

Learning and Calibrating Per-Location Classifiers for Visual Place Recognition

Petr Gronát, Josef Sivic, Guillaume Obozinski & Tomas Pajdla

International Journal of Computer Vision

ISSN 0920-5691
Volume 118
Number 3

Int J Comput Vis (2016) 118:319-336
DOI 10.1007/s11263-015-0878-x



Your article is protected by copyright and all rights are held exclusively by Springer Science +Business Media New York. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your article, please use the accepted manuscript version for posting on your own website. You may further deposit the accepted manuscript version in any repository, provided it is only made publicly available 12 months after official publication or later and provided acknowledgement is given to the original source of publication and a link is inserted to the published article on Springer's website. The link must be accompanied by the following text: "The final publication is available at link.springer.com".

Learning and Calibrating Per-Location Classifiers for Visual Place Recognition

Petr Gronát^{1,3} · Josef Sivic¹ · Guillaume Obozinski² · Tomas Pajdla³

Received: 19 February 2015 / Accepted: 21 December 2015 / Published online: 23 April 2016
© Springer Science+Business Media New York 2016

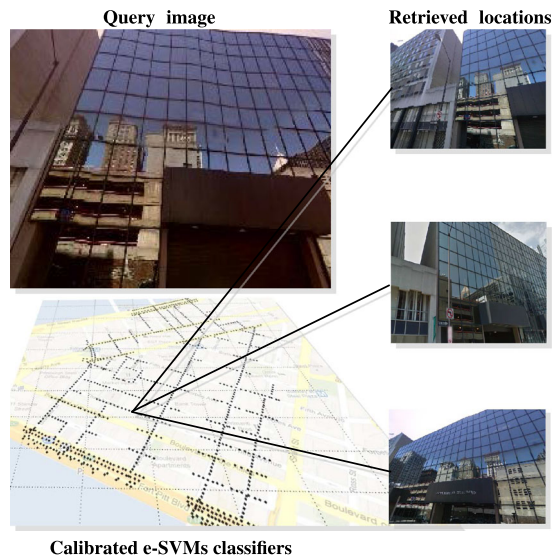


Fig. 1 The goal of this work is to localize a query photograph (*left top*) by finding other images of the same place in a large geotagged image database (*right column*). We cast the problem as a classification task and learn a classifier for each location in the database. We develop two procedures to calibrate the outputs of the large number of per-location classifiers without the need for additional labeled training data

Communicated by Svetlana Lazebnik.

✉ Petr Gronát
petr.gronat@inria.fr

¹ Inria, Willow project, Département d'Informatique de l'Ecole Normale Supérieure, ENS/INRIA/CNRS UMR 8548, 72012 Paris, France

² Ecole des Ponts ParisTech, 6, av. Blaise Pascal, 77455 Marne-la-Vallée, France

³ Czech Technical University in Prague, Faculty of Electrical Engineering, Department of Cybernetics, 12135 Praha 2, Czech Republic

Abstract The aim of this work is to localize a query photograph by finding other images depicting the same place in a large geotagged image database. This is a challenging task due to changes in viewpoint, imaging conditions and the large size of the image database. The contribution of this work is two-fold. First, we cast the place recognition problem as a classification task and use the available geotags to train a classifier for each location in the database in a similar manner to per-exemplar SVMs in object recognition. Second, as only one or a few positive training examples are available for each location, we propose two methods to calibrate all the per-location SVM classifiers without the need for additional positive training data. The first method relies on p-values from statistical hypothesis testing and uses only the available negative training data. The second method performs an affine calibration by appropriately normalizing the learnt classifier hyperplane and does not need any additional labelled training data. We test the proposed place recognition method with the bag-of-visual-words and Fisher vector image representations suitable for large scale indexing. Experiments are performed on three datasets: 25,000 and 55,000 geotagged street view images of Pittsburgh, and the 24/7 Tokyo benchmark containing 76,000 images with varying illumination conditions. The results show improved place recognition accuracy of the learnt image representation over direct matching of raw image descriptors.

Keywords Place recognition · Exemplar SVM · Geo-localization · Classifier calibration

1 Introduction

Visual place recognition (Cummins and Newman 2009; Knopp et al. 2010; Schindler et al. 2007) is a challenging task

as the query and database images may depict the same 3D structure (e.g. a building) from a different camera viewpoint, under different illumination, or the building can be partially occluded. In addition, the geotagged database may be very large. For example, we estimate that Google Street View of France alone contains more than 60 million panoramic images. It is, however, an important problem as automatic, accurate and fast visual place recognition would have many practical applications in robotics, augmented reality or navigation.

Similar to other work in large scale place recognition (Cummins and Newman 2009; Knopp et al. 2010; Schindler et al. 2007; Torii et al. 2013) and image retrieval (Nister and Stewenius 2006; Philbin et al. 2007; Sivic and Zisserman 2003; Jégou et al. 2012), we describe each image by a set of local invariant features (Bay et al. 2006; Lowe (2004)) that are encoded and aggregated into a fixed-length single vector descriptor for each image. In particular, in this work we consider the sparse tf-idf weighted bag-of-visual-words representation (Sivic and Zisserman 2003; Philbin et al. 2007) and the compact Fisher vector descriptors (Jégou et al. 2012).

The resulting vectors are then normalized to have unit L_2 norm and the similarity between the query and a database vector is measured by their dot product. This representation has some desirable properties such as robustness to background clutter and partial occlusion. Efficient retrieval can then be achieved using inverted file indexing (Jégou et al. 2011).

While in image retrieval databases are typically unstructured collections of images, place recognition databases are usually structured: images have geotags, are localized on a map and depict a consistent 3D world. Knowing the structure of the database can lead to significant improvements in both speed and accuracy of place recognition. Examples include: (i) building an explicit 3D reconstruction of the scene (Irschara et al. 2009; Li et al. 2010; Li et al. 2012); (ii) constructing an image graph (Cao and Snavely 2013; Philbin et al. 2010; Turcot and Lowe 2009), where images are nodes and edges connect close-by images on the map (Torii et al. 2011), or (iii) using the geotagged data as a form of supervision to select local features that characterize a certain location (Knopp et al. 2010; Schindler et al. 2007) or re-rank retrieved images (Zamir and Shah 2010).

In this work, we also take advantage of geotags as an available form of supervision and investigate whether the place recognition problem can be cast as a classification task. Learning visual classifiers has been investigated for landmark recognition (Li et al. 2009) where consumer photographs were clustered into landmark classes based on geo-tags. In this work we wish to recognize individual street locations rather than a small number of landmarks, and as a consequence have only a few (1–5) photographs capturing the same location. In particular, we train a classifier for each location on the map in a similar manner to per-exemplar

classification in object recognition (Malisiewicz et al. 2011). (By *location* we mean a scene captured by a camera, hence using our terminology, a camera heading to the north captures different location than a camera heading to the south, even though these two cameras share the same GPS location. In the following text we use terms *place* and *location* interchangeably.)

This is beneficial as each classifier can learn which features are discriminative for a particular place. The classifiers are learnt offline. At query time, the query photograph is localized by transferring the GPS tag of the best scoring location classifier.

While learning classifiers for each location may be appealing, calibrating outputs of the individual classifiers is a critical issue. In object recognition (Malisiewicz et al. 2011), it is addressed in a separate calibration stage on a held-out set of training data. This is not possible in the place recognition set-up as only a small number, typically one to five, of positive training images are available for each location (e.g. street view images viewing the same building facade). To address this issue, we propose two calibration methods. The first method relies on p-values from statistical hypothesis testing and uses only the available negative training data. The second method performs a simple affine calibration by appropriately normalizing the learnt classifiers and does not need any additional labelled calibration examples.

2 Related Work

The task of geo-localizing a given input query photograph has recently received considerable attention. The output can be a coarse geo-localization on the level of continents and cities (Doersch et al. 2012; Hays and Efros 2008; Kalogerakis et al. 2009) or a name of the depicted landmark (Li et al. 2009). In this work we focus on visually recognizing the “same place” by finding an image in geo-tagged database that depicts the same building facade or street-corner as shown in the query (Chen et al. 2011; Cummins and Newman 2009; Knopp et al. 2010; Schindler et al. 2007; Torii et al. 2013; Zamir and Shah 2010).

This *visual place recognition* problem is typically treated as large-scale instance-level retrieval (Cummins and Newman 2009; Chen et al. 2011; Knopp et al. 2010; Schindler et al. 2007; Torii et al. 2013; Zamir and Shah 2010), where images are represented using local invariant features (Lowe 2004) encoded and aggregated into the bag-of-visual-words (Csurka et al. 2004; Sivic and Zisserman 2003) or Fisher vector (Jégou et al. 2012) representations. The image database can be further augmented by 3D point clouds (Klingner et al. 2013), automatically reconstructed by large-scale structure from motion (SfM) (Agarwal et al.

2009; Klingner et al. 2013), which enables accurate prediction of query image camera position (Li et al. 2012; Sattler et al. 2012). In contrast, in this work we investigate learning a discriminative place-specific image representation. A similar idea has been recently explored in (Cao and Snavely 2013) who learn a graph-based discriminative representation for landmark image collections where typically many images are available for each landmark. In this work, we focus on street-level images such as Google Street View, which have greater coverage, but typically only one or a small number of images are available for each location. To address this issue we learn a discriminative re-weighting of the descriptor specific to each image in the database using per-exemplar support vector machine (Malisiewicz et al. 2011).

2.1 Per-Exemplar Support Vector Machines

The exemplar support vector machines (e-SVM) has been used in a number of visual recognition tasks including category-level recognition (Malisiewicz et al. 2011), cross-domain retrieval (Shrivastava et al. 2011), scene parsing (Tighe and Lazebnik 2013) or as an initialization for more complex discriminative clustering models (Doersch et al. 2012; Singh et al. 2012). The main idea is to train a linear support vector machine (SVM) classifier from a single positive example (or very few positive examples) and a large number of negatives. The intuition is that the resulting weight vector will give a higher weight to the discriminative dimensions of the positive training data point and will down weight dimensions that are non-discriminative with respect to the negative training data.

The exemplar support vector machine can be learnt at query time where the weight vector is used as a new query image representation (Shrivastava et al. 2011). However, this requires training a new classifier afresh for each query that is computationally very demanding. In this work, similar to (Malisiewicz et al. 2011) who learn per-exemplar object category representation, we learn per-exemplar classifiers for each place in the database off-line. A key advantage is that each per-exemplar classifier is trained independently and hence the learning can be heavily parallelized. The per-exemplar training brings, however, also an important drawback. As each classifier is trained independently a careful calibration of the resulting classifier scores is required (Malisiewicz et al. 2011).

2.2 Calibrating Classifier Scores

Several calibration approaches have been proposed in the literature (see Gebel and Weihs (2007) and references therein for a review). The most known consists of fitting a logistic regression to the output of the SVM (Platt 1999). This approach, however, has a major drawback as it imposes a

parametric form (the logistic a.k.a. sigmoid function) of the likelihood ratio of the two classes, which typically leads to biased estimates of the calibrated scores. Another important calibration method is the isotonic regression (Zadrozny and Elkan 2002), which allows for a non-parametric estimate of the output probability. Unfortunately, the fact that we have only a single positive example (or only very few of them which are almost identical, and which are all used for training) essentially prevents us from using any of these methods. To address these issues, we develop two classifier calibration methods that do not need additional labelled positive examples. Related to ours is also the recent work of Scheirer et al. (Scheirer et al. 2012) who develop a classifier calibration method for face attribute similarity search. Their method (discussed in more detail in Sect. 4) also does not require labelled positive examples but, in contrast to us, uses a parametric model (the Weibull distribution) for the scores of negative examples.

2.3 Linear Discriminant Analysis and Whitening

Our work is also related to linear discriminative transformations of feature space that have shown good performance in object recognition (Gharbi et al. 2012; Hariharan et al. 2012) and 2D–3D alignment (Aubry et al. 2014; Aubry et al. 2014). While conceptually the idea of finding a discriminative projection of the original feature space is similar to our work, the main difference is in the used loss function. While we use hinge loss (Scholkopf and Smola 2002) to train the new discriminative representation of each place, (Aubry et al. 2014; Gharbi et al. 2012; Hariharan et al. 2012) use the Euclidean loss. The advantage of using the Euclidean loss is that the discriminative projection can be computed in closed form. The resulting projection is tightly related to Linear Discriminant Analysis and whitening the feature space (Aubry et al. 2014; Gharbi et al. 2012; Hariharan et al. 2012). Such whitened representations have shown promise for image retrieval (Jégou and Chum 2012) or matching HOG (Dalal and Triggs 2005) descriptors (Doersch et al. 2013), however, we have found they do not perform well for place recognition.

2.4 Contributions

This paper has two main contributions. First, we cast the place recognition problem as a classification task where we use the available geo-tags as a weak form of supervision to train a classifier for each location in the database (Sect. 3). These classifiers are subsequently used for ranking the database images at query time.

Second, as only one or a few positive training examples are available for each location, we propose two methods to calibrate all the per-location SVM classifiers without the need for additional positive training data. The first method (Sect. 4)

relies on p-values from statistical hypothesis testing. The second method (Sect. 5) performs an affine calibration by appropriately normalizing the learnt decision hyperplane. We also describe a memory efficient classifier representation for the sparse bag-of-visual-word vectors (Sect. 6) and experimentally demonstrate benefits of the proposed approach (Sects. 7 and 8).

3 Per-Location Classifiers for Place Recognition

We are given an image descriptor \mathbf{x}_j , one for each database image j . This representation can be a sparse tf-idf weighted bag-of-visual-words vector (Sivic and Zisserman 2003) or a dense compact descriptor such as the Fisher vector (FV) (Jégou et al. 2012). The goal is to learn a score function f_j for each database image j , so that, at test time, given the descriptor \mathbf{q} of the query image, we can either retrieve the correct target image as the image j^* with the highest score

$$j^* = \arg \max_j f_j(\mathbf{q}) \tag{1}$$

or use these scores to rank candidate images and use geometric verification to identify the correct location in an n -best list. Instead of approaching the problem directly as a large multiclass classification problem, we tackle the problem by learning a per-exemplar linear SVM classifier (Malisiewicz et al. 2011) for each database image j . Similar to (Knopp et al. 2010), we use the available geotags to construct the negative set \mathcal{N}_j for each image j . The negative set is constructed so as to concentrate difficult negative examples, i.e. from images that are far away from the location of image j and similar to the target image as measured by the dot product between their feature vectors. The details of the construction procedure will be given in Sect. 7. The positive set \mathcal{P}_j is represented by a single positive example, which is \mathbf{x}_j itself. For some locations the positive set can be expanded, details of this procedure will be given in 7.3.

Each SVM classifier produces a classification score s_j which is a priori not comparable with the score of the other classifiers. A calibration of these classification scores will therefore be key to convert them to comparable scores f_j . This calibration problem is more difficult than usual given that we only have a single positive example and will be addressed in Sect. 4.

3.1 Learning Per-Location SVM Classifiers

Each linear SVM classifier generates a classification score s_j of the form

$$s_j(\mathbf{q}) = \mathbf{q}^T \mathbf{w}_j + b_j \tag{2}$$

where \mathbf{w}_j is a weight vector re-weighting contributions of individual visual words and b_j is the bias specific for image j . Given the training sets \mathcal{P}_j and \mathcal{N}_j , the aim is to find a vector \mathbf{w}_j and bias b_j such that the score difference between \mathbf{x}_j and the closest neighbor from its negative set \mathcal{N}_j is maximized. Learning the weight vector \mathbf{w}_j and bias b_j is formulated as a minimization of the convex objective

$$\begin{aligned} \Omega(\mathbf{w}_j, b_j) = & \|\mathbf{w}_j\|^2 + C_1 \sum_{\mathbf{x} \in \mathcal{P}_j} h(\mathbf{w}_j^T \mathbf{x} + b_j) \\ & + C_2 \sum_{\mathbf{x} \in \mathcal{N}_j} h(-\mathbf{w}_j^T \mathbf{x} - b_j), \end{aligned} \tag{3}$$

where the first term is the regularizer, the second term is the loss on the positive training data weighted by scalar parameter C_1 , and the third term is the loss on the negative training data weighted by scalar parameter C_2 . This is a standard SVM formulation (3), also used in exemplar-SVM (Malisiewicz et al. 2011). In our case h is the squared hinge loss, which we found to work better in our setting than the standard hinge-loss. Parameters \mathbf{w}_j and b_j are learned separately for each database image j in turn.

3.2 The Need for Calibrating Classifier Scores

Since the classification scores s_j are learned independently for each location j , they cannot be directly used for place recognition as in Eq. (1). As illustrated in Fig. 2, for a given query \mathbf{q} , a classifier from an incorrect location (b) can have a higher score (Eq. (2)) than the classifier from the target location (a). Indeed, the SVM score is a signed distance from the discriminating hyperplane and is a priori not comparable between different classifiers. This issue is addressed by calibrating scores of the learnt classifiers. The goal of the calibration is to convert the output of each classifier into a probability (or in general a “universal” score), which can be meaningfully compared across classifiers. In the following two sections we develop two classifier calibration methods that do not need additional labelled positive examples.

4 Non-parametric Calibration of the SVM-Scores from Negative Examples Only

In this section we describe a classifier calibration method that exploits the availability of large amounts of negative data, i.e. images from other far away locations in the database. In particular, the method estimates the significance of the classification score of a test example compared to the typical classification score of the (plentifully available) negative examples. Intuitively, we will use a large dataset of negative examples to calibrate the individual classifiers so that they reject the same number of negative examples at each level of

the calibrated score. We will expand this idea in detail using the concepts from hypothesis testing.

4.1 Calibration via Significance Levels

In the following, we view the problem of deciding whether a query image matches a given location based on the corresponding SVM classification score as a hypothesis testing problem. In particular, we appeal to ideas from the traditional frequentist hypothesis testing framework also known as Neyman–Pearson (NP) framework (see e.g. Casella and Berger (2001), chap. 8).

We define the null hypothesis as $H_0 = \{\text{the image is a random image}\}$ and the alternative as $H_1 = \{\text{the image matches the particular location}\}$. The NP framework focuses on the case where the distribution of the data under H_0 is well known, whereas the distribution under H_1 is not accessible or too complicated to model, which matches perfectly our setting.

In the NP framework, the *significance level* of a score is measured by the p-value or equivalently by the value of the cumulative density function (cdf) of the distribution of the negatives at a given score value. The cdf is the function F_0 defined by $F_0(s) = \mathbb{P}(S_0 \leq s)$, where S_0 is a random variable corresponding to the classification scores of negative data (see Fig. 3 for an illustration of the relation between the cdf and the density of the function). The cdf (or the corresponding p value¹) is naturally estimated by the empirical cumulative density function \hat{F}_0 , which is computed as:

$$\hat{F}_0(s) = \frac{1}{N_c} \sum_{n=1}^{N_c} 1_{\{s_n \leq s\}}, \tag{4}$$

where $(s_n)_{1 \leq n \leq N_c}$ are the SVM classification scores associated with N_c negative examples used for calibration. Note that no positive examples are involved in the construction of the cumulative density function. $\hat{F}_0(s)$ is the fraction of the negative examples used for calibration (ideally held out negative examples) that have a score below a given value s . Computing \hat{F}_0 exactly would require to store all the SVM scores for all the calibration data for all classifiers, so in practice, we only keep a fraction of the larger scores. We also interpolate the empirical cdf between consecutive datapoints so that instead of being a staircase function it is a continuous piecewise linear function such as illustrated in

¹ The notion most commonly used in statistics is in fact the p value. The p value associated to a score is the quantity $\alpha(s)$ defined by $\alpha(s) = 1 - F_0(s)$; so the more significant the score is, the closer to 1 the cdf value is, and the closer to 0 the p-value is. To keep the presentation simple, we avoid the formulation in terms of p-values and we only talk of the probabilistic calibrated values obtained from the cdf F_0 .




Fig. 2 An illustration of the proposed normalization of SVM scores for database images. In each plot, the x axis shows the raw SVM score. The y axis shows the calibrated output. For the given query, the raw SVM score of image (b) is lower than for image (a), but the calibrated score of image (b) is higher than for image (a)




Fig. 3 Cumulative density function. Illustration of the relation between **a** the probability density of the random variable S_0 modeling the scores of the negative examples and **b** the corresponding cumulative density function $F_0(s) = \mathbb{P}(S_0 \leq s)$.

Fig. 2. Given a query, we first compute its SVM classification score s_q and then compute the calibrated probability, the score function $f(q) = \hat{F}_0(s_q)$. We obtain a similar calibrated score function $f_j(q)$ for each of the SVMs associated with each of the target locations, which can now be ranked. Two other examples of score calibration functions are shown in

Fig. 6 in Sect. 8. Note that while Fig. 2 illustrates only few points on the cdf, the two plots in Fig. 6 show a complete cdf that contains on the order of 25k data points. Note also that the two cumulative density functions in Fig. 6 are similar but not identical.

4.2 Summary of the calibration procedure

For each trained place-specific classifier s_j we construct the empirical cumulative density function (4) of classification scores of the negative examples and keep only its top K values. This can be done offline and the procedure is summarized in Algorithm 1. At query time, given a query image descriptor \mathbf{q} , we compute the uncalibrated classification score $s_j(\mathbf{q})$ and then use the stored cdf values to compute the calibrated score $f_j(\mathbf{q})$. This procedure is performed for each database image j and is summarized in Algorithm 2. Finally, the best candidate database image is selected by equation (1). Alternatively, candidate database images can be also ranked according to the calibrated score.

Algorithm 1 P-value calibration: offline stage

Require: \mathcal{X} ...column wise matrix of image descriptors
 \mathbf{w}_j, b_j ...learnt SVM weights and biases
Ensure: \hat{F}_{0j} ...calibration functions
1: **procedure** P-VALUE CALIBRATION
2: $N \leftarrow$ database size
3: $\mathcal{X} \leftarrow$ descriptor matrix of negative examples
4: **for** $\forall j \in 1 \dots N$ **do**
5: $N_c \leftarrow$ number of negative examples
6: $\mathbf{w} \leftarrow$ learned SVM weight for image j
7: $b \leftarrow$ learned SVM bias for image j
8: $\sigma \leftarrow \mathbf{w}^T \mathcal{X} + b$
9: *Compute the cdf:*
10: $\mathbf{s}_j \leftarrow$ sorted σ in descending order
11: $\hat{F}_{0j} \leftarrow [N_c \dots 0]/N_c$

Algorithm 2 P-value calibration: online stage

Require: \mathbf{q} ...query image descriptor
 \mathbf{w}_j, b_j ...learnt SVM weights and biases
 \hat{F}_{0j} ...learnt calibration function
Ensure: $f_j(\mathbf{q})$...calibrated score
1: **procedure** CALIBRATING SCORES
2: $\mathbf{q} \leftarrow$ query image descriptor
3: $N \leftarrow$ database size
4: **for** $\forall j \in 1 \dots N$ **do** // for each database image
5: $\mathbf{w} \leftarrow$ learned SVM weight for image j
6: $b \leftarrow$ learned SVM bias for image j
7: $\hat{F}_0 \leftarrow \hat{F}_{0j}$ // Empirical cdf
8: $\mathbf{s} \leftarrow \mathbf{s}_j$ // Corresponding sorted scores
9: $s_q \leftarrow \mathbf{q}^T \mathbf{w} + b$ // compute uncalibrated classifier score
10: Find n such that $s_n \leq s_q < s_{n+1}$
11: Compute the interpolated empirical cdf value:
 $\hat{F}_0(s_q) \approx \hat{F}_0(s_n) + \frac{s_q - s_n}{s_{n+1} - s_n} (\hat{F}_0(s_{n+1}) - \hat{F}_0(s_n))$.
12: $f_j(\mathbf{q}) = \hat{F}_0(s_q)$ // output the calibrated score

4.3 Discussion

It should be noted that basing the calibration only on the negative data has the advantage that we privilege precision over recall, which is justified given the imbalance of the available training data (many more negatives than positives). Indeed, since we are learning with a single positive example (or a very few), intuitively, we cannot guarantee that the learned partition of the space will generalize well to other positives, whose scores in the test set can potentially drop significantly. By contrast, since we are learning from a comparatively large number of negative examples, we can trust the fact that new negative examples will stay in the half-space containing the negative training set, so that their scores are very unlikely to be large. Our method is therefore based on the fact that we can measure reliably how surprising a high classification score would be if it was the score of a negative example. This exactly means that we can control false positives (type I error) reasonably well but not false negatives (type II error or equivalently the power of our test/classifier), exactly as in the Neyman-Pearson framework.

An additional reason for not relying on positive examples for the calibration in our case is that (even if we had sufficiently many of them) the positive examples that we collect using location and geometric verification from the geotagged database (details in Sect. 7.3) typically have illumination conditions that are extremely similar to each other and not representative of the distribution of test positives which can have very different illuminations. This is because of the controlled nature of the capturing process of geotagged street-level imagery (e.g. Google Street View) used for experiments in this work. Close-by images are typically captured at a similar time (e.g. on the same day) and under similar imaging conditions.

Scheirer et al. (2012) propose a method, which is related to ours, and calibrate SVM scores by computing the corresponding cdf value of a Weibull distribution fitted to the top negative scores. The main difficulty is that the Weibull model should be fitted only to the tail of the distribution of the negatives, which is in general difficult to identify. As a heuristic, Scheirer et al. propose to fit the Weibull model to false positives (i.e. the negative samples classified incorrectly as positives). But in our case, most of the exemplar SVMs that we are training have zero false positives in a held out set, which precludes the application of their method.

Finally, we should remark that we are not doing here calibration in the same sense of the word as the calibration based on logistic regression (or isotonic regression), since logistic regression estimates the probability of making a correct prediction by assigning a new data to class 1, while we are estimating how unlikely it would be for a negative example to have such a high score. The calibration with either methods yields “universal” scores in the sense that they are compa-

rable from one SVM to another, but the calibrated values obtained from logistic regression are not comparable to the values obtained from our approach.

5 Affine Calibration by Normalizing the Classification Hyperplane

The non-parametric calibration method described in the previous section has two computational disadvantages, which make it hard to scale-up to very large datasets. First, the method requires storing the non-parametric model of the calibration function for each learned classifier. This has memory complexity of $O(NK)$, where N is the number of images (classifiers) in the database and K the number of stored elements of the non-parametric model. For typical values of $K = 1000$ and $N = 1M$ this would require additional 4GB of memory, comparable to the size of the inverted index itself. Second, computing the cumulative density function requires applying all N learnt classifiers to the entire set of negative examples, which has also size N . As a result computing the cdf has complexity $O(N^2)$, which becomes quickly infeasible already for datasets with N larger than 100,000.

To address these issues we first describe an affine calibration model that calibrates the classifier score with a simple linear function defined by only two parameters: its slope and offset, greatly reducing the required storage. Second, we show that the parameters of the affine calibration function can be obtained by normalizing the learnt classification hyperplane without applying the classifiers on the negative data and thus bringing down the computational complexity to $O(N)$. As a result, computing and storing the calibration functions becomes feasible for very large datasets with 1M images.

5.1 Affine Calibration Model

Using the affine calibration model we transform the uncalibrated classification score $s_j(\mathbf{q})$ of query \mathbf{q} with a linear function

$$f_j(\mathbf{q}) = \alpha_j s_j(\mathbf{q}) + \beta_j, \tag{5}$$

where $f_j(\mathbf{q})$ is the output calibrated score, and α_j and β_j are scalar calibration parameters specific to each classifier j . In this work we use linear classifiers, hence substituting for $s_j(\mathbf{q})$ the linear classifier from (2) results also in a linear calibrated classifier

$$f_j(\mathbf{q}) = \tilde{\mathbf{w}}_j^T \mathbf{q} + \tilde{b}_j, \tag{6}$$

where $\tilde{\mathbf{w}}_j = \alpha_j \mathbf{w}_j$ and $\tilde{b}_j = \alpha_j b_j + \beta_j$. Note that the calibrated classifier (6) has the same form as the original classifier (2) and hence this representation does not require

any additional storage compared to storing the original classifier. The question remains how to set the parameters α_j and β_j of the calibration function (5), which is discussed next.

5.2 Calibration by Normalization

Parameters of the affine calibration function (5) could be learnt from negative training data in a similar manner to, for example, (Aubry et al. 2014). We have tried to estimate the parameters in a similar manner by fitting a line to the tail of the cdf, however this procedure did not yield satisfactory results. In addition, as discussed above, in our case this requires running all N classifiers on all N images, which is prohibitive for large datasets. Instead, we have found that a good calibration can be obtained by normalizing the learnt hyperplane \mathbf{w} . In particular, we set

$$\alpha_j = \frac{1}{\|\mathbf{w}_j\|}, \tag{7}$$

$$\beta_j = -b_j \alpha_j, \tag{8}$$

where \mathbf{w}_j and b_j are the parameters of the learnt SVM hyperplane for location j and $\|\mathbf{w}\|$ is the L_2 norm of \mathbf{w} .

Given this choice of α_j and β_j the bias term in Eq. (6) cancels out and the calibrated classification score (6) reduces to

$$f_j(\mathbf{q}) = \frac{1}{\|\mathbf{w}_j\|} \mathbf{w}_j^T \mathbf{q} = \tilde{\mathbf{w}}_j^T \mathbf{q}. \tag{9}$$

Since \mathbf{q} is L_2 normalized, the outcome of the particular choice of α_j and β_j is that the Eq. (9) is equivalent to computing the normalized dot-product between vectors \mathbf{q} and \mathbf{w} . This was found to work well in image retrieval (Sivic and Zisserman 2003) or matching whitened HOG descriptors (Doersch et al. 2013). In this work we investigate whether this intuition about descriptor matching can be used as a form of calibration for the learnt place-specific classifier.

Note that this form of calibration by normalization is scalable to very large datasets as it (i) requires only $O(N)$ computations offline to pre-compute the calibration parameters for each of the N learnt classifiers (Eqs. (7) and (8)), and (ii) does not need any additional storage or computation at query time as the calibration parameters can be included in the classifier (6). In Appendix we examine the per-exemplar SVM cost and give an additional intuition why calibration by re-normalization works.

6 Memory Efficient Classifier Representation

We learn a (calibrated) linear discriminative classifier with weight vector \mathbf{w}_j and bias b_j for each image j in the

database.² These classifier parameters become the new representation for each image. In this section we discuss how the classifier parameters can be stored in a memory efficient manner that is amenable for indexing. The goal is to apply all the learnt classifiers to the query descriptor \mathbf{q}

$$\mathbf{s} = \mathbf{q}^T \mathcal{W} + \mathbf{b}, \quad (10)$$

where \mathcal{W} is $d \times N$ matrix storing all the learnt \mathbf{w}_j classifiers as columns, \mathbf{b} is a $1 \times N$ vector storing all the learnt bias values b_j , \mathbf{q} is the input query descriptor, \mathbf{s} is a $1 \times N$ vector of output scores for all classifiers in the database, N is the number of images in the database and d is the dimensionality of the image representation. As discussed in detail in Sect. 7 we investigate two different image representations: (i) the compact Fisher vectors (Jégou et al. 2012) and (ii) the bag-of-visual-word vectors (Sivic and Zisserman 2003). The learnt classifiers for these two image descriptors have different statistics and require different methods for storing and indexing. Next, we discuss the classifier representations for the two types of image representations.

6.1 Fisher Vectors

The Fisher vector descriptors are not sparse, but have a relatively low-dimension $d \in \{128, 512, 2048\}$ hence it is possible to store directly the (non-sparse) matrix \mathcal{W} containing the learnt classifier parameters \mathbf{w} . In this work we exhaustively compute the classifier scores for all images in the database (given by Eq. (10)) using efficient (but exact) matrix-vector multiplication routines. However, this computation can be further sped-up using product quantization indexing as described in Jégou et al. (2011).

6.2 Bag-of-Visual-Words

In the bag-of-visual-words representation, each image is represented by a high dimensional vector \mathbf{x} , where the dimensionality d is typically 100,000, but the vector is very sparse with only about 2000 non-zero entries. The learnt \mathbf{w}_j are of the same (high) dimension d but are *not sparse*. As a result, directly storing the learnt classifiers becomes quickly infeasible. To illustrate this, consider a database of $N = 1,000,000$ images. Storing the original descriptors with about 2,000 non-zero entries for each image would take around 8GB. However, directly storing the learnt non-sparse $100,000 \times 1,000,000$ matrix \mathcal{W} would require 400GB of memory. To address this issue we have developed an alternative indexing structure taking advantage of the dual form of

the linear classifier as a sparse linear combination of a small number of support vectors (Scholkopf and Smola 2002). The key observation is that the number of support vectors k is significantly lower than dimensionality d of the original image descriptor. In the following we omit index j for clarity. In detail, we represent each \mathbf{w} by its corresponding coefficients α_i of the linear combination of the support vectors (individual image descriptors) \mathbf{x}_i such that

$$\mathbf{w} = \sum_i \alpha_i \mathbf{x}_i = \mathcal{X} \cdot \boldsymbol{\alpha}, \quad (11)$$

where α_i , the elements of vector $\boldsymbol{\alpha}$, are coefficients of the linear combination of the training data points \mathbf{x}_i and the matrix \mathcal{X} contains (as columns) descriptors of the entire database. Note that the vector $\boldsymbol{\alpha}$ is sparse and the number of non-zero elements depends on the number of support vectors k .

As a result, matrix \mathcal{W} containing all learned classifier weights can be expressed in the dual form as

$$\mathcal{W} = \mathcal{X} \mathcal{A}, \quad (12)$$

where \mathcal{X} is the (sparse) matrix of the bag-of-visual-words image descriptors and \mathcal{A} is the (sparse) matrix of α coefficients, where each column corresponds to vector $\boldsymbol{\alpha}$ from (11). Instead of storing all (non-sparse) weight vectors \mathcal{W} , which has memory complexity $O(dN)$ where $d (=100,000)$ is the dimensionality of the image representation and N is the size of the database, we store two sparse matrices \mathcal{X} and \mathcal{A} , which has memory complexity $O(mN + kN)$ where $m (=2,000)$ is the number on non-zero elements in the original bag-of-visual-word descriptors, and k is the typical number of support vectors. In our case k is about the size of the training data which is around 500. As a result, the storage requirements are significantly reduced. For example, for a database of 1M images the dual representation requires only about 10 GB of storage compared to 400GB for directly storing classifiers \mathcal{W} . Note that sparsity can be imposed directly on the learnt classifiers \mathbf{w} by appropriate regularization (Scholkopf and Smola 2002). However, we found this approach did not yield competitive results in terms of accuracy.

7 Experimental Setup and Implementation Details

In this section we describe the experimental datasets, outline the two types of used image descriptors, and finally give implementation details of the classifier learning procedure.

7.1 Image datasets

Experiments are performed on two datasets, the Pittsburgh place recognition dataset (Gronát et al. 2013) and the Tokyo 24/7 dataset (Torii et al. 2015).

² When the calibration by re-normalization method is used the \mathbf{w}_j contains the re-normalized weights and the bias b_j is zero. However, to cover both calibration methods we include the bias term in the derivations in this section.

7.1.1 Pittsburgh Dataset

The first dataset contains Google Street View panoramas downloaded Gronát (2015) from the Internet covering an area of $1.3 \times 1.2 \text{ km}^2$ of the city of Pittsburgh (U.S.). Similar to Chen et al. (2011), we generate for each panorama 12 overlapping perspective views corresponding to two different elevation angles 4° and 28° to capture both the street-level scene and the building façades. This results in a total of 24 perspective views each with 90° FOV and resolution of 960×720 pixels. In this manner we generate two versions of this dataset. The first version covers a smaller area and contains 25k perspective images. The second larger dataset contains 55k images.

As a query set with known ground truth GPS positions, we use 8999 panoramas from the Google Street View research dataset (Google: ICMLA 2011 <http://www.icmla-conference.org/icmla11/challenge.htm>), which cover approximately the same area, but were captured at a different time. The query images typically depict the same places from different viewpoints and under different illumination conditions. We generate a test query set such that we first select a panorama at random, and second, we generate a perspective image with a random orientation and random elevation pitch. This way we synthesize 4,000 query test images. Both the query and database images are available upon request at Gronát <http://www.di.ens.fr/willow/research/perlocation/>.

7.1.2 24/7 Tokyo dataset

The 24/7 Tokyo dataset (Torii et al. 2015) contains Google Street View panoramas downloaded from the Internet covering an area of $1.6 \times 1.6 \text{ km}^2$ of the city of Tokyo. The dataset contains 76k perspective views. The query set contains 315 query images from 105 distinct locations captured by different types of camera phones. This dataset is very challenging as each location is captured at three different times: during day, at sunset and during night. The dataset is available upon request at Torii (2015).

7.2 Image Descriptors

We perform experiments with two types of image descriptors: the sparse high-dimensional bag-of-visual-word vectors (Sivic and Zisserman 2003) and the compact (not-sparse) Fisher vectors (Jégou et al. 2012). Details of each are given next.

7.2.1 Bag-of-Visual-Word Representation

We extract SURF descriptors (Bay et al. 2006) for each image and learn a vocabulary of 100k visual words by approximate k-means clustering (Philbin et al. 2007) from a subset of fea-

tures from 5,000 randomly selected database images. Then, a tf-idf weighted vector (Sivic and Zisserman 2003) is computed for each image by assigning each descriptor to the nearest cluster center. Finally, all database vectors are normalized to have unit L_2 norm.

7.2.2 Fisher Vectors

Following (Jégou et al. 2012) we project the extracted 128-dimensional rootSIFT (Arandjelović and Zisserman 2012) descriptors to 64 dimensions using PCA. The projection matrix is learnt on a set of descriptors from 5,000 randomly selected database images. This has also the effect of decorrelating the rootSIFT descriptor. The 64-dimensional descriptors are then aggregated into Fisher vectors using a Gaussian mixture model with $N = 256$ components, which results in a $2 \times 256 \times 64 = 32,768$ -dimensional descriptor for each image. The Gaussian mixture model is learnt from descriptors extracted from 5,000 randomly sampled database images. The high-dimensional Fisher vector descriptors are then projected down to dimension using PCA learnt from all available images in the database. The resulting low dimensional Fisher vectors are then normalized to have unit L_2 -norm, which we found to be important in practice.

7.3 Parameters of per-location classifier learning

To learn the exemplar support vector machine for each database image j , the positive and negative training data are constructed as follows. The *negative training set* \mathcal{N}_j is obtained by: (i) finding the set of images with geographical distance greater than 200 m; (ii) sorting the images by decreasing value of similarity to image j measured by the dot product between their respective descriptors; (iii) taking the top $N = 500$ ranked images as the negative set. In other words, the negative training data consists of the hard negative images, i.e. those that are similar to image j but are far away from its geographical position, hence, cannot have the same visual content. The *positive training set* \mathcal{P}_j consist of the descriptor \mathbf{x}_j of the target image j .

We found that for the bag-of-visual-words representation it was useful to further expand (Chum et al. 2007) positive training set by close by images that view the same scene structures. These images can be identified by geometric verification (Philbin et al. 2007) as follows. We first build a graph where each image in the database represents a node and an edge represents a spatial adjacency in the world. An edge is present if the positions of the two images are within $50m$ of each other. Then, we score each edge by the number of geometrically verified matches (Philbin et al. 2007). Finally, we remove edges with score below a threshold of $t_m = 40$ matches. It is worth noting that the graph contains many isolated nodes. This typically indicates that the viewpoint

change between two adjacent panoramas is large. For each image in the database, we include between zero and five extra positive examples that are directly connected in the graph.

For the support vector machine classifier (SVM) training we use `libsvm` (Fan et al. 2008). We use the same global C_1 and C_2 parameters for all per-exemplar classifiers, but find the optimal value of the parameters for each image representation by a cross-validation evaluating performance on a held out query set.

For the calibration by re-normalization, we L_2 normalize the learned \mathbf{w}_j using Eq. (9) and use this normalized vector as the new image descriptor \mathbf{x}'_j for image j . At query time we compute the descriptor \mathbf{q} of the query image and measure its similarity score to the learnt descriptors \mathbf{x}'_j for each database image by Eq. (1).

For the p -value calibration, we take the learnt classifier for each database image j and compute its SVM classification score for all other database images to construct its empirical cumulative density function (4). We keep only the top 1,000 values that, in turn, represent the calibration function. At query time, given the query descriptor \mathbf{q} , we compute the SVM score (2) for each database image j , and compute its calibrated SVM score f_j (4).

8 Results

We evaluate the proposed per-location classifier learning approach on two different image descriptors: the bag-of-visual-words model (Sect. 8.1) and Fisher vectors (Sect. 8.2). We also compare the recognition accuracy of the two learnt representations relative to their compactness measured by their memory footprint (Sect. 8.3). Finally, we compare results to linear discriminant analysis (LDA) and whitening baselines (Sect. 8.4), outline the main failure modes (Sect. 8.5) and discuss the scalability of our method (Sect. 8.6).

Since the ground truth GPS position for each query image is available, for each method we measure performance using the percentage of correctly recognized queries (Recall) similarly to, (e.g., Chen et al. 2011; Knopp et al. 2010; Sattler et al. 2012). We deem the query as correctly localized if at least one of the top K retrieved database images is within 20 meters from the ground truth position of the query.

8.1 Bag-of-Visual-Words Model

Results for the bag-of-visual-words image representation are shown in Table 1. Learning per-location classifiers with either calibration method (p -val and w -norm) clearly improves over the standard bag-of-visual-words baseline (BOW) that does not perform any learning. In addition, both calibration methods significantly improve over the learnt SVM classifiers without any calibration (BOW SVM no calib) underscoring

the importance of calibration for the independently learnt per-location classifiers.

In Table 1, we also compare performance to our implementation of the confuser suppression approach (Conf. supp.) of Knopp et al. (2010) that, in each database image, detects and removes features that frequently appear at other far-away locations (using parameters $t = 3.5$ and $w = 70$). Our results show an improvement, specially at recall@1.

Inspecting the detailed plots in Fig. 4 we further note that the p -val calibration performs slightly better than the w -norm calibration for shorter top K shortlists but this effect is reversed for larger K . This could be attributed to the fact that the p -val calibration uses the negative data to control false positive errors, but has less control over false negatives, as discussed in Sect. 4.3.

In Fig. 6 we visualize the learnt SVM weights on BOW for p -val. We visualize the contribution of each feature to the SVM score for the corresponding query image. Red circles represent features with negative weights while green circles correspond to features with positive weights. The area of each circle is proportional to the contribution of the corresponding feature to the SVM score. For instance for the left figure notice that the correctly localized queries (c) contain more green colored features than queries from other places (b) and (a). Query (b) gets a high score because the building has orange and white stripes similar to the the sun-blinds of the bakery, which are features that also have large positive weights in the query image (c) of the correct place. In the top row we visualize the calibration of raw SVM score for three different queries. The calibration function of the target image j is shown in the blue and the corresponding SVM scores of the three queries are denoted by red circles. Notice that both images (b) and (c) have high calibrated score even their respective SVM score was different.

Finally, examples of correctly and incorrectly localized queries are shown in Fig. 9.

8.2 Fisher Vectors

Results of the proposed per-location learning method for the Fisher vector image representation for different dimensions are shown in Table 2 and Fig. 5. Similar to bag-of-visual-words, the learnt representation (w -norm) significantly improves the place recognition performance over the baseline Fisher vector (FV) matching without learning. The improvements are consistent across different lengths of shortlist K and for different dimensionality of the Fisher vector representation. We report results only for the w -norm calibration as we found that the p -val calibration did not perform well for the learnt Fisher vector classifiers (top 1 recall of 25.3% compared to baseline performance of 33.6% for dimension 128). When examining the results we have observed that for bag-of-visual-words the cdf estimated on

the database well represents the scores of (unseen) negative query images at test time. However, this is not the case for Fisher vectors where estimated cdf on the database does not represent well the scores of negative query images at test time. The scores of (unseen) negative query images often fall outside of the estimated cdf or at the very tail that is only sparsely sampled. As a result the estimated query image p values for Fisher vectors are often over-confident and incorrect (Figs. 6, 7).

Notice that the proposed per-location learning method consistently improves performance over the raw Fisher vector descriptors on the larger Pittsburgh 55k dataset and the challenging 24/7 Tokyo dataset (76k images). Examples of correctly and incorrectly localized queries are shown in Fig. 8. Next, we compare the performance of the two learnt representations relative to their memory footprints.

8.3 Analysis of Recognition Accuracy Versus Compactness

Here we analyze the recognition accuracy of the learnt representations vs. their compactness measured by their memory footprint on the Pittsburgh 25k image dataset. Ideally, we wish to learn a more compact representation, that still improves the recognition accuracy. However, usually there is a trade-off between the discriminative power of the representation and its size, where having a more compact representation reduces the recognition accuracy (Jégou et al. 2012). We observe a similar behavior but our learnt representation results in a higher recognition accuracy for a given size, or alternatively, significantly reduces the size of the representation for a given accuracy. The results are summarized in Fig. 7. The figure shows the recognition performance (y-axis) for the different dimensionality of the Fisher vector representation, which corresponds to different memory footprints (x-axis). For example, for $d = 128$ the memory footprint is about 24 MB, whereas for $d = 2048$ the memory footprint is about 384 MB. Note that the x-axis is in log-scale. The bag-of-visual-words representation has a fixed dimensionality (and fixed memory footprint) and hence each bag-of-visual-words method is shown only as a single point on the graph. For Fisher vectors, the results demonstrate that for a given level of accuracy (y-axis) the proposed method learns a more compact (lower-dimensional) representation (x-axis). For example, our learnt 128-dimensional descriptor (memory footprint of 24 MB) achieves a similar accuracy (around 65%) as the 256-dimensional raw Fisher descriptor (memory footprint of 51MB, interpolated from Fig. 7). This corresponds to 50% memory savings for the same level of recognition performance. Note that similar to (Jégou et al. 2012), we observe decrease in performance at high-dimensions for both the FV baseline and our method. The results also demonstrate the benefits of using the compact

Table 1 Evaluation of the learnt bag-of-visual-words representation on the Pittsburgh 25k dataset

Method:	25k Pittsburgh				
	1	2	5	10	20
recall@K [%]					
BOW SVM no calib.	6.4	8.1	13.5	17.5	20.5
BOW	28.7	35.7	45.8	53.7	61.5
BOW Conf. supp Knopp et al. (2010)	29.6	37.3	48.9	59.3	69.2
BOW w-norm	31.8	38.7	49.7	60.2	69.4
BOW p val	33.0	40.3	50.2	58.7	66.4

The table shows the fraction of correctly recognized queries (recall@K) for the different values of $K \in \{1, 2, 5, 10, 20\}$ retrieved database images. The learnt representations (BOW w-norm and BOW p val) outperform the raw bag-of-visual-words baseline (BOW) as well as the learnt representation without calibration (BOW SVM no calib)

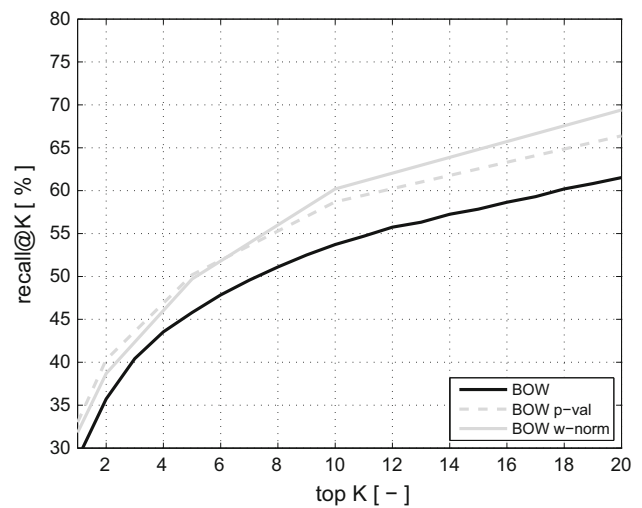


Fig. 4 Evaluation of the learnt bag-of-visual-words representation on the Pittsburgh 25k Gronát et al. (2013) dataset. The graph shows the fraction of correctly recognized queries (recall@K, y axis) versus the number of top K retrieved database images for the raw bag-of-visual-words baseline (BOW) and the learnt representation with two different calibration methods (p-val and w-norm)

FV descriptors compared to the bag-of-visual-words baseline achieving significantly better recognition accuracy for a similar memory footprint (Fig. 9).

8.4 Comparison to Linear Discriminant Analysis (LDA) and Whitening Baselines

We have compared our method to the linear discriminant analysis (LDA) (Aubry et al. 2014; Hariharan et al. 2012; Gharbi et al. 2012) and whitening (Jégou and Chum 2012) baselines. Results are reported on the Pittsburgh 25k dataset. The LDA baseline finds a discriminative linear projection of the feature space by minimizing an Euclidean loss, rather than the hinge loss used in our work. In detail, following (Aubry et al. 2014) we have used all available data to learn

Table 2 Evaluation of the learnt Fisher vector representation on the Pittsburgh (Gronát et al. 2013) and 24/7 Tokyo (Torii et al. 2015) datasets

recall@K [%]	1	2	5	10	20
Method/dataset	<i>25k Pittsburgh</i>				
FV128	33.6	41.8	52.0	59.8	67.7
FV128 w-norm	38.3	47.5	57.7	65.8	72.7
FV512	44.3	51.7	61.4	68.7	75.2
FV512 w-norm	47.6	55.4	65.1	72.4	78.8
FV2048	46.9	54.1	63.8	70.5	76.8
FV2048 w-norm	50.2	57.3	67.0	73.8	78.0
FV16384	45.3	54.1	63.8	69.4	75.3
FV16384 w-norm	49.3	56.0	65.9	72.5	76.8
Method/dataset	<i>55k Pittsburgh</i>				
FV128	10.9	14.1	20.2	26.4	33.2
FV128 w-norm	13.5	17.7	25.0	31.8	39.0
FV512	17.3	21.1	28.4	34.2	40.3
FV512 w-norm	19.8	25.1	32.7	38.7	46.0
FV2048	19.2	23.5	29.9	35.2	41.9
FV2048 w-norm	20.8	25.9	33.1	38.7	45.9
Method/dataset	<i>24/7 Tokyo</i>				
FV128	14.2	20.0	27.9	34.2	41.5
FV128 w-norm	16.9	22.0	29.6	37.2	44.8
FV512	35.2	40.3	43.8	48.2	57.1
FV512 w-norm	36.1	42.0	46.8	52.8	61.4
FV2048	37.4	42.5	48.5	53.9	58.7
FV2048 w-norm	42.9	46.7	52.8	58.8	66.7
FV4096	42.9	46.3	54.0	59.0	64.8
FV4096 w-norm	44.3	47.1	54.7	61.1	66.5

The table shows the fraction of correctly recognized queries (recall@K) for the different values of $K \in \{1, 2, 5, 10, 20\}$ retrieved database images. The learnt Fisher vector representation (FV *w-norm*) consistently improves over the standard Fisher vector matching baseline (FV) for all target dimensions

the covariance matrix and used the calibrated LDA score (see Aubry et al. (2014) Eq. 11) to obtain a classifier for each database image. We have applied the LDA method on the 128-dimensional Fisher vector descriptor but have obtained significantly worse performance (31.9% for recall@1) than our method (recall@1 of 38.3%). We believe the better performance of our method can be attributed to (i) the use of hinge-loss and (ii) training using the top scoring hard negative examples that are specific for each place.

Next, we compare results to PCA compression followed by whitening as suggested in Jégou and Chum (2012). For bag-of-visual-words, we follow Jégou and Chum (2012) and compare performance to PCA whitening to a target dimension of 4096. We have observed performance drop compared




Fig. 5 Evaluation of the learnt Fisher vector representation on the Pittsburgh 25k (Gronát et al. 2013) dataset. The graph shows the fraction of correctly recognized queries (recall@K, y axis) versus the number of top K retrieved database images for the raw Fisher vector baseline (FV) for different dimensions compared to the learnt representation (*w-norm*). Note the consistent improvements over all lengths of shortlist K for all dimensions

to the raw bag-of-visual-words baseline (28.7% to 26.1% for recall@1). We hypothesize this could be attributed to the large dictionary size used in our work (100k), whereas Jégou and Chum (2012) report improved results for single dictionary whitening only for dictionaries of up to 32k visual words. Finally, we have also applied PCA whitening on Fisher vector descriptors of dimensions 128, 512 and 2048, but have not observed significant improvements over the baseline raw descriptors. In fact, for the highest dimension (2048) we have observed a performance drop (49.6% to 41.3%), which could be attributed to amplification of low-energy noise as also reported in Jégou and Chum (2012).

8.5 Analysis of Improvement and Failure Cases

We have examined the improvement and failures of the *w-norm* method w.r.t. the Fisher vector baseline on the Pittsburgh 25k dataset. We analyzed the cases for which the *w-norm* method improves the rank of the first true positive compared to the baseline and for which the rank of the first true positive is made worse. In detail, considering a shortlist of the size 20 we want to identify when: (i) an image with the rank of >20 is attracted into the short list (improvement), and (ii) an image with the rank of ≤ 20 is pushed out of the short list using our method (failure).

We observe that in 237 cases a low-ranked true positive image by the baseline (ranked 40–70) is attracted into the short list by the *w-norm* method, resulting in an improvement. Note that in many other cases our method improves ranking but here we only count the cases for which the baseline




Fig. 6 A visualization of learnt feature weights for two database images. In each panel: *first row (Right)* Target database image j . *(Left)* Cumulative density function (or calibrated score) learnt for the SVM scores of the corresponding classifier f_j ; three query images displayed on the *second row* are represented by their SVM scores and cdf values $F_0(s)$, denoted (a–c) on the graph. *Third row* A visualization of the contribution of each feature to the SVM score for the corresponding query image. *Red circles* represent features with negative weights while *green circles* correspond to features with positive weights. The area of each circle is proportional to the contribution of the corresponding feature to the SVM score. Notice that the correctly localized queries **c** contain

more green colored features than queries from other places **b** and **a**. Please note also that the calibration *cdfs* in the left and right panel are similar but not identical. *Left panel* Query **b** gets a high score because the building has orange and white stripes similar to the the sun-blinds of the bakery, which are features that also have large positive weights in the query image **c** of the correct place. *Right panel* Query **b** is in fact also an image of the same location with a portion of the left skyscraper in the target image detected in the *upper left corner* and the side of the rightmost building in the target image detected in the *top right corner*. Both are clearly detected by the method as indicated by a large quantity of *green circles* in the corresponding regions (Color figure online)

method does not have any true positives in the top 20 short-list. On the other hand, in 39 cases our method downranks correct images appearing in the top 20 short-list. However, these images are typically downranked only slightly and still occur in the top 40 shortlist.

Finally, we observe that the downranking typically occurs on hard examples where the baseline performance is already bad. When visually inspecting the failure cases we have observed that our method typically fails on queries containing a large amount of clouds or vegetation and images containing narrow streets and tunnels. Our method also sometimes retrieves images capturing the same building from a different side or a large distance.

8.6 Scalability

In the offline stage, our method collects hard negative examples for each location in the database, which are consequently used to train exemplar SVM classifiers. As only a constant number of examples (1–5 positives and 500 negatives) is used to train each per-location classifier the overall complexity of training is linear, $O(N)$, i.e. we need to train one classifier (with constant training time) for each of N images in the database. The bottleneck of the offline stage is collecting

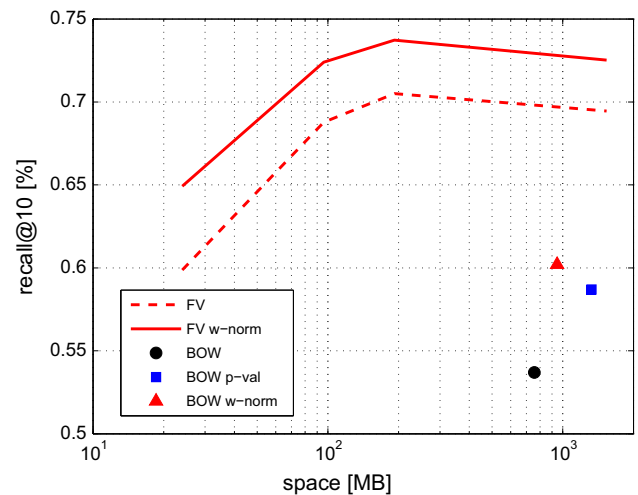


Fig. 7 The recognition performance versus the memory requirements for the Pittsburgh 25k dataset. The fraction of correctly localized queries at the top 10 retrieved images (*y axis*) versus the memory footprint (*x axis*) for the different representations. For Fisher vectors, the learnt descriptor (FV *w-renorm*) clearly outperforms the raw Fisher vector descriptor (FV) for all dimensions corresponding to different memory footprints (*x axis*). Learnt per-location representations for the bag-of-visual-words model (BOW *p-val* and BOW *w-norm*) also improve performance over the raw bag-of-visual-words (BOW). However, the Fisher vectors provide much better recognition performance for the same memory footprint



Fig. 8 Examples of correctly and incorrectly localized queries for the learnt bag-of-visual-words representation. Each example shows a query image (*left*) together with correct (*green*) and incorrect (*red*) matches from the database obtained by learnt bag-of-visual-words representation *p-val* method (*top*) and the standard bag-of-visual-words baseline

(*bottom*). Note that the proposed method is able to recognize the place depicted in the query image despite changes in viewpoint, illumination and partial occlusion by other objects (*trees, lamps*) and buildings. Note also that bag-of-visual-words baseline is often confused by repeating patterns on facades and walls (Color figure online)

the negative examples that is quadratic $O(N^2)$ in the database size. In other words, for each of N database images, we need to find the top 500 most similar negatives among all N database images. However, we believe that even finding

negatives can be scaled-up to very large datasets with standard compression techniques such as product quantization (PQ) (Jegou et al. 2011) combined with sub-linear approximate nearest neighbor search (Muja and Lowe 2014).




Fig. 9 Examples of correctly and incorrectly localized queries for the learnt Fisher vector representation. Each example shows a query image (left) together with correct (green) and incorrect (red) matches from the database obtained by the learnt Fisher vector representation w -norm method (top) and the standard Fisher vector baseline (bottom) for dimension 128. Note that the proposed method is able to recognize

the place depicted in the query image despite changes in viewpoint, illumination and partial occlusion by other objects (trees, lamps) and buildings. Note that the baseline methods often finds images depicting the same buildings but in a distance whereas our learnt representation often finds a closer view better matching the content of the query (Color figure online)

At query time our method needs to compute the calibrated e-SVM score (Eq. (2)) of the query for each image in the database. In the case of w -norm method the calibration weights can be included in the classifier weight

matrix, as discussed in Sect. 6. For the p-val calibration method, each e-SVM score must be calibrated using K stored values of the non-parametric CDF model. This requires a search for the two closest values and subsequent interpola-

tion, which yields complexity of $O(N \log K)$. Since K is only a constant both the w-norm and p-val methods have a linear time complexity (in the size of the database) at query time but with different constants. However, in practice the constant in the p-val method can be quite large. The actual running time per query is 340ms for the bag-of-visual-words representation with p-val calibration and 3ms for the FV128 descriptor with w-norm calibration. Both timings are on the 25k Pittsburgh dataset on a desktop with an Intel Xenon E5 CPU using a single thread. Hence in practice, the p-val method may be scalable only to medium size datasets. For the w-norm method, the query time can be further sped-up using sub-linear approximate nearest neighbor search (Muja and Lowe 2014) on compressed descriptors (Jegou et al. 2011), making the method scalable to very large datasets.

9 Conclusions

We have shown that place recognition can be cast as a classification problem and have used geotags as a readily-available supervision to train an ensemble of classifiers, one for each location in the database. As only few positive examples are available for each location, we have developed two procedures to calibrate the output of each classifier without the need for additional positive training data. We have shown that learning per-location representations improves the place recognition performance over the raw bag-of-visual-words and Fisher vector matching baselines. The developed calibration methods are not specific to place recognition and can be useful for other per-exemplar classification tasks, where only a small number of positive examples are available (Malisiewicz et al. 2011).

Acknowledgments This work was supported by the MSR-INRIA laboratory, the EIT-ICT labs, Google, the ERC project LEAP and the EC project RVO13000 - Conceptual development of research organization. Supported by the Intelligence Advanced Research Projects Activity (IARPA) via Air Force Research Laboratory, contract FA8650-12-C-7212. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, AFRL or the U.S. Government.

Appendix

In Sect. 8 we show that the simple calibration by normalization often results in surprisingly good place recognition performance without the need for any additional positive or negative calibration data. In this appendix, we give a possible explanation why this simple calibration works. We focus

on the case of a single positive training example, i.e. when training set $\mathcal{P} = \mathbf{x}^+$, which is the typical case for place recognition where only one positive example is available for each place. The analysis holds also for the case of multiple expanded positive examples as in our case the positive examples are coming from the same database of Street View images, and hence have very similar statistics (illumination, capturing conditions, the same camera, etc.).

In particular, we first analyze the SVM objective and show that the learnt hyperplane \mathbf{w} can be interpreted as a new descriptor \mathbf{x}^* that replaces the original positive example \mathbf{x}^+ and is re-weighted to increase its separation from the negative data. Second, we show that when \mathbf{x}^* is normalized, i.e. $\mathbf{x}^* = \frac{\mathbf{w}}{\|\mathbf{w}\|}$, the dot-product $\mathbf{q}^T \mathbf{x}^*$ corresponds to measuring the cosine of the angle between the (normalized) query descriptor \mathbf{q} and the new descriptor \mathbf{x}^* , which was found to work well in the literature for descriptor matching, as discussed in Sect. 5.2. The two steps are given next.

Analysis of Per-Exemplar SVM Objective

For a single positive example $\mathcal{P} = \mathbf{x}^+$, the per-exemplar SVM objective (3) can be written as

$$\Omega(\mathbf{w}, b) = \|\mathbf{w}\|^2 + C_1 \cdot h(\mathbf{w}^T \mathbf{x}^+ + b) + C_2 \sum_{\mathbf{x} \in \mathcal{N}} h(-\mathbf{w}^T \mathbf{x} - b). \tag{13}$$

In the following, we analyze the objective (13) and provide intuition why *re-normalized* weight vector \mathbf{w} can be interpreted as a new descriptor. In particular, we show first that when the weight C_2 of the negative data in objective (13) goes to zero the learnt normalized $\tilde{\mathbf{w}}$ is identical to the original positive training data point \mathbf{x}^+ . Second, when $C_2 > 0$, the learnt vector $\tilde{\mathbf{w}}$ moves away from the positive vector \mathbf{x}^+ to increase its separation from the negative data. The two cases are detailed next.

Case 1 $C_2 \rightarrow 0$. The goal is to show that when the weight C_2 of the negative data in objective (13) goes towards zero, the resulting hyperplane vector \mathbf{w} is parallel with the vector of positive training descriptor \mathbf{x}^+ . When \mathbf{w} is normalized to have unit L2 norm the two vectors are identical. First, let us decompose \mathbf{w} into parallel and orthogonal part with respect to the positive training data point \mathbf{x}^+ , i.e. $\mathbf{w} = \mathbf{w}^\perp + \mathbf{w}^\parallel$, where $(\mathbf{w}^\perp)^T \mathbf{x}^+ = 0$. Next, we observe that when the weight of the negative data diminishes ($C_2 \rightarrow 0$), any non-zero component \mathbf{w}^\perp will increase the value of the objective. As a result, for $C_2 \rightarrow 0$ the objective is minimized by \mathbf{w}^\parallel , i.e. the optimal \mathbf{w} is parallel with \mathbf{x}^+ .

In detail, for $\mathbf{w} = \mathbf{w}^\perp + \mathbf{w}^\parallel$, the objective (3) can be written as




Fig. 10 An illustration of the effect of decreasing parameter C_2 in the exemplar support vector machine objective. The positive exemplar \mathbf{x}^+ is shown in green. The negative data points are shown in red. All training data is L2 normalized to lie on a hyper-sphere. **a** For $C_2 > 0$, the normal \mathbf{w} of the optimal hyper-plane moves away from the direction given by the positive example \mathbf{x}^+ in a manner that reduces the loss on the negative data. **b** As the parameter C_2 decreases the learnt \mathbf{w} becomes parallel to the positive training example \mathbf{x}^+ and its magnitude $\|\mathbf{w}\|$ goes to 0 (Color figure online)

$$\|\mathbf{w}^\perp + \mathbf{w}^\parallel\|^2 + C_1 \cdot h\left((\mathbf{w}^\perp + \mathbf{w}^\parallel)^T \mathbf{x}^+ + b\right) + C_2 \sum_{\mathbf{x} \in \mathcal{N}} h\left(-(\mathbf{w}^\perp + \mathbf{w}^\parallel)^T \mathbf{x} - b\right). \tag{14}$$

Note that the orthogonal part \mathbf{w}^\perp does not change the value of the second term in (14) because $(\mathbf{w}^\perp + \mathbf{w}^\parallel)^T \mathbf{x}^+ = (\mathbf{w}^\parallel)^T \mathbf{x}^+$, and hence (14) reduces to

$$\|\mathbf{w}^\perp + \mathbf{w}^\parallel\|^2 + C_1 \cdot h\left(\mathbf{w}^\parallel^T \mathbf{x}^+ + b_j\right) + C_2 \sum_{\mathbf{x} \in \mathcal{N}} h\left(-(\mathbf{w}^\perp + \mathbf{w}^\parallel)^T \mathbf{x} - b\right). \tag{15}$$

In the limit case as $C_2 \rightarrow 0$ any non-zero component \mathbf{w}^\perp will increase the value of the objective (15). This can be seen by noting that the third term vanishes when $C_2 \rightarrow 0$ and hence the objective is dominated by the first two terms. Further, the second term in (15) is independent of \mathbf{w}^\perp . Finally, the first term will always increase for any non-zero value of w^\perp as $\|\mathbf{w}^\perp + \mathbf{w}^\parallel\|^2 \geq \|\mathbf{w}^\parallel\|^2$ for any $\mathbf{w}^\perp \neq 0$.

As a result, in the limit case when $C_2 \rightarrow 0$ the optimal \mathbf{w} is parallel with \mathbf{x}^+ . Note also, that when C_2 is exactly equal to zero, $C_2 = 0$, the optimal \mathbf{w} vanishes, i.e. the objective (15) is minimized by trivial solution $\|\mathbf{w}\| = 0$ and $b = -1$. The effect of decreasing the parameter C_2 is illustrated in Fig. 10.

Case II $C_2 > 0$. When the weight C_2 of the negative data in the objective (15) increases the direction of the optimal w will be different from \mathbf{w}^\parallel and will change to take into account the loss on the negative data points. Explicitly writing the hinge-loss $h(x) = \max(1 - x, 0)$ in the last term of (15), we see that \mathbf{w} will move in the direction that reduces $\sum_{\mathbf{x} \in \mathcal{N}} \max(1 + \mathbf{w}^T \mathbf{x} + b, 0)$, i.e. that reduces the dot prod-

uct $\mathbf{w}^T \mathbf{x}$ on the negative examples that are active (support vectors).

The Need for Normalization of \mathbf{w}

Above we have shown that the learnt hyperplane \mathbf{w} moves away from the positive example \mathbf{x}^+ in a manner that reduces the loss on the negative data. The aim is to use this learnt vector \mathbf{w} as a new descriptor \mathbf{x}^* replacing the original positive example \mathbf{x}^+ . However, we wish to measure the cosine of the angle between the the new descriptor \mathbf{x}^* and the query image \mathbf{q} . This is equivalent to the normalized dot product, hence the vector \mathbf{w} needs to be normalized.

References

Agarwal, S., Snavely, N., Simon, I., Seitz, S. & Szeliski, R. (2009). Building Rome in a day. In *ICCV* (pp. 72–79).

Arandjelović, R. & Zisserman, A. (2012). Three things everyone should know to improve object retrieval. In *IEEE PAMI*.

Aubry, M., Maturana, D., Efros, A., Russell, B. & Sivic, J. (2014). Seeing 3D chairs: exemplar part-based 2D–3D alignment using a large dataset of CAD models. In *CVPR*.

Aubry, M., Russell, B. & Sivic, J. (2014) Painting-to-3D model alignment via discriminative visual elements. *ACM Transactions on Graphics*.

Bay, H., Tuytelaars, T. & Van Gool, L. (2006). SURF: Speeded up robust features. In *ECCV*.

Cao, S. & Snavely, N. (2013). Graph-based discriminative learning for location recognition. In *IEEE Conference on CVPR* (pp. 700–707).

Casella, G. & Berger, R. (2001). *Statistical inference*.

Chen, D., Baatz, G., Köser, Tsai, S., Vedantham, R., Pylvanainen, T., Roimela, K., Chen, X., Bach, J., Pollefeys, M., Girod, B. & Grzeszczuk, R. (2011). City-scale landmark identification on mobile devices. In *CVPR*.

Chum, O., Philbin, J., Sivic, J., Isard, M. & Zisserman, A. (2007). Total recall: Automatic query expansion with a generative feature model for object retrieval. In *ICCV*.

Csurka, G., Bray, C., Dance, C., & Fan, L. (2004). Visual categorization with bags of keypoints. In *Workshop on Statistical Learning in Computer Vision, ECCV* (pp. 1–22).

Cummins, M. & Newman, P. (2009). Highly scalable appearance-only SLAM - FAB-MAP 2.0. In *Proceedings of Robotics: Science and Systems*, Seattle, USA.

Dalal, N. & Triggs, B. (2005). Histogram of oriented gradients for human detection. In *CVPR*.

Doersch, C., Gupta, A. & Efros, A.A. (2013). Mid-level visual element discovery as discriminative mode seeking. In *NIPS*.

Doersch, C., Singh, S., Gupta, A., Sivic, J., & Efros, A. A. (2012). What makes Paris look like Paris? *SIGGRAPH*, 31(4), 101.

Fan, R. E., Chang, K. W., Hsieh, C. J., Wang, X. R., & Lin, C. J. (2008). LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research*, 9, 1871–1874.

Gebel, M., & Weihs, C. (2007). Calibrating classifier scores into probabilities. *Advances in Data Analysis* (pp. 141–148). Berlin: Springer.

Gharbi, M., Malisiewicz, T., Paris, S., & Durand, F. (2012). A Gaussian approximation of feature space for fast image similarity. Technical Report, MIT.

Google: ICMLA 2011 streetview recognition challenge. <http://www.icmla-conference.org/icmla11/challenge.htm>.

- Gronát, P.: Project webpage: Learning and calibrating per-location classifiers for visual place recognition. <http://www.di.ens.fr/willow/research/perlocation/>.
- Gronát, P. (2015). streetget. <http://www.di.ens.fr/willow/research/streetget/>.
- Gronát, P., Obozinski, G., Sivic, J. & Pajdla, T. (2013). Learning and calibrating per-location classifiers for visual place recognition. In *CVPR*.
- Hariharan, B., Malik, J. & Ramanan, D. (2012). Discriminative decorrelation for clustering and classification. In *ECCV*.
- Hays, J. & Efros, A.A. (2008). im2gps: estimating geographic information from a single image. In *CVPR*.
- Irschara, A., Zach, C., Frahm, J.M. & Bischof, H. (2009). From structure-from-motion point clouds to fast location recognition. In *CVPR*.
- Jégou, H. & Chum, O. (2012). Negative evidences and co-occurrences in image retrieval: The benefit of PCA and whitening. In *ECCV* (pp. 774–787).
- Jegou, H., Douze, M., & Schmid, C. (2011). Product quantization for nearest neighbor search. *IEEE Transactions on PAMI*, 33(1), 117–128.
- Jégou, H., Perronnin, F., Douze, M., Sánchez, J., Pérez, P., & Schmid, C. (2012). Aggregating local image descriptors into compact codes. *IEEE Transactions on PAMI*, 34, 1704–1716.
- Kalogerakis, E., Vesselova, O., Hays, J., Efros, A. & Hertzmann, A. (2009). Image sequence geolocation with human travel priors. In *ICCV* (pp. 253–260).
- Klingner, B., Martin, D. & Roseborough, J. (2013). Street view motion-from-structure-from-motion. In *ICCV*.
- Knopp, J., Sivic, J. & Pajdla, T. (2010). Avoiding confusing features in place recognition. In *ECCV*.
- Li, Y., Crandall, D. & Huttenlocher, D. (2009). Landmark classification in large-scale image collections. In *ICCV*.
- Li, Y., Snavely, N. & Huttenlocher, D. (2010). Location recognition using prioritized feature matching. In *ECCV*.
- Li, Y., Snavely, N., Huttenlocher, D. & Fua, P. (2012). Worldwide pose estimation using 3d point clouds. In *ECCV*.
- Lowe, D. (2004). Distinctive image features from scale-invariant keypoints. *IJCV*, 60(2), 91–110.
- Malisiewicz, T., Gupta, A. & Efros, A.A. (2011). Ensemble of exemplar-svm for object detection and beyond. In *ICCV*.
- Muja, M. & Lowe, D.G. (2014). Scalable nearest neighbor algorithms for high dimensional data. In *IEEE Transactions on PAMI* 36.
- Nister, D. & Stewenius, H. (2006). Scalable recognition with a vocabulary tree. In *CVPR*.
- Philbin, J., Chum, O., Isard, M., Sivic, J. & Zisserman, A. (2007). Object retrieval with large vocabularies and fast spatial matching. In *CVPR*.
- Philbin, J., Sivic, J. & Zisserman, A. (2010). Geometric latent dirichlet allocation on a matching graph for large-scale image datasets. In *IJCV*.
- Platt, J. (1999). Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in Large Margin Classifiers*, 10(3), 61–74.
- Sattler, T., Leibe, B. & Kobbelt, L. (2012). Improving image-based localization by active correspondence search. In *ECCV*.
- Sattler, T., Weyand, T., Leibe, B., & Kobbelt, L. (2012). Image retrieval for image-based localization revisited. In *Proceedings of BMVC*.
- Scheirer, W., Kumar, N., Belhumeur, P.N. & Boult, T.E. (2012). Multi-attribute spaces: Calibration for attribute fusion and similarity search. In *CVPR*.
- Schindler, G., Brown, M. & Szeliski, R. (2007). City-scale location recognition. In *CVPR*.
- Scholkopf, B., & Smola, A. (2002). *Learning with kernels*. Cambridge: MIT press.
- Shrivastava, A., Malisiewicz, T., Gupta, A. & Efros, A.A. (2011). Data-driven visual similarity for cross-domain image matching. In *SIGGRAPH ASIA*.
- Singh, S., Gupta, A. & Efros, A.A. (2012). Unsupervised discovery of mid-level discriminative patches. In *ECCV*.
- Sivic, J., Zisserman, A. (2003). Video Google: A text retrieval approach to object matching in videos. In *ICCV*. <http://www.robots.ox.ac.uk/vgg>.
- Tighe, J. & Lazebnik, S. (2013). Finding things: Image parsing with regions and per-exemplar detectors. In *CVPR*.
- Torii, A. (2015). Project webpage: 24/7 place recognition by view synthesis. <http://www.ok.ctrl.titech.ac.jp/torii/project/247/>.
- Torii, A., Arandjelović, R., Sivic, J., Okutomi, M. & Pajdla, T. (2015). 24/7 place recognition by view synthesis. In *CVPR*.
- Torii, A., Sivic, J. & Pajdla, T. (2011). Visual localization by linear combination of image descriptors. In *IEEE Workshop on Mobile Vision*.
- Torii, A., Sivic, J., Pajdla, T. & Okutomi, M. (2013) Visual place recognition with repetitive structures. In *CVPR*.
- Turcot, P., & Lowe, D. (2009). Better matching with fewer features: The selection of useful features in large database recognition problem. In *WS-LAVD, ICCV*.
- Zadrozny, B. & Elkan, C. (2002) Transforming classifier scores into accurate multiclass probability estimates. In *ACM SIGKDD*.
- Zamir, A. & Shah, M. (2010) Accurate image localization based on google maps street view. In *ECCV*.