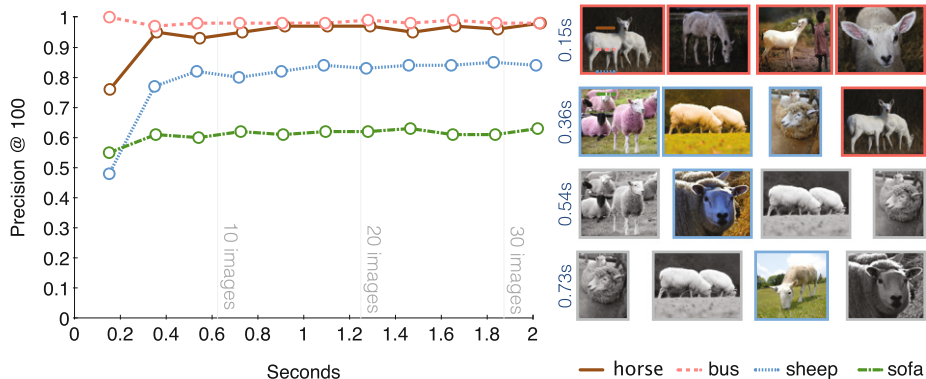


Fig. 4. Precision @ 100 against training time for four queries using our on-the-fly architecture. The number of images in the dynamically expanding positive image training pool over time is also marked on the plot. The top-4 returned images for the ‘sheep’ query at the first four time-steps (up to 0.73s) is shown to the right. False positives are outlined in red, and new images in the top-4 at each time step are outlined in blue. Even for this moderately challenging query, the model settles in under a second (Color figure online).

back-end. At regular time intervals, the front-end receives a ranked list of dataset images from the back-end, and displays them in the user interface.

On-line (GPU back-end). The GPU back-end runs in parallel with the front-end, and is responsible for both training the ranking model and applying it to the dataset. *Training* an L_2 -regularised linear SVM model is carried out using the mini-batch SGD with Pegasos updates [35]: at iteration t , the learning rate is $\frac{1}{\lambda t}$, where λ is the L_2 -norm regularisation constant, set to 1 in our experiments. Each batch contains an equal amount of positive and negative samples; the total batch size was set to $B = 32$ in our experiments. The training commences as soon as the first positive image has been downloaded and is received from the front-end, after which B random crops are taken each iteration from the pool of positive training images downloaded so far. The front-end in the meantime will continue downloading new images from the Internet, constantly increasing the size of the positive image pool and the diversity of the extracted crops. We note that while the positive image features need to be computed on-the-fly, this is very quick in the case of ConvNets. *Ranking* takes place using the current SVM model every τ seconds (we used $\tau \sim 0.18$). As mentioned above, the pre-computed dataset features are pre-stored on a GPU, so the scores for 1M images are computed in ≈ 0.01 s. The 1M scores are then ranked (also on GPU, ≈ 0.002 s) and the list of the top-ranked images is passed to the front-end to be displayed to the user. All components of the GPU back-end are implemented within the same framework, derived from Caffe [34].

4.1 System Performance

In order to evaluate the real-world performance of the system, we ran queries for several PASCAL VOC classes and tracked how the performance (measured in terms of Precision @ 100) evolved over time. To simulate the latency introduced by downloading images from the Internet, we limited the rate of positive images entering the network to 12 images/second (which is what we found to be a typical average real-world rate on our test system). These images were sampled randomly from the top-50 image URLs returned from Google Image search.

The results of these experiments for four classes are shown in Fig. 4. Even for some of the most challenging PASCAL VOC classes ‘sheep’ and ‘sofa’, the performance converged to its final value in ~ 0.6 s, and as can be seen from the evolving ranking at each time-step the ordering at the top of the ranking generally stabilizes within a second, showing a good diversity of results. For easier classes such as ‘airplane’, convergence and stabilization occurs even faster.

In real terms, this results in a typical query time for our on-the-fly architecture, from entering the text query to viewing the ranked retrieved images, of **1–2 s and often less** to complete convergence and stabilization of results. However, one of the advantages of our proposed architecture is that it is adaptable to differing query complexity, and we can return good results early whilst still continuing to train in the background if necessary, exposing the classifier to an expanding pool of training data as it is downloaded from the web and updating the ranked list on-the-fly.

Novel On-the-fly Queries. Although experimental results have thusfar only been presented for the PASCAL VOC classes, the advantage of an on-the-fly architecture is that no limitation is imposed on the object categories which can be queried for, as a new classifier can be trained on demand (in our case using

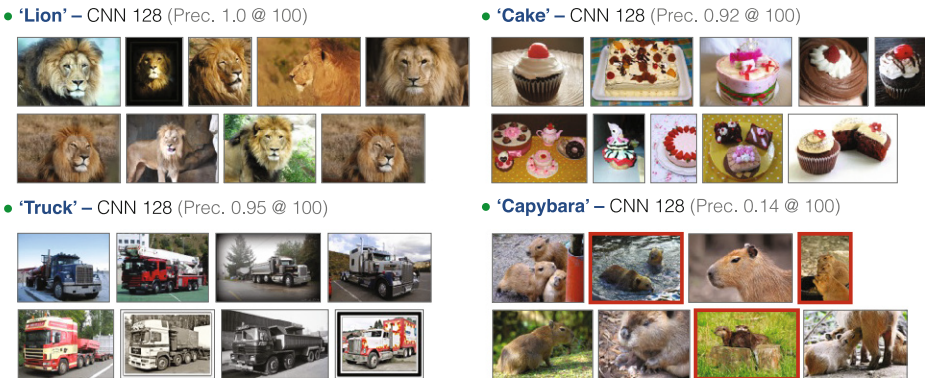


Fig. 5. Sample results for queries outside of the twenty PASCAL VOC classes. False positives are outlined in red (Color figure online).

Google Image search as a ‘live’ source of training data). We present some additional selected results of the on-the-fly system in Fig. 5, using the same setup as in *Scenario 3b* and query terms disjunct from the twenty PASCAL VOC classes. It can be seen that the architecture is very much generalisable to query terms outside of the PASCAL category hierarchy.

5 Conclusion

In this paper we have presented a system for on-the-fly object category retrieval, which builds upon the recent advances in deep convolutional image representations. We demonstrated how such representations can be efficiently compressed and used in a novel incremental learning architecture, capable of retrieval across datasets of 1M+ images within seconds and running entirely on a single GPU.

For larger datasets the CPU, or multiple GPU cards, could be employed for ranking once the classifier has been learnt on the GPU. Along with further investigation of how the diversity of the ranked results changes over time, this is the subject of future work.

Acknowledgements. This work was supported by the EPSRC and ERC grant Vis-Rec no. 228180. We gratefully acknowledge the support of NVIDIA Corporation with the donation of the GPUs used for this research.

References

1. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: a large-scale hierarchical image database. In: Proceedings of CVPR (2009)
2. Soomro, K., Zamir, A.R., Shah, M.: UCF101: a dataset of 101 human actions classes from videos in the wild. CoRR abs/1212.0402 (2012)
3. Sivic, J., Zisserman, A.: Video Google: a text retrieval approach to object matching in videos. In: Proceedings of ICCV, vol. 2, pp. 1470–1477 (2003)
4. Nister, D., Stewenius, H.: Scalable recognition with a vocabulary tree. In: Proceedings of CVPR, pp. 2161–2168 (2006)
5. Philbin, J., Chum, O., Isard, M., Sivic, J., Zisserman, A.: Object retrieval with large vocabularies and fast spatial matching. In: Proceedings of CVPR (2007)
6. Jegou, H., Douze, M., Schmid, C.: Hamming embedding and weak geometric consistency for large scale image search. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) ECCV 2008, Part I. LNCS, vol. 5302, pp. 304–317. Springer, Heidelberg (2008)
7. Jégou, H., Douze, M., Schmid, C., Pérez, P.: Aggregating local descriptors into a compact image representation. In: Proceedings of CVPR (2010)
8. Chatfield, K., Zisserman, A.: VISOR: towards on-the-fly large-scale object category retrieval. In: Lee, K.M., Matsushita, Y., Rehg, J.M., Hu, Z. (eds.) ACCV 2012, Part II. LNCS, vol. 7725, pp. 432–446. Springer, Heidelberg (2013)
9. Fernando, B., Tuytelaars, T.: Mining multiple queries for image retrieval: on-the-fly learning of an object-specific mid-level representation. In: Proceedings of ICCV (2013)

10. Liu, Y., Xu, D., Tsang, I.W., Luo, J.: Using large-scale web data to facilitate textual query based retrieval of consumer photos. In: Proceedings of the 17th ACM International Conference on Multimedia, MM 2009, pp. 55–64 (2009)
11. Parkhi, O.M., Vedaldi, A., Zisserman, A.: On-the-fly specific person retrieval. In: International Workshop on Image Analysis for Multimedia Interactive Services. IEEE (2012)
12. Chatfield, K., Lempitsky, V., Vedaldi, A., Zisserman, A.: The devil is in the details: an evaluation of recent feature encoding methods. In: Proceedings of BMVC (2011)
13. Perronnin, F., Sánchez, J., Mensink, T.: Improving the fisher kernel for large-scale image classification. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) ECCV 2010, Part IV. LNCS, vol. 6314, pp. 143–156. Springer, Heidelberg (2010)
14. Perronnin, F., Akata, Z., Harchaoui, Z., Schmid, C.: Towards good practice in large-scale learning for image classification. In: Proceedings of CVPR, pp. 3482–3489 (2012)
15. Jégou, H., Douze, M., Schmid, C.: Product quantization for nearest neighbor search. IEEE PAMI **33**(1), 117–128 (2011)
16. Torralba, A., Fergus, R., Freeman, W.T.: 80 million tiny images: a large dataset for non-parametric object and scene recognition. IEEE PAMI **30**, 1958–1970 (2008)
17. Raginsky, M., Lazebnik, S.: Locality sensitive binary codes from shift-invariant kernels. In: NIPS (2009)
18. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. Proc. IEEE **86**, 2278–2324 (1998)
19. Krizhevsky, A., Sutskever, I., Hinton, G.E.: ImageNet classification with deep convolutional neural networks. In: NIPS, pp. 1106–1114 (2012)
20. Zeiler, M.D., Fergus, R.: Visualizing and understanding convolutional networks. CoRR abs/1311.2901 (2013)
21. Chatfield, K., Simonyan, K., Vedaldi, A., Zisserman, A.: Return of the devil in the details: delving deep into convolutional nets. In: Proceedings of BMVC (2014)
22. Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The PASCAL Visual Object Classes (VOC) challenge. IJCV **88**, 303–338 (2010)
23. Huiskes, M.J., Lew, M.S.: The MIR flickr retrieval evaluation. In: MIR 2008: Proceedings of the 2008 ACM International Conference on Multimedia Information Retrieval (2008)
24. Huiskes, M.J., Thomee, B., Lew, M.S.: New trends and ideas in visual concept detection: the MIR flickr retrieval evaluation initiative. In: MIR 2010: Proceedings of the 2010 ACM International Conference on Multimedia Information Retrieval, pp. 527–536 (2010)
25. Donahue, J., Jia, Y., Vinyals, O., Hoffman, J., Zhang, N., Tzeng, E., Darrell, T.: DeCAF: a deep convolutional activation feature for generic visual recognition. CoRR abs/1310.1531 (2013)
26. Razavian, A., Azizpour, H., Sullivan, J., Carlsson, S.: CNN Features off-the-shelf: an Astounding Baseline for Recognition. CoRR abs/1403.6382 (2014)
27. Sánchez, J., Perronnin, F.: High-dimensional signature compression for large-scale image classification. In: Proceedings of CVPR (2011)
28. Jégou, H., Furon, T., Fuchs, J.J.: Anti-sparse coding for approximate nearest neighbor search. In: Proceedings of ICASSP, pp. 2029–2032 (2012)
29. Simonyan, K., Vedaldi, A., Zisserman, A.: Learning local feature descriptors using convex optimisation. IEEE PAMI **36**, 1573–1585 (2014)
30. Parkhi, O.M., Simonyan, K., Vedaldi, A., Zisserman, A.: A compact and discriminative face track descriptor. In: Proceedings of CVPR. IEEE (2014)

31. Arandjelović, R., Zisserman, A.: Three things everyone should know to improve object retrieval. In: Proceedings of CVPR (2012)
32. Sánchez, J., Perronnin, F., Emidio de Campos, T.: Modeling the spatial layout of images beyond spatial pyramids. *Pattern Recogn. Lett.* **33**, 2216–2223 (2012)
33. Arandjelović, R., Zisserman, A.: All about VLAD. In: Proceedings of CVPR (2013)
34. Jia, Y.: Caffe: an open source convolutional architecture for fast feature embedding (2013). <http://caffe.berkeleyvision.org/>
35. Shalev-Shwartz, S., Singer, Y., Srebro, N.: Pegasos: primal estimated sub-gradient SOLver for SVM. In: Proceedings of ICML, vol. 227 (2007)