

# Cross-Dimensional Weighting for Aggregated Deep Convolutional Features

Yannis Kalantidis<sup>(✉)</sup>, Clayton Mellina, and Simon Osindero

Computer Vision and Machine Learning Group Flickr, Yahoo, San Francisco, USA  
ykal@yahoo-inc.com, clayton@yahoo-inc.com

**Abstract.** We propose a simple and straightforward way of creating powerful image representations via cross-dimensional weighting and aggregation of deep convolutional neural network layer outputs. We first present a generalized framework that encompasses a broad family of approaches and includes cross-dimensional pooling and weighting steps. We then propose specific non-parametric schemes for both spatial- and channel-wise weighting that boost the effect of highly active spatial responses and at the same time regulate burstiness effects. We experiment on different public datasets for image search and show that our approach outperforms the current state-of-the-art for approaches based on pre-trained networks. We also provide an easy-to-use, open source implementation that reproduces our results.

## 1 Introduction

Visual image search has been evolving rapidly in recent years with hand-crafted local features giving way to learning-based ones. Deep Convolutional Neural Networks (CNNs) were popularized by the seminal work of Krizhevsky *et al.* [19] and have been shown to “effortlessly” improve the state-of-the-art in multiple computer vision domains [29], beating many highly optimized, domain-specific approaches. It comes as no surprise that such features, based on deep networks, have recently also dominated the field of visual image search [3–5, 29].

Many recent image search approaches are based on deep features, *e.g.*, Babenko *et al.* [4, 5] and Razavian *et al.* [3, 29] proposed different pooling strategies for such features and demonstrated state-of-the-art performance in popular benchmarks for *compact* image representations, *i.e.*, representations of up to a few hundred dimensions.

Motivated by these advances, in this paper we present a simple and straightforward way of creating powerful image representations via cross-dimensional weighting and aggregation. We place our approach in a general family of approaches for multidimensional aggregation and weighting and present a specific instantiation that we have thus far found to be most effective on benchmark tasks.

We base our cross-dimensional weighted features on a generic deep convolutional neural network. Since we aggregate outputs of convolutional layers before the fully connected ones, the data layer can be of arbitrary size [20]. We therefore avoid resizing and cropping the input image, allowing images of different

aspect ratios to keep their spatial characteristics intact. After extracting deep convolutional features from the last spatial layer of a CNN, we apply weighting both spatially and per channel before sum-pooling to create a final aggregation. We denote features derived after such cross-dimensional weighting and pooling as *CroW* features.

Our contributions can be summarized as follows:

- We present a generalized framework that sketches a family of approaches for aggregation of convolutional features, including cross-dimensional weighting and pooling steps.
- We propose non-parametric weighting schemes for both spatial- and channel-wise weighting that boost the effect of highly active spatial responses and regulate the effect of channel burstiness respectively.
- We present state-of-the-art results on three public datasets for image search without any fine-tuning.

With a very small computational overhead, we are able to improve the state-of-the-art in visual image search. For the popular *Oxford* [26] and *Paris* [27] datasets, the mean average precision for our *CroW* feature is over 10% higher than the previous state-of-the-art for compact visual representations. Additionally, our features are trivially combined for simple query expansion, enjoying even better performance. *We provide an easy-to-use, open source implementation that reproduces our results on GitHub*<sup>1</sup>.

The paper is structured as follows: In Sect. 2 we present and discuss related work, while in Sect. 3 we present a general framework for weighted pooling to orient past work and our own explorations. In Sect. 4 we describe two complimentary feature weighting schemes, and we present experimental results for visual search in Sect. 5. The paper concludes with Sect. 6.

## 2 Related Work

Until recently, the vast majority of image search approaches were variants of the bag-of-words model [32] and were based on local features, typically SIFT [21]. Successful extensions include soft assignment [27], spatial matching [2, 26], query expansion [1, 6, 7, 35], better descriptor normalization [1], feature selection [36, 38], feature burstiness [15] and very large vocabularies [22]. All the aforementioned strategies perform very well for object retrieval but are very hard to scale, as each image is represented by hundreds of patches, causing search time and memory to suffer.

The community therefore recently turned towards global image representations. Starting from local feature aggregation strategies like VLAD [16] or Fisher Vectors [24] multiple successful extensions have arisen [9, 12, 33, 34], slowly increasing the performance of such aggregated features and closing the gap between global and bag-of-word representations for image search. Triangulation embedding with democratic aggregation [17] was shown to give state-of-the-art results for SIFT-based architectures, while handling problems related to

<sup>1</sup> <https://github.com/yahoo/crow>.

burstiness and interactions between unrelated descriptors prior to aggregation. Recently, Murray and Perronnin [23] generalized max-pooling from bag-of-words to Fisher Vector representations achieving high performance in search as well as classification tasks.

After the seminal work of Krizhevsky *et al.* [19], image search, along with the whole computer vision community, embraced the power of deep learning architectures. Out-of-the-box features from pre-trained Convolutional Neural Networks (CNNs) were shown to effortlessly give state-of-the-art results in many computer vision tasks, including image search [29].

Among the first to more extensively study CNN-based codes for image search were Babenko *et al.* [5] and Razavian *et al.* [3, 29]. They experimented with aggregation of responses from different layers of the CNN, both fully connected and convolutional. They introduced a basic feature aggregation pipeline using max-pooling that, in combination with proper normalization and whitening was able to beat all aggregated local feature based approaches for low dimensional image codes. Gong *et al.* [10] used orderless VLAD pooling of CNN activations on multiple scales and achieved competitive results on classification and search tasks.

Very recently Tolias *et al.* [37] proposed max-pooling over multiple image regions sampled on the final convolutional layer. Their approach achieves state-of-the-art results and is complementary to our cross-dimensional weighting. Cimpoi *et al.* [8] also recently proposed using Fisher Vector aggregation of convolutional features for texture recognition. Their approach achieves great performance, it is however computationally demanding; PCA from 65K dimensions alone requires multiplication with a very large matrix. Our approach is training- and parameter-free, with only a very small computational overhead.

In another very recent related work, Babenko and Lempitsky proposed the SPoC features [4] with slightly different design choices from the pipeline of [5] and sum- instead of max-pooling. As the latter approach is very related to ours, we discuss the differences of the two approaches in the following sections and explain SPoC in terms of the proposed aggregation framework.

The first approaches that *learn* features for landmark retrieval [11, 28] are presented at the current ECCV conference. Both approaches use clean annotated data and fine-tune a deep CNN for feature extraction using a pairwise [28] or ranking [11] loss. These approaches are now state-of-the-art in the most common benchmarks. Still, our proposed features are not far behind, without requiring training or clean annotated data.

## 3 Framework for Aggregation of Convolutional Features

### 3.1 Framework Overview

In this section we present a simple and straightforward way of creating powerful image representations. We start by considering a general family of approaches that can be summarized as proceeding through the following steps. Greater details and motivations for these steps will be given in subsequent sections, along with the specific instantiation that we have thus far found to be most effective on benchmark tasks.

- 1: Perform spatially-local pooling.** Sum-pooling or max-pooling over a spatially local neighborhood within each channel of a convolutional layer, with neighborhood size  $w \times h$  and stride  $s$ . Some limiting cases include: (1) a pooling neighborhood that occupies the full spatial extent of each channel (i.e. global pooling); and (2) a  $1 \times 1$  pooling neighborhood (effectively not doing pooling at all). After pooling, we have a three-dimensional tensor of activities.
- 2: Compute spatial weighting factors.** For each location  $(i, j)$  in the locally pooled feature maps we assign a weight,  $\alpha_{ij}$ , that is applied to each channel at that location.
- 3: Compute channel weighting factors.** For each channel  $k$ , we assign a weight,  $\beta_k$  that is applied to each location in that channel.
- 4: Perform weighted-sum aggregation.** We apply the previously derived weights location-wise and channel-wise before using a channel-wise sum to aggregate the full tensor of activities into a single vector.
- 5: Perform vector normalization.** The resulting vector is then normalized and power-transformed. A variety of norms can be used here.
- 6: Perform dimensionality reduction.** We then reduce the dimensionality of the normed-vector. PCA is a typical choice here, and we may also choose to perform whitening or other per-dimension scalings on entries of the dimensionality reduced vector.
- 7: Perform final normalization.** We then apply a second and final normalization step.

Algorithm 1 summarises these steps as pseudocode.

---

**Algorithm 1.** Framework for Aggregation of Convolutional Features

---

```


input : 3d feature tensor  $\mathcal{X}$ , pooling nhood size  $w \times h$ , stride  $s$ , and type  $p$ , spatial weight
         generation function,  $\Omega_s$ , channel weight generation function,  $\Omega_c$ , initial norm type,
          $a$ , and power scaling,  $b$ , pre-trained whitening parameters  $W$ , final feature
         dimensionality  $K'$ , final norm type,  $c$ 
output:  $K'$ -dimensional aggregate feature vector  $\mathcal{G} = \{g_1, \dots, g_K\}$ 
1  $\tilde{\mathcal{X}} = \text{pool}(\mathcal{X}; w, h, s, p)$  // Initial local pooling
2  $\Omega_s(\tilde{\mathcal{X}}) \rightarrow \alpha_{ij} \forall i, j$  // Spatial weighting
3  $\Omega_c(\tilde{\mathcal{X}}) \rightarrow \beta_k \forall k$  // Channel weighting
4  $f_k = \sum_{i=1}^W \sum_{j=1}^H \alpha_{ij} \beta_k \mathcal{X}_{kij} \forall k$ 
5  $\hat{\mathcal{F}} = \text{pnorm}(\mathcal{F}; a, b)$  // Normalize and powerscale
6  $\tilde{\mathcal{F}} = \text{PCA}(\hat{\mathcal{F}}; W, K')$  // dim. reduction and whitening
7  $\mathcal{G} = \text{norm}(\tilde{\mathcal{F}}, c)$  // Normalize again

```

---

### 3.2 Cross-Dimensional Weighting

Let  $\mathcal{X} \in \mathbb{R}^{(K \times W \times H)}$  be the 3-dimensional *feature* tensor from a selected layer  $l$ , where  $K$  is the total number of channels and  $W, H$  the spatial dimensions



**Fig. 1.** Prior to aggregation, the convolutional features can be weighted channel-wise by a weight vector  $\beta$  and weighted location-wise by a weight matrix  $\alpha$  such that  $\mathcal{X}'_{kij} = \alpha_{ij}\beta_k\mathcal{X}_{kij}$ . The weighted features  $\mathcal{X}'$  are sum-pooled to derive an aggregate feature.

of that layer. As mentioned above, the spatial dimensions may vary per image depending on its original size, but we omit image-specific subscripts here for clarity.

We denote the entry in  $\mathcal{X}$  corresponding to channel  $k$ , at spatial location  $(i, j)$  as  $\mathcal{X}_{kij}$ . For notational convenience, we also denote the channel-wise matrices of  $\mathcal{X}$  as  $\mathcal{C}^{(k)}$ , where  $\mathcal{C}^{(k)}_{ij} = \mathcal{X}_{kij}$ . Similarly, we use  $\lambda^{(ij)}$  to denote the vector of channel responses at location  $(i, j)$ , where  $\lambda_k^{(ij)} = \mathcal{X}_{kij}$ .

A weighted feature tensor  $\mathcal{X}'$  is produced by applying per-location weights,  $\alpha_{ij}$ , and per-channel weights,  $\beta_k$ , to feature tensor  $\mathcal{X}$  as illustrated in Fig. 1:

$$\mathcal{X}'_{kij} = \alpha_{ij}\beta_k\mathcal{X}_{kij} \tag{1}$$

The weighted feature tensor is aggregated by sum-pooling per channel. Let *aggregated feature vector*  $\mathcal{F} = \{f_1, \dots, f_k\}$  associated with the layer  $l$  be the vector of weight-summed activations per channel:

$$f_k = \sum_{i=1}^W \sum_{j=1}^H \mathcal{X}'_{kij} \tag{2}$$

After aggregation, we follow what was shown to be the best practice [3,29] and L2-normalize  $\mathcal{F}$ , then whiten using parameters learnt from a separate dataset and L2-normalize again. We denote the features that are derived from the current framework as *Cross-dimensional Weighted* or *CroW* features.

## 4 Feature Weighting Schemes

In this section we present our non-parametric spatial and channel weighting for Steps 2 and 3 of the framework. We propose a spatial weighting derived from the spatial activations of the layer outputs themselves and a channel weighting derived from channel sparsity.

### 4.1 Response Aggregation for Spatial Weighting

We propose a method to derive a spatial weighting based on the normalized total response across all channels. Let  $\mathbf{S}' \in \mathbb{R}^{(W \times H)}$  be the matrix of aggregated responses from all channels *per spatial location*, which we compute by *summing* feature maps  $\mathbf{C}^{(k)}$ :

$$\mathbf{S}' = \sum_k \mathbf{C}^{(k)}. \tag{3}$$


After normalization and power-scaling we get aggregated spatial response map  $\mathbf{S}$ , whose value at spatial location  $(i, j)$  is given by:

$$S_{ij} = \left( \frac{S'_{ij}}{\left(\sum_{m,n} S'_{mn} \right)^{1/a}} \right)^{1/b}, \tag{4}$$


After computing the 2d spatial aggregation map  $\mathbf{S}$  for feature  $\mathbf{X}$ , we can apply it independently on every channel, setting  $\alpha_{ij} = S_{ij}$  and using  $\alpha_{ij}$  as in Eq. 1.

We experimented with different norms for normalizing the aggregate responses  $\mathbf{S}'$ , *i.e.*, L1, L2, *inf*, power normalization with  $a = 0.5$  [25]. We found that image search performance remains very high in all cases and the differences are very small, usually less than 0.01 in mAP. We therefore choose to use the L2 norm and  $b = 2$  for our spatial aggregation maps, before applying them to the features.

We visualize highly weighted spatial locations in Fig. 3 with images from the *Paris* [27] dataset. Our spatial weighting boosts features at locations with salient visual content and down weights non-salient locations. Notably, similar visual elements are boosted under our weighting despite large variation in lighting and perspective.



**Fig. 2.** Mean  $\tilde{S}_{ij}$  plotted against channel sparsity at the corresponding location. The correlation of channel-wise sparsity for the 55 images in the query-set of the *Paris* dataset. Images are sorted by landmark class in both dimensions.



**Fig. 3.** Visualization of spatial weighting by aggregate response. On the left we show original images in the *Paris* dataset along with their spatial weights. On the right we visualize the receptive fields of the 7 highest weighted locations and the 7 lowest weighted locations for each image. The top two images are of Notre Dame and the bottom two are of the Panthéon.

In Fig. 2a we show the relationship between our spatial weights  $S_{ij}$  and the sparsity of the channel responses  $\lambda^{(ij)}$ . We compute the spatial weight  $S_{ij}$  of every location in the *Paris* dataset and normalize each by the maximum spatial weight for the image in which it occurs, which we denote  $\tilde{S}_{ij}$ . The mean  $\tilde{S}_{ij}$  for each level of channel sparsity at the corresponding location is plotted as cyan in Fig. 2a.

It can be seen that our spatial weighting tends to boost locations for which multiple channels are active relative to other spatial locations of the same image. This suggests that our spatial weighting is a non-parametric and computationally cheap way to favor spatial locations for which features co-occur while also accounting for the strength of feature responses. We speculate that these locations are more discriminative as there are combinatorially more configurations at mid-ranges of sparsity.

## 4.2 Sparsity Sensitive Channel Weighting

We now propose a method to derive a channel weighting based on the sparsity of feature maps. We expect that similar images will have similar occurrence rates for a given feature. For each channel  $k$  we find  $Q_k$ , the proportion of non-zero responses, and compute the *per-channel sparsity*,  $\bar{\Xi}_k$ , as:

$$\bar{\Xi}_k = 1 - Q_k, \quad (5)$$



where  $\mathcal{Q} = \frac{1}{WH} \sum_{ij} \mathbb{1}[\lambda^{(ij)} > 0]$ . In Fig. 2b we visualize the pair-wise correlation of the vectors of channel sparsities  $\mathcal{E} \in \mathbb{R}^K$  for images in the query-set of the *Paris* dataset. The query-set for the *Paris* dataset contains 55 images total, 5 images each for 11 classes of Paris landmarks. We order the images by class. It is apparent that channel sparsities  $\mathcal{E}$  are highly correlated for images of the same landmark and less correlated for images of different landmarks. It appears that the sparsity pattern of channels contains discriminative information.

Since we sum-pool features  $\lambda^{(ij)}$  over spatial locations when we derive our aggregated feature, channels with frequent feature occurrences are already strongly activated in the aggregate feature. However, infrequently occurring features could provide important signal if, for example, the feature consistently occurs though only a small number of times in images of the same class. Motivated by this insight, we devise a channel weighting scheme similar to the concept of *inverse document frequency*. That is, we boost the contribution of rare features in the overall response by using the per-channel weight,  $\mathcal{I}_k$ , defined as:

$$\mathcal{I}_k = \log \left( \frac{K\epsilon + \sum_h \mathcal{Q}_h}{\epsilon + \mathcal{Q}_k} \right), \tag{6}$$

where  $\epsilon$  is a small constant added for numerical stability.

Our sparsity sensitive channel weighting is also related to and motivated by the notion of intra-image visual burstiness [15]. Channels with low sparsity correspond to filters that give non-zero responses in many image regions. This implies some spatially recurring visual elements in the image, that were shown to negatively affect matching [15]. Although we don't go as far as [17] and try to learn a “democratic” matching kernel, our sparsity sensitive weights do down-weight channels of such bursty convolutional filters.

To provide further insight into the effect of our sparsity-sensitive channel weights (SSW), we visualize the receptive fields of active locations in channels that our weights boost.



**Fig. 4.** Regions corresponding to locations that contribute (are non-zero) to the 10 channels with the highest sparsity-sensitive weights for the four images of Fig. 3.



In Fig. 4 we show all receptive fields that are non-zero in (one or more) of the channels with the highest sparsity-sensitive channel weights. As values from these channels are increased before aggregation, our approach gives more weight to CNN outputs that correspond to the image regions shown on the right.

### 4.3 Discussion

Using the framework described in Sect. 3, we can explain different approaches in terms of their pooling, weighting and aggregation steps; we illustrate some interesting cases in Table 1. For example, approaches that aggregate the output of a max-pooling layer of the convolutional neural network are essentially performing max-pooling in Step 1.

In terms of novelty, it is noteworthy to restate that the spatial weighting presented in Sect. 4.1 corresponds to a well known principle, and approaches like [8, 17, 23] have addressed similar ideas. Our spatial weighting is notable as a simple and strong baseline. Together with the channel weighting, the *CroW* features are able to deliver state-of-the-art results at practically the same computational cost as off-the-self features.

**Uniform Weighting.** If we further uniformly set both spatial and channel weights and then perform sum-pooling per channel we end up with a simpler version of *CroW* features, that we denote as *uniform CroW* or *uCroW*.

**Relation to SPoC [4] Features.** SPoC [4] can be described in terms of our framework as illustrated in Table 1. *CroW* and SPoC features differ in their spatial pooling, spatial weighting, and channel weighting. For the first spatially-local pooling step, *CroW* (and *uCroW*) max-pool (we are essentially using the outputs of the last pooling layer of the deep convolutional network rather than the last convolutional one as in SPoC). SPoC uses a centering prior for spatial weighting to boost features that occur near the center of the image, whereas we propose a spatial weighting derived from the spatial activations of the layer outputs themselves. Lastly, SPoC uses a uniform channel weighting, whereas we propose a channel weighting derived from channel sparsity. We demonstrate improvements for each of these design choices in Sect. 5.

**Table 1.** The pooling and weighting steps for three instantiations of our aggregation framework, *i.e.*, the proposed *CroW*, the simplified *uCroW* and SPoC [4]. SW refers to the spatial weighting presented in 4.1, while SSW to the sparsity sensitive channel weighting presented in Sect. 4.2.

Step	SPoC [4]	<i>uCroW</i>	<i>CroW</i>
1: local pooling	None	Max	Max
2: spatial weighting	Centering prior	Uniform	<i>SW</i>
3: channel weighting	Uniform	Uniform	<i>SSW</i>
4: aggregation	Sum	Sum	Sum

## 5 Experiments

### 5.1 Evaluation Protocol

**Datasets.** We experiment on four publicly available datasets. For image search we report results on *Oxford* [26] and *Paris* [27], further combining them with the *Oxford100k* [26] dataset as distractors. We also present results on the *Holidays* [14] dataset. For Oxford we used the common protocol as in all other methods reported, *i.e.* the cropped queries. We regard the cropped region as input for the CNN and extract features. For Holidays we use the “upright” version of the images.

**Evaluation Metrics.** For image search experiments on *Oxford*, *Paris* and *Holidays* we measure mean average precision (mAP) over all queries. We use the evaluation code provided by the authors. For deep neural networks we use Caffe<sup>2</sup> [18] and the publicly available pre-trained VGG16 model [31]. As usual with Caffe, we zero-center the input image by *mean pixel subtraction*. In all cases, table rows including citations present results reported in the cited papers.

**Query Expansion.** One can trivially use simple query expansion techniques [7] with *CroW* features. Given the ranked list of database images by ascending distance to the query, we sum the aggregated feature vectors of the top  $M$  results, L2-normalize and re-query once again. Despite its simplicity, we show that this consistently improves performance, although it does come at the cost of one extra query.

### 5.2 Preliminary Experiments

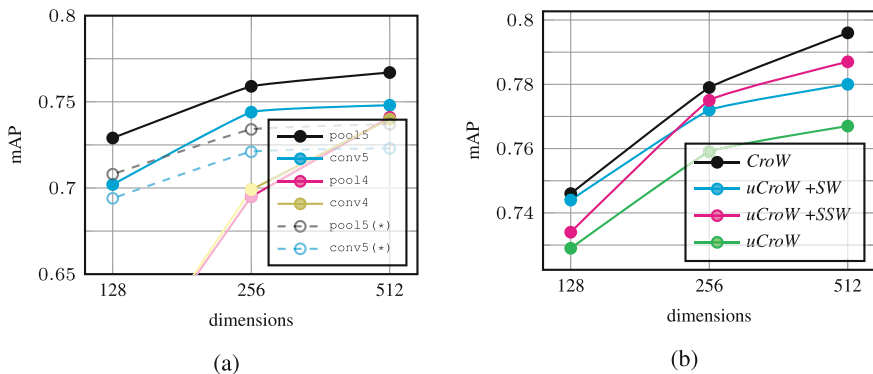
**Image Size and Layer Selection.** In Fig. 5a we investigate the performance of *uCroW* features when aggregating responses from different layers of the network.

Our *uCroW* features are in essence similar to the very recently proposed SPoC features of [4], but have some different design choices that make them more generic and powerful. Firstly, SPoC features are derived from the VGG19 model while our *uCroW* features are derived from the VGG16 model; in this section we show that our *uCroW* features performs much better even if we are using a smaller deep network. Secondly, we do not resize the input image to  $586 \times 586$  as in [4] and instead keep it at its original size. SPoC is therefore comparable to the dotted cyan line in Fig. 5a.

Choosing the last pooling and convolutional layers of the network significantly improves performance over the fourth, especially as the final dimension decreases. Moreover, the `pool5` layer consistently outperforms `conv5-3`, showing that max pooling in Step 1 is indeed beneficial.

Regarding image size, we see that keeping the original size of the images is another factor that contributes to higher performance.

<sup>2</sup> <http://caffe.berkeleyvision.org/>.



**Fig. 5.** Mean average precision on Paris. Different lines denote *uCroW* features from the corresponding layers of VGG16; **conv4** (**conv5**) corresponds to **conv4.3** (**conv5.3**). Solid lines denote that the original image size is kept, while for dashed lines the images were resized to  $586 \times 586$  as in [4]. Both **conv4** and **pool4** layers have very poor performance in low dimensions, with 0.58 mAP for  $d = 128$ . SPoC features [4] correspond to the dotted cyan line. Mean average precision on *Paris* when varying the dimensionality of the final features.

**Effect of the Final Feature Dimensionality.** In Fig. 5b we present mAP on *Paris* when varying the dimensionality of the final features. We present results for all weighting combinations of the proposed approach. *uCroW* refers to uniform or no weighting. *uCroW +SW* refers to using the only the spatial weighting of Sect. 4.1 on top of *uCroW*, *uCroW +SSW* to using the sparsity sensitive channel weighting of Sect. 4.2 on top of *uCroW*, while *CroW* refers to our complete approach with both weighting schemes. As we see, the *uCroW +SSW* combination is affected more by dimensionality reduction than the rest. This can be interpreted as an effect of the subsequent dimensionality reduction. When calculating the sparsity sensitive weights all dimensions are taken into account, however, in the final reduced vector many of those were discarded.

**Notes on Max-pooling.** In preliminary experiments we also tested max-pooling instead of sum pooling for feature aggregation. Consistently with [4] we found it to be always inferior to sum-pooling when whitening was used.

**Table 2.** Mean average precision on Paris when learning the whitening parameters on Oxford, Holidays and Oxford100k for different values of  $d$ .

$d$		Oxford	Holidays	Oxford100k
512	<i>uCroW</i>	0.786	0.752	0.803
	<i>CroW</i>	0.797	0.792	<b>0.810</b>
256	<i>uCroW</i>	0.739	0.728	0.732
	<i>CroW</i>	<b>0.765</b>	0.784	0.762

Interestingly, max pooling performs better than sum-pooling in the non-whitened space, but mAP without whitening the features is much inferior (sometimes more than 10% less) in all datasets tested.

**Whitening.** We learn the whitening parameters from a separate set of images. In Table 2 we present results on the *Paris* dataset when using 3 other datasets for whitening: the semantically related *Oxford* dataset, the *Holidays* dataset and the larger *Oxford100k* dataset. As we reduce the dimensionality, we see overfitting effects for the case where we learn on *Oxford* and this comes as no surprise: as dimensions are reduced, more dimensions that are selective for buildings are

**Table 3.** Mean average precision on *Paris*, *Oxford* and *Holidays* against the state-of-the-art for different values of  $d$ . QE denotes query expansion with the top  $M = 10$  results. The fourth (sixth) column presents results when augmenting the *Paris* (*Oxford*) dataset with the 100k distractors from *Oxford100k*. Results in the lowest set of rows correspond to methods with local features, followed by spatial verification.

Method	d	Paris	+Oxf100k	Oxford	+Oxf100k	Holidays
Tr. Embedding [17]	1024	—	—	0.560	0.502	0.720
Tr. Embedding [17]	512	—	—	—	—	0.700
Gong <i>et al.</i> [5]	512	—	—	—	—	0.783
Neural Codes [5]	512	—	—	0.435	0.392	—
R-MAC [37]	512	<b>0.830</b>	<b>0.757</b>	0.669	0.616	—
<i>uCroW</i>	512	0.786	0.710	0.697	0.641	0.839
<i>CroW</i>	512	0.797	0.722	<b>0.708</b>	<b>0.653</b>	<b>0.851</b>
Tr. Embedding [17]	256	—	—	—	—	0.657
Neural Codes [5]	256	—	—	0.435	0.392	0.749
Razavian <i>et al.</i> [30]	256	0.670	—	0.533	0.489	0.716
SPoC [4]	256	—	—	0.531	0.501	0.802
R-MAC [37]	256	0.729	0.601	0.561	0.470	—
<i>uCroW</i>	256	0.739	0.658	0.667	0.612	0.815
<i>CroW</i>	256	<b>0.765</b>	<b>0.691</b>	<b>0.684</b>	<b>0.637</b>	<b>0.851</b>
Tr. Embedding [17]	128	—	—	0.433	0.353	—
Neural Codes [5]	128	—	—	0.433	0.384	—
<i>uCroW</i>	128	0.699	0.610	0.625	0.559	—
<i>CroW</i>	128	<b>0.746</b>	<b>0.670</b>	<b>0.641</b>	<b>0.590</b>	<b>0.828</b>
<i>CroW</i> + QE	128	0.793	0.728	0.670	0.641	—
<i>CroW</i> + QE	256	0.815	0.753	0.718	0.676	—
<i>CroW</i> + QE	512	<b>0.848</b>	<b>0.794</b>	<b>0.749</b>	<b>0.706</b>	—
Tolias <i>et al.</i> [34]	—	0.770	—	0.804	0.750	—
Total Recall II [6]	—	0.805	0.710	0.827	0.767	—
Mikulik <i>et al.</i> [22]	—	0.824	0.773	0.849	0.795	—

kept when we learn the reduction parameters on a semantically similar dataset like *Oxford*.

To be directly comparable with related works, we learn the whitening parameters on *Oxford* when testing on *Paris* and vice versa, as accustomed. We use the *Oxford100k* dataset for whitening on the *Holidays*.

### 5.3 Image Search

In Table 3 we present comparisons of our approach with the state-of-the-art in image search on *Paris*, *Oxford* and *Holidays*. Both *uCroW* and *CroW* consistently outperform all other aggregation methods for different representation sizes, apart from R-MAC [37], which exhibits very high performance for Paris in 512 dimensions.

*uCroW* is a very strong baseline that gives state-of-the-art performance by itself. It therefore makes sense that improvement over *uCroW* is hard to get. *CroW* improves performance in all cases, with the gain increasing as the dimensionality of the final features decreases. For comparison, if we apply our weighting (instead of the centering prior) to SPoC features, the gain on *Paris* is around 3.9% and 4.6% for 256 and 512 dimensions respectively.

In Fig. 6 we present some interesting results *using just  $d = 32$  dimensional features*. They demonstrate the invariance of *CroW* features to viewpoint and lighting variations even after heavy compression.

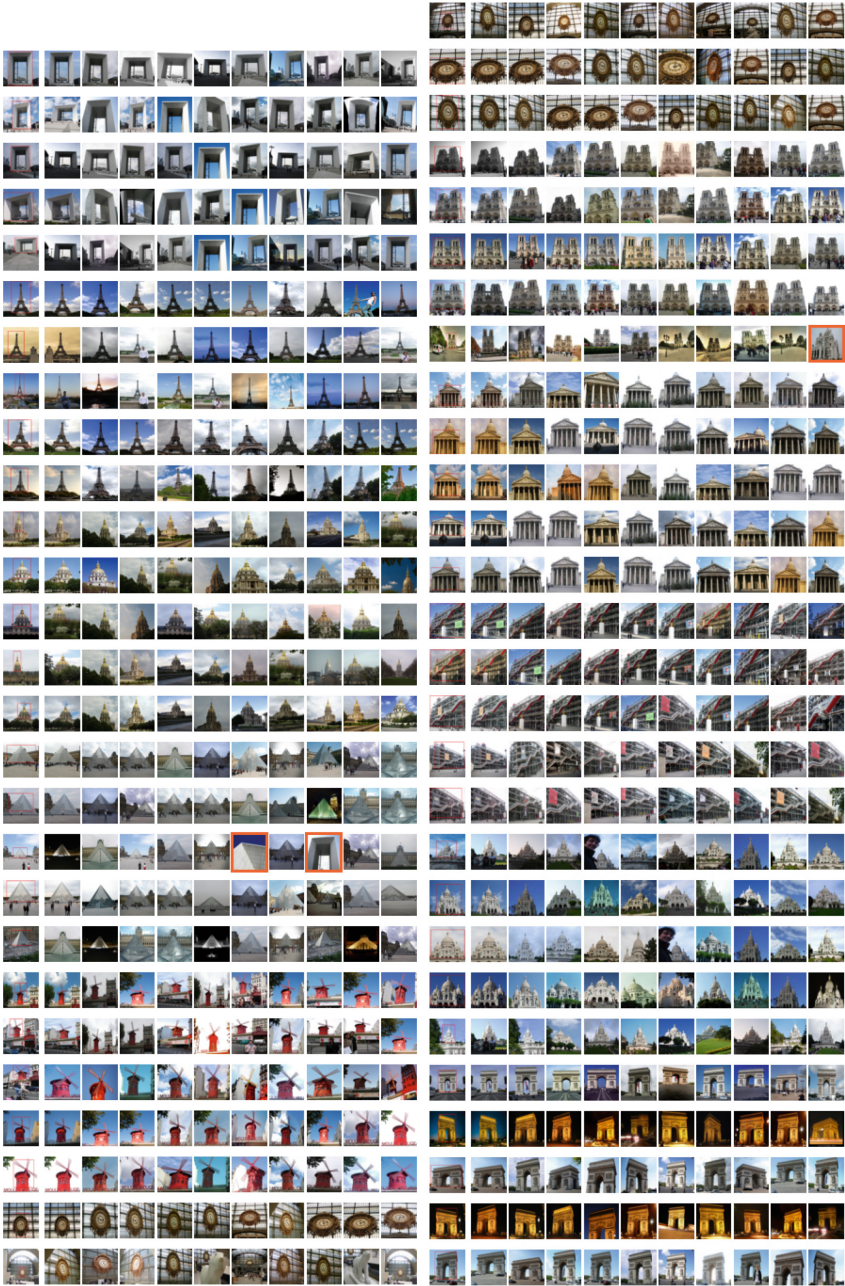
When further combining our approach with query expansion, we get even better results that compare to (or surpass on *Paris*) far more sophisticated approaches like [6, 22, 34] that are based on local features and include spatial verification steps.

In Fig. 7 we show the top-10 results for *all* 55 queries on the paris dataset using the uncompressed *CroW* features ( $d = 512$ ). **We only have 3 false results in total** for precision@10. This illuminates why query expansion is so effective: the top ranked results are already of high quality.



**Fig. 6.** Sample search results using *CroW* features compressed to just  $d = 32$  dimensions. The query image is shown at the leftmost side with the query bounding box marked in a red rectangle. (Color figure online)





**Fig. 7.** Top-10 results returned for all 55 queries of the Paris dataset, using the 512-dimensional *CroW* features (and no query expansion). The query image is shown on the leftmost place, with the query bounding box marked with a red rectangle. Our features produce just 3 false results in total, which are marked with an orange border. (Color figure online)

Although our approach is consistently better, the performance gap between *CroW* and the state-of-the-art is smaller in *Holidays*, where it outperforms the best competing method by about 4.9% and 1.2% for  $d = 256, 512$ , respectively.

## 6 Conclusions

In this paper we outline a generalized framework for aggregated deep convolutional features with cross-dimensional weighting which encompasses recent related works such as [4]. We propose simple, non-parametric weighting schemes for spatial- and channel-wise weighting and provide insights for their behavior by visualizing and studying the distributional properties of the layer output responses. Using this approach, we report results that outperform the state-of-the-art in popular image search benchmarks.

The *CroW* features are one instantiation of our generic aggregation framework. Still, it gives the current state-of-the-art results in image retrieval with minimal overhead and has intuitive qualities that offer insights on the nature of convolutional layer features.

Our aggregation framework is a valuable scaffold within which to discuss and explore new weighting schemes. The framework gave us a clear way to investigate channel and spatial weights independently. Learning weights for a particular task is a promising future direction. Likewise, with sufficient ground truth data, it is possible to fine-tune the entire end-to-end process within our proposed framework using, say, a rank-based loss as in [11,28], together with *attentional mechanisms* and spatial deformations [13].

## References

1. Arandjelovic, R., Zisserman, A.: Three things everyone should know to improve object retrieval. In: CVPR (2012)
2. Avrithis, Y., Tolias, G.: Hough pyramid matching: speeded-up geometry re-ranking for large scale image retrieval. *Int. J. Comput. Vision (IJCV)*, 1–19 (2013)
3. Azizpour, H., Razavian, A.S., Sullivan, J., Maki, A., Carlsson, S.: From generic to specific deep representations for visual recognition. In: DeepVision Workshop, CVPR (2015)
4. Babenko, A., Lempitsky, V.: Aggregating deep convolutional features for image retrieval. In: ICCV (2015)
5. Babenko, A., Slesarev, A., Chigorin, A., Lempitsky, V.: Neural codes for image retrieval. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8689, pp. 584–599. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-10590-1\\_38](https://doi.org/10.1007/978-3-319-10590-1_38)
6. Chum, O., Mikulik, A., Perdoch, M., Matas, J.: Total recall II: query expansion revisited. In: CVPR (2011)
7. Chum, O., Philbin, J., Sivic, J., Isard, M., Zisserman, A.: Total recall: automatic query expansion with a generative feature model for object retrieval. In: ICCV (2007)
8. Cimpoi, M., Maji, S., Vedaldi, A.: Deep filter banks for texture recognition and segmentation. In: CVPR (2015)



9. Delhumeau, J., Gosselin, P., Jegou, H., Perez, P.: Revisiting the VLAD image representation. In: ACM Multimedia (2013)
10. Gong, Y., Wang, L., Guo, R., Lazebnik, S.: Multi-scale orderless pooling of deep convolutional activation features. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8695, pp. 392–407. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-10584-0\\_26](https://doi.org/10.1007/978-3-319-10584-0_26)
11. Gordo, A., Almazán, J., Revaud, J., Larlus, D.: Deep image retrieval: learning global representations for image search. In: ECCV (2016)
12. Gosselin, P.H., Murray, N., Jégou, H., Perronnin, F.: Revisiting the fisher vector for fine-grained classification. *Pattern Recogn. Lett.* **49**, 92–98 (2014)
13. Jaderberg, M., Simonyan, K., Zisserman, A., Kavukcuoglu, K.: Spatial transformer networks. In: NIPS (2015)
14. 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. LNCS, vol. 5302, pp. 304–317. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-88682-2\\_24](https://doi.org/10.1007/978-3-540-88682-2_24)
15. Jégou, H., Douze, M., Schmid, C.: On the burstiness of visual elements. In: CVPR (2009)
16. Jégou, H., Douze, M., Schmid, C., Perez, P.: Aggregating local descriptors into a compact image representation. In: CVPR (2010)
17. Jégou, H., Zisserman, A.: Triangulation embedding and democratic aggregation for image search. In: CVPR (2014)
18. Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., Darrell, T.: Caffe: convolutional architecture for fast feature embedding. arXiv preprint [arXiv:1408.5093](https://arxiv.org/abs/1408.5093) (2014)
19. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NIPS (2012)
20. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. arXiv preprint [arXiv:1411.4038](https://arxiv.org/abs/1411.4038) (2014)
21. Lowe, D.: Local feature view clustering for 3D object recognition. In: CVPR (2001)
22. Mikulík, A., Perdoch, M., Chum, O., Matas, J.: Learning a fine vocabulary. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) ECCV 2010. LNCS, vol. 6313, pp. 1–14. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-15558-1\\_1](https://doi.org/10.1007/978-3-642-15558-1_1)
23. Murray, N., Perronnin, F.: Generalized max pooling. In: CVPR (2014)
24. Perronnin, F., Liu, Y., Sanchez, J., Poirier, H.: Large-scale image retrieval with compressed Fisher vectors. In: CVPR (2010)
25. 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. LNCS, vol. 6314, pp. 143–156. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-15561-1\\_11](https://doi.org/10.1007/978-3-642-15561-1_11)
26. Philbin, J., Chum, O., Isard, M., Sivic, J., Zisserman, A.: Object retrieval with large vocabularies and fast spatial matching. In: CVPR (2007)
27. Philbin, J., Chum, O., Sivic, J., Isard, M., Zisserman, A.: Lost in quantization: improving particular object retrieval in large scale image databases. In: CVPR (2008)
28. Radenović, F., Tolias, G., Chum, O.: CNN image retrieval learns from BoW: unsupervised fine-tuning with hard examples. In: ECCV (2016)
29. Razavian, A.S., Azizpour, H., Sullivan, J., Carlsson, S.: CNN features off-the-shelf: an astounding baseline for recognition. In: DeepVision Workshop, CVPR (2014)
30. Razavian, A.S., Sullivan, J., Maki, A., Carlsson, S.: Visual instance retrieval with deep convolutional networks. arXiv preprint [arXiv:1412.6574](https://arxiv.org/abs/1412.6574) (2014)

31. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition (2014). CoRR abs/1409.1556
32. Sivic, J., Zisserman, A.: Video Google: a text retrieval approach to object matching in videos. In: ICCV, pp. 1470–1477 (2003)
33. Tolias, G., Avrithis, Y., Jégou, H.: To aggregate or not to aggregate: selective match kernels for image search. In: ICCV (2013)
34. Tolias, G., Avrithis, Y., Jégou, H.: Image search with selective match kernels: aggregation across single and multiple images. *Int. J. Comput. Vision*, 1–15 (2015)
35. Tolias, G., Jégou, H.: Visual query expansion with or without geometry: refining local descriptors by feature aggregation. *Pattern Recogn.* **47**(10), 3466–3476 (2014)
36. Tolias, G., Kalantidis, Y., Avrithis, Y., Kollias, S.: Towards large-scale geometry indexing by feature selection. *Comput. Vis. Image Underst. (CVIU)* **120**, 31–45 (2014)
37. Tolias, G., Sivic, R., Jégou, H.: Particular object retrieval with integral max-pooling of CNN activations. In: ICLR (2016)
38. Turcot, P., Lowe, D.: Better matching with fewer features: the selection of useful features in large database recognition problems. In: ICCV (2009)