

Harvesting Image Databases from the Web

F. Schroff, Dept. of Engineering Science, University of Oxford, UK

A. Criminisi, Microsoft Research Ltd., Cambridge, UK

A. Zisserman, Dept. of Engineering Science, University of Oxford, UK

Abstract—

The objective of this work is to automatically generate a large number of images for a specified object class. A multi-modal approach employing both text, meta data and visual features is used to gather many high-quality images from the web.

Candidate images are obtained by a text based web search querying on the object identifier (e.g. the word penguin). The web pages and the images they contain are downloaded. The task is then to remove irrelevant images and re-rank the remainder. First, the images are re-ranked based on the text surrounding the image and meta data features. A number of methods are compared for this re-ranking. Second, the top-ranked images are used as (noisy) training data and a SVM visual classifier is learnt to improve the ranking further. We investigate the sensitivity of the cross-validation procedure to this noisy training data. The principal novelty of the overall method is in combining text/meta-data and visual features in order to achieve a completely automatic ranking of the images.

Examples are given for a selection of animals, vehicles and other classes totalling 18 classes. The results are assessed by precision/recall curves on ground truth annotated data and by comparison to previous approaches including those of Berg *et al.* [5] and Fergus *et al.* [12].



1 INTRODUCTION

The availability of image databases has proved invaluable for training and testing object class models during the recent surge of interest in object recognition. However, producing such databases containing a large number of images and with high precision is still an arduous manual task. Image search engines apparently provide an effortless route, but currently are limited by poor precision of the returned images and restrictions on the total number of images provided. For example, with Google Image Search the precision is as low as 32% on one of the classes tested here (shark) and averages at 39%, and downloads are restricted to 1000 images.

Fergus *et al.* [12, 11], Lin *et al.* [20], Li *et al.* [19], Collins *et al.* [8], Vijayanarasimhan *et al.* [33] and Fritz&Schiele [14] dealt with the precision problem by re-ranking the images downloaded from an image search. The method in [12] involved visual clustering of the images by using probabilistic Latent Semantic Analysis (pLSA) [15] over a visual vocabulary, [19] used a Hierarchical Dirichlet Process instead of pLSA, and [33] use multiple instance learning to learn the visual models. Lin *et al.* [20] re-ranked using the text on the original page from which the image was obtained. However, for all these methods the yield is limited by the restriction on the total number of images provided by the image search.

Berg *et al.* [5] overcome the download restriction by starting from a *web* search instead of an *image* search. This search can generate thousands of images. Their method then proceeds in two stages: first, topics are discovered based on words occurring on the web pages using Latent Dirichlet Allocation (LDA) [6] on *text* only. Image clusters for each topic are formed by selecting images where nearby text is top ranked by the topic. A user then partitions the clusters into positive and negative for the class. Second, images and the associated text from these clusters are used as exemplars to train a classifier

based on voting on visual (shape, colour, texture) and text features. The classifier is then used to re-rank the downloaded dataset. Note, the user labelling of clusters avoids the problem of polysemy, as well as providing good training data for the classifier. The method succeeds in achieving a greater yield, but at the cost of manual intervention.

Our objective in this work is to harvest a large number of images of a particular class *automatically*, and to achieve this with high precision. Our motivation is to provide training databases so that a new object model can be learnt effortlessly. Following [5] we also use web search to obtain a large pool of images and the web pages that contain them. The low precision does not allow us to learn a class model from such images using vision alone. The challenge then is how best to combine text, meta-data and visual information in order to achieve the best image re-ranking.

The two main contributions are: first, we show in section 3 that meta-data and text attributes on the web page containing the image provide a useful estimate of the probability that the image is in class, and thence can be used to successfully rank images in the downloaded pool. Second, we show in section 4 that this probability is sufficient to provide (noisy) training data for a visual classifier, and that this classifier delivers a superior re-ranking to that produced by text alone. Figure 1 visualises this two stage improvement over the initially downloaded images. The class independent text ranker significantly improves this unranked baseline and is itself improved by quite a margin when the vision based ranker (trained on the text ranker results) is employed. We compared our proposed discriminative framework (SVM) to unsupervised methods (topic models) concluding that the discriminative approach is better suited for this task and thus the focus of this work.

Others have used text and images together, however in a slightly different setting. For example, Barnard *et al.* [2] use groundtruth annotated images as opposed to noisy annotation

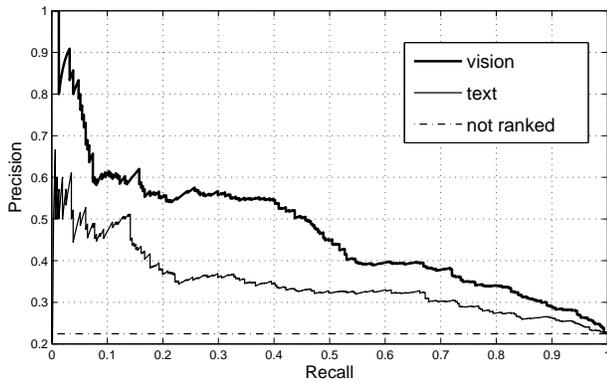


Fig. 1. **Text & visual ranking vs. unranked-baseline:** precision recall plot for the text re-ranking, the visual ranking trained on the text ranking, and the (unranked) original downloaded images, for the “shark” query.

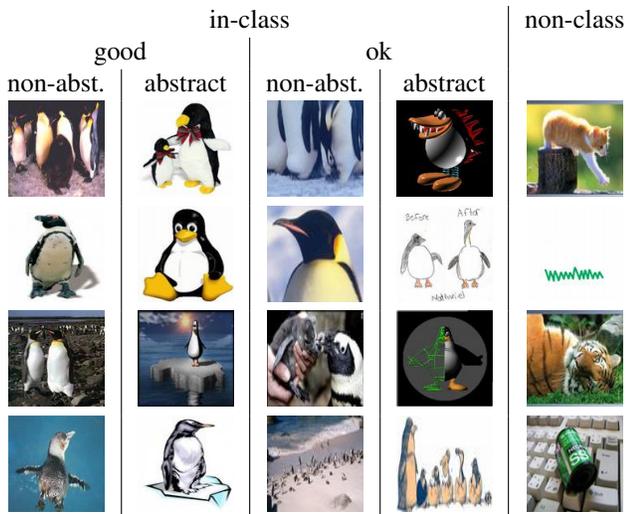


Fig. 2. **Image annotations:** Example images corresponding to annotation categories for the class penguin.

stemming from web pages, as in our case. Other work of Berg *et al.* [4] uses text from the Internet, but focused on identifying a specific class rather than general object classes.

We show in section 5.1 that our automatic method achieves superior ranking results to those produced by the method of Berg *et al.* [5] and also to that of Google Image Search.

This paper is an extended version of [28]. The extensions include: a comparison of different text-ranking methods, additional visual features (HOG), an investigation of the cross-validation to noise in the training data, and a comparison of different topic models (for the visual features).

2 THE DATABASES

This section describes the methods for downloading the initial pool of images (together with associated meta-data) from the Internet, and the initial filtering that is applied. For the purposes of training classifiers and for assessing precision and recall the downloaded images are annotated manually for 18

classes: airplane (ap), beaver (bv), bikes (bk), boat (bt), camel (cm), car (cr), dolphin (dp), elephant (ep), giraffe (gf), guitar (gr), horse (hs), kangaroo (kg), motorbikes (mb), penguin (pg), shark (sk), tiger (tr), wristwatch (ww), zebra (zb).

Data collection: We compare three different approaches to downloading images from the web. The first approach, named WebSearch, submits the query word to Google web search and all images that are linked within the returned web pages are downloaded. Google limits the number of returned web pages to 1000, but many of the web pages contain multiple images, so in this manner thousands of images are obtained. The second approach, ImageSearch, starts from Google image search (rather than web search). Google image search limits the number of returned images to 1000, but here each of the returned images is treated as a “seed” – further images are downloaded from the web page from where the seed image originated. The third approach, GoogleImages, includes only the images directly returned by Google image search (a subset of those returned by ImageSearch). The query can consist of a single word or more specific descriptions such as “penguin animal” or “penguin OR penguins”. Images smaller than 120×120 are discarded. In addition to the images, text surrounding the image HTML tag is downloaded together with other meta-data such as the image filename.

Groundtruth annotation: In a similar manner to Fergus *et al.* [12], images are divided into three categories:

in-class-good: Images that contain one or many class instances in a clearly visible way (without major occlusion, lighting deterioration or background clutter and of sufficient size).

in-class-ok: Images that show parts of a class instance, or obfuscated views of the object due to lighting, clutter, occlusion and the like.

non-class: Images not belonging to *in-class*.

The good and ok sets are further divided into two subclasses: **abstract:** Images that do not resemble realistic natural objects (*e.g.* drawings, non realistic paintings, comics, casts or statues).

non-abstract: Images not belonging to the previous class.

Example annotations for the class penguin are shown in figure 2. The full dataset is published in [29]. As is usual in annotation there are ambiguous cases, *e.g.* deciding when occlusion is sufficiently severe to classify as *ok* rather than *good*, or when the objects are too small. The annotations were made as consistent as possible by a final check from one person. Note, the *abstract vs. non-abstract* categorisation is not general but is suitable for the object classes we consider in this paper. For example, it would not be useful if the class of interest was “graph” or “statue” or a similar more abstract category.

Table 1 details the statistics for each of the three retrieval techniques (WebSearch, ImageSearch and GoogleImages). Note that some images are common between the methods. ImageSearch gives a very low precision (only about 4%) and is not used for the harvesting experiments. This low precision is probably due to the fact that Google selects many images from web-gallery pages, which contain images of all sorts. Google is able to select the in-class

Service	in-class	non-class	precision
WebSearch	8773	25252	26%
ImageSearch	5963	135432	4%
GoogleImages	4416	6766	39%

TABLE 1

Statistics by source: The statistics of downloaded images for different retrieval techniques for 18 classes.

images from those pages, *e.g.* the ones with the object-class in the filename, however if we use those web-pages as seeds the overall precision greatly decreases. Therefore, we only use WebSearch and GoogleImages, which are merged into one dataset per object class. Table 2 lists the 18 categories downloaded and the corresponding statistics for *in-class* and *non-class* images. The overall precision of the images downloaded for all 18 classes is about 29%.

Due to the great diversity of images available on the Internet and because of how we retrieve the images, it is difficult to make general observations on how these databases look. However, it is clear that polysemy affects the returned images. Interestingly, this is not a problem that could be predicted directly from the English word, since most of the classes we search for don’t have direct polysemous meanings, *i.e.* they are not polysemous in the sense of “bank” (as in place to get money, or river bank) for example. It is rather that the words correspond to brands or product names (“leopard tank”) or team names (the NHL ice hockey team “San Jose Sharks”) or are used as attributes (“tiger shark”). Apart from that, the *in-class* images occur in almost all variations imaginable, as sharks crashed into houses or other oddities. Even though context [31] can clearly be important in re-ranking the images (*e.g.* camel and kangaroo in desert-like images), it will have its limitations due to the variety of occurrences of the object.

2.1 Removing Drawings & Symbolic Images

Since we are mostly interested in building databases for natural image recognition, we ideally would like to remove all *abstract* images from the downloaded images. However, separating *abstract* images from all others automatically is very challenging for classifiers based on visual features. Instead we tackle the easier visual task of removing drawings & symbolic images. These include: comics, graphs, plots, maps, charts, drawings and sketches, where the images can be fairly simply characterised by their visual features (see below). Example images are shown in figure 3. Their removal significantly reduces the number of *non-class* images, improving the resulting precision of the object class datasets as shown in table 2 (overall precision goes from 29% to 35%). Filtering out such images also has the aim of removing this type of *abstract* image from the *in-class* images.

Learning the filter: We train a radial basis function Support Vector Machine (SVM) on a hand labelled dataset (examples in figure 3). After the initial training no further user interaction is required. In order to obtain this dataset, images were downloaded using ImageSearch with one level of recursion

Class	downloaded images			after drawing&symbolic filtering			
	in-cl.	non-cl.	prec.	in-cl.	non-cl.	prec.	false pos.
airplane (ap)	904	1659	35.27%	635	1007	38.67%	91
beaver (bv)	236	3121	7.03%	160	2195	6.79%	4
bikes (bk)	1268	1931	39.64%	983	1082	47.60%	111
boat (bt)	856	2175	28.24%	726	1354	34.90%	70
camel (cm)	594	1808	24.73%	485	1274	27.57%	46
car (cr)	1128	1042	51.98%	938	568	62.28%	92
dolphin (dp)	791	1416	35.84%	533	906	37.04%	81
elephant (ep)	937	1558	37.56%	763	1007	43.11%	11
giraffe (gf)	945	1267	42.72%	802	763	51.25%	32
guitar (gr)	1219	2035	37.46%	873	832	51.20%	248
horse (hs)	1229	1720	41.68%	975	1043	48.32%	78
kangaroo (kg)	418	1763	19.17%	329	1161	22.08%	14
motorbikes (mb)	732	953	43.44%	607	582	51.05%	86
penguin (pg)	748	1400	34.82%	447	794	36.02%	33
shark (sk)	583	1710	25.43%	413	1089	27.50%	60
tiger (tr)	379	2068	15.49%	311	1274	19.62%	17
wristwatch (ww)	941	957	49.58%	710	549	56.39%	220
zebra (zb)	483	1662	22.52%	416	987	29.65%	19
total	14391	30245	32.24%	11106	18467	37.55%	1313

TABLE 2

Image class statistics of the original downloaded images using WebSearch&GoogleImages only, and after applying the drawing&symbolic images removal filter.

(*i.e.* web pages linked from “seed” web pages are also used) with queries such as “sketch” or “drawing” or “draft”. The goal was to retrieve many images and then select suitable training images manually. The resulting dataset consists of approximately 1400 drawings&symbolic images, and 2000 non drawings&symbolic images.

Three simple visual only features are used: (i) a colour-histogram; (ii) a histogram of the L2-norm of the gradient; (iii) a histogram of the angles ($0 \dots \pi$) weighted by the L2-norm of the corresponding gradient. In all cases 1000 equally spaced bins are used. The motivation behind this choice of features is that drawings&symbolic images are characterised by sharp edges in certain orientations and or a distinctive colour distribution (*e.g.* only few colours in large areas). The method achieves around 90% classification accuracy on the drawings&symbolic images database (using two-fold cross-validation).

This classifier is applied to the entire downloaded image dataset to filter out drawing&symbolic images, before further processing. The total number of images that are removed for each class is shown in table 2. In total 42% of *non-class* images are removed over all classes. The remaining images are those used in our experiments. As well as successfully removing *non-class* images, the filter also succeeds in removing an average of 60% (123 images) *in-class abstract* images, with a range between 45% (for motorbikes, 40 images) and 85% (for wristwatch, 11 images). There is some loss of the desired *in-class non-abstract* images, with on average 13% (90

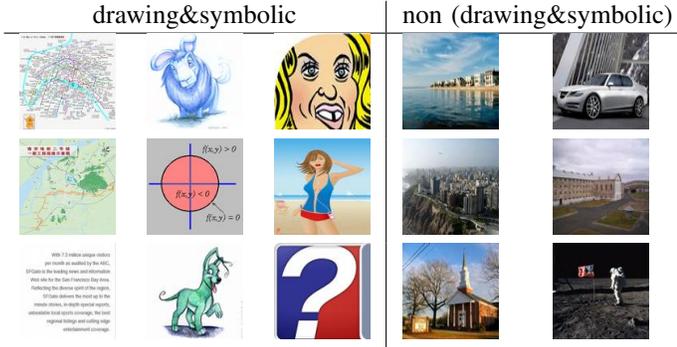


Fig. 3. **Drawings&symbolic images:** Examples of positive and negative training images.

images) removed, though particular classes lose a relatively high percentage (28% for shark and wristwatch). Even though this seems to be a high loss the precision of the resulting datasets is improved in all cases except for the class shark.

3 RANKING ON TEXTUAL FEATURES

We now describe the re-ranking of the returned images based on text and meta-data alone. Here we follow and extend the method proposed by Frankel *et al.* [13] in using a set of textual attributes whose presence is a strong indication of the image content.

Textual features: We use seven features from the text and HTML-tags on the web page: *contextR*, *context10*, *filedir*, *filename*, *imagealt*, *imagetitle*, *websitetitle*.

Filedir, *filename* and *websitetitle* are self-explanatory. *Context10* includes the ten words on either side of the image-link. *ContextR* describes the words on the web-page between eleven and 50 words away from the image-link. *Imagealt* and *imagetitle* refer to the “alt” and “title” attribute of the image-tag. The features are intended to be conditionally independent, given the image content (we address this independence below). It is difficult to compare directly with the features in [13], since no precise definition of the features actually used is given.

Context here is defined by the HTML source, not by the rendered page since the latter depends on screen resolution and browser type and is an expensive operation. In the text processing a standard stop list [24] and the Porter stemmer [25] are used. In addition HTML-tags and domain specific stop words (such as “html” or “ ”) are ignored.

We also experimented with a number of other features, such as the image MIME type (‘gif’, ‘jpeg’ *etc.*), but found that they did not help discrimination.

3.1 Image ranking

Using these seven textual features, the goal is to re-rank the retrieved images. Each feature is treated as binary: “True” if it contains the query word (*e.g.* penguin) and “False” otherwise. The seven features define a binary feature vector for each image $\mathbf{a} = (a_1, \dots, a_7)$, and the ranking is then based on the posterior probability, $P(y = \textit{in-class} | \mathbf{a})$, of the image being *in-class*, where $y \in \{\textit{in-class}, \textit{non-class}\}$ is the class-label of an image.

We learn a class *independent* ranker in order to re-rank the images based on the posterior $P(y | \mathbf{a})$. To re-rank images for one particular class (*e.g.* penguin) we do not employ the groundtruth data for that class. Instead, we train the Bayes classifier (specifically we learn $P(\mathbf{a} | y)$, $P(y)$ and $P(\mathbf{a})$) using all available annotations, *except* the class we want to re-rank. This way we evaluate performance as a *completely automatic class independent* image ranker; *i.e.* for any new and unknown class, the images can be re-ranked without *ever* using labelled groundtruth knowledge of that class.

3.1.1 Ranking models

We compare different Bayesian posterior models for $P(y = \textit{in-class} | \mathbf{a})$. Specifically we looked at the suitability of the following decompositions:

Chow-Liu dependence tree decomposition [7]:

$$P(\mathbf{a} | y) \propto \prod_1^8 P(x_i | x_{m(j)}) \quad (1)$$

with $x = (a_1, \dots, a_7, y)$ and m being a permutation of $(1, \dots, 8)$. The Chow-Liu model approximates the full joint dependency graph as a tree by retaining the edges between variables with the highest mutual information.

Naïve Bayes model:

$$P(\mathbf{a} | y) \propto \prod_1^7 P(a_i | y) \quad (2)$$

Given the class label for an image the text features are assumed to be independent. For our application this is “obviously” not the case, *e.g.* filename and image alternative tag are highly correlated.

Pairwise dependencies:

$$P(\mathbf{a} | y) \propto \prod_{i,j=1}^7 P(a_i, a_j | y) \quad (3)$$

Only pairwise dependencies are modelled. This is similar to the Chow-Liu model, but less sparse.

Full joint:

$$P(\mathbf{a} | y) \propto P(a_1, \dots, a_7 | y) \quad (4)$$

The full joint probability distribution is learnt. If the amount of available training data is too small the learnt model can be inaccurate.

Mixed naïve Bayes:

$$P(\mathbf{a} | y) \propto P(a_1, \dots, a_4 | y) \prod_5^7 P(a_i | y) \quad (5)$$

where $P(a_1, \dots, a_4 | y)$ is the joint probability of the first four textual features (*contextR*, *context10*, *filedir*, *filename*).

Logistic regression: Additionally we evaluate the performance using the discriminative logistic regression model, where

$$P(y | \mathbf{a}) = \frac{1}{1 + e^{-\mathbf{w}^T \mathbf{a}}} \quad (6)$$

prec.	ap	bv	bk	bt	cm	cr	dp	ep	gf	gr	hs	kg	mb	pg	sk	tr	ww	zb	avg.
Mixed Naïve Bayes																			
15%	40.61	30.00	70.44	52.22	49.65	61.16	70.27	70.70	82.39	63.18	57.72	53.93	67.94	68.42	54.05	43.81	78.03	48.41	59.05
100	39.00	31.00	76.00	52.00	47.00	66.00	69.00	77.00	87.00	58.00	58.00	53.00	67.00	69.00	55.00	43.00	79.00	50.00	59.78
avg.	41.67	31.17	58.66	48.97	48.17	66.46	59.18	64.64	76.10	65.71	57.21	40.95	66.62	61.53	42.55	37.61	70.59	52.97	55.94
Pairs																			
15%	41.52	34.29	67.45	54.64	50.00	61.43	68.42	64.91	82.98	61.95	57.72	57.14	66.42	69.15	43.17	43.40	81.10	53.04	58.82
100	43.00	36.00	66.00	51.00	53.00	61.00	69.00	62.00	79.00	59.00	61.00	53.00	66.00	69.00	38.00	41.00	81.00	49.00	57.61
avg.	46.64	33.47	60.53	45.19	45.63	65.24	59.83	62.28	74.40	65.23	55.18	43.35	65.55	62.74	39.30	38.26	71.69	53.52	54.89
Chow-Liu tree																			
15%	38.11	19.35	62.45	53.27	47.97	61.16	61.42	64.91	72.22	68.65	55.91	32.88	69.53	57.52	41.10	38.33	66.03	40.67	52.86
100	43.00	19.00	69.00	55.00	50.00	59.00	63.00	65.00	84.00	70.00	58.00	33.00	72.00	56.00	41.00	34.00	80.00	33.00	53.78
avg.	41.11	21.08	56.48	45.30	45.50	66.55	57.18	58.36	73.47	61.44	55.29	38.65	66.13	55.48	42.57	36.01	63.39	45.58	51.64
Naïve Bayes																			
15%	41.89	38.71	67.45	54.36	50.00	61.71	70.27	64.91	81.25	61.95	57.49	57.14	66.42	70.65	43.17	41.07	81.10	58.65	59.34
100	43.00	36.00	66.00	51.00	54.00	61.00	68.00	62.00	79.00	61.00	53.00	66.00	69.00	37.00	40.00	81.00	58.00	58.00	58.00
avg.	46.67	34.21	60.58	45.22	45.49	65.11	59.85	62.32	74.36	65.52	55.24	43.46	65.55	62.93	39.79	38.43	71.69	52.13	54.92
Full joint																			
15%	41.33	32.00	69.76	53.54	47.02	59.83	71.56	72.55	85.40	61.95	58.92	39.67	65.44	67.71	45.45	41.07	79.84	55.96	58.28
100	37.00	32.00	77.00	56.00	45.00	62.00	71.00	74.00	84.00	57.00	55.00	42.00	63.00	66.00	45.00	41.00	82.00	53.00	58.00
avg.	42.13	32.19	60.59	49.26	49.09	66.64	62.60	63.18	75.88	65.75	57.64	37.64	65.59	62.12	44.92	38.64	71.33	54.13	55.52

TABLE 3

Precision of textual re-ranking: The performance of the textual re-ranking for all 18 classes over different models: precision at 15% recall, precision at 100 images recall, and average precision. The precision is given as a percentage. The last column gives the average over all classes. Mixed naïve Bayes performs best and was picked for all subsequent experiments.

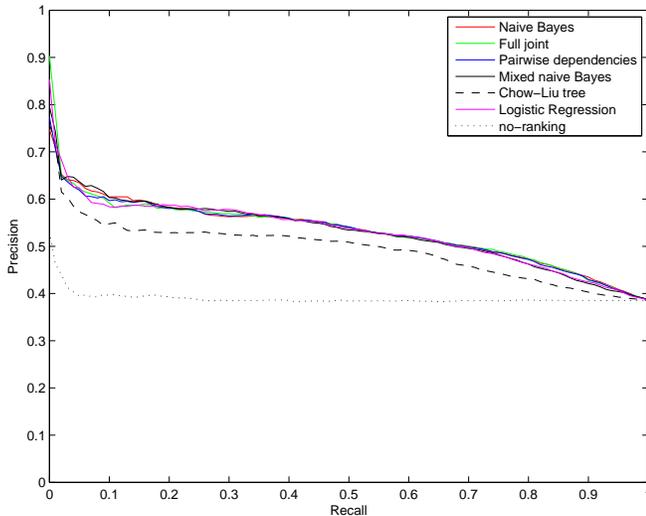


Fig. 4. **Average precision-recall of text-rankers.** The precision-recall curve averaged over all 18 classes for each of the described ranking methods.

3.1.2 Text re-ranking results

We assess the performance by reporting precision at various points of recall as well as average precision in table 3. Figure 4 and 5 give an overview of the different methods and their performance. Figure 6 shows the precision-recall curves for selected classes using the mixed naïve model. It can clearly be seen that precision is highly increased at the lower recall levels, compared to the average precision of table 2.

The mixed model (5) gave slightly better performance than other factorisations *on 100 images recall*, and reflects the fact that the first four features are less independent from each

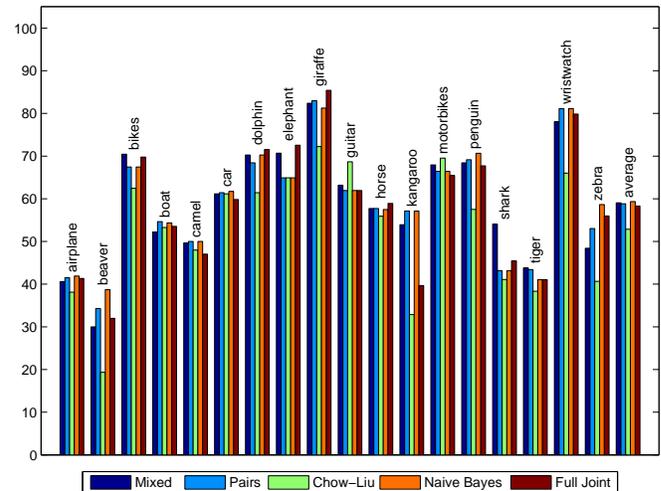


Fig. 5. **Comparison of text-rankers.** The precision at 15% recall for all 18 classes and four different models. See table 3 for details and the average precision over all classes.

other than the remaining three. Overall all models perform comparably and the differences are negligible, except for the Chow-Liu dependence tree which performs slightly worse. We chose the mixed naïve model for our experiments as the 100 images recall is more related to our selection of training data, than average precision or precision at 15% recall.

A separate set of experiments was carried out to measure how the performance of the text ranker varies with the number and choice of classes used for training. Ideally we would like to compare $P(\mathbf{a}|y)$, $P(y)$ and $P(\mathbf{a})$ learnt using different numbers of training classes. However, given our goal of ranking

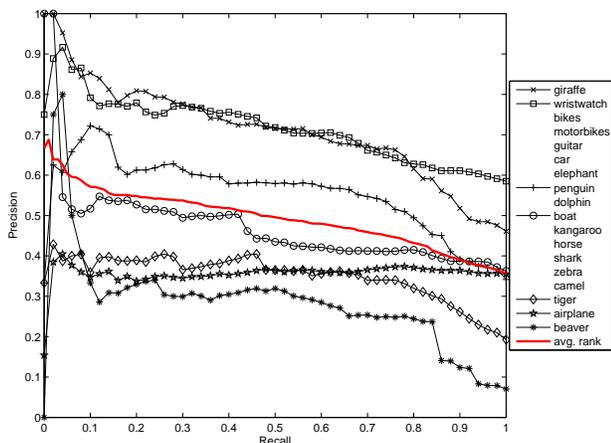


Fig. 6. **Text based re-ranking:** precision vs. recall estimated for each class with *abstract* images considered *in-class* using the mixed naïve Bayes model. The labels are shown in decreasing order of precision at 15% recall. The recall precision curves are only shown for selected classes for clarity. The average over *all* 18 classes is also shown.

images we instead compare these probabilities indirectly by assessing precision at 15% recall. We find that the performance is almost unaltered by the choice of training classes provided more than five classes (chosen randomly) are used for training.

Discussion: As can be seen from figure 6 the text re-ranker performs well on average, and significantly improves the precision up to quite a high recall level (see figure 8 and 9 for top ranked images). In section 4 we will show that this is sufficient to train a visual classifier. For some classes the text ranker performs very well (*e.g.* wristwatch, giraffe) for others it performs rather poorly (*e.g.* airplane, beaver, camel, tiger). Visual inspection of the highly ranked “outlier” (*non-class*) images in the text ranked lists gives some explanation for these performances. Classes that perform well (wristwatch, giraffe) generally have outliers that are unrelated to each other. In contrast for the classes that perform poorly the outlier images are related and result from lack of discriminativity of the query word – for example for airplanes there are images of airplane food, airports, toy airplanes, paper airplanes, airplane interiors, and advertisements with comic airplanes. Other classes suffer from the type of polysemy described in section 2: for camels there are brand and cigarette related outliers; and for tiger there is the attribute problem with images of helicopters, tanks, fish (sharks), boxing, golf, stones, and butterflies.

These examples cover very different classes from animals to various man-made objects. Generalisation to other well defined object classes is demonstrated in the experiments of section 5.1 on the object classes of [5] and [12].

We investigated two alternative text-based classifiers, pLSA and SVM, in addition to the Bayes estimator finally adopted, but found they had inferior performance. For the SVM the same binary text features \mathbf{a} were used. It is possible that the binary features lead to the poor performance of the SVM. For the pLSA we used *context10* and *contextR* (similar to

[5]). Due to problems in the pLSA clustering, the problem of how to select the right topic without user interaction as in [5], and the question of how to use the additional binary features (*e.g.* *filename*) in a principled manner, we adopted the Bayes estimator instead.

The text re-ranking is not meant to compete with the performance of Internet image search engines, which could also be used in our algorithm and given their recent performance improvements would be a reasonable choice. Instead we decided to use this simple setup to gather the training data for our visual models, with the focus of independence and a more controlled setup, considering the fast changing quality of image search engines.

4 RANKING ON VISUAL FEATURES

The text re-ranking of section 3 associates a posterior probability with each image as to whether it contains the query class or not. The problem we are now faced with is how to use this information to train a visual classifier which would improve the ranking further. The problem is one of training from noisy data: we need to decide which images to use for positive and negative training data and how to select a validation set in order to optimise the parameters of the classifier.

We first describe the visual features used, and then how the classifier is trained.

Visual features: We follow the approach of [12] and use a variety of region detectors with a common visual vocabulary in the bag of visual words model framework (BOW). All images are first resized to 300 pixels in width. Regions are detected using: difference of Gaussians, Multiscale-Harris [22], Kadir’s saliency operator [18], and points sampled from Canny edge points. Each image region is represented as a 72 dimensional SIFT [21] descriptor. A separate vocabulary consisting of 100 visual words is learnt for each detector using k-means, and these vocabularies are then combined into a single one of 400 words. Finally, the descriptor of each region is assigned to the vocabulary. The software for the detectors is obtained from [32]. Fuller implementation details are given in [12], and are reproduced in our implementation. The 72 dimensional SIFT is based on the work in [12] and driven by the motivation that a coarser spatial binning can improve generalisation as opposed to a finer binning that is more suitable for particular object matching.

In addition to the descriptors used in [12] we add the widely used HOG descriptor [9], computed over the whole image to incorporate some spatial layout into the model. It was also used in a similar setting in [14]. We use a cellsize of 8 pixels, a blocksize of one cell, and 9 contrast invariant gradient bins. This results in a 900 dimensional feature vector, as the images were resized to 80×80 pixels. The two descriptors are concatenated resulting in a 1300 dimensional feature vector per image.

4.1 Training the visual classifier

At this point we can select n_+ positive training images from the top of the text ranked list, or those that have a posterior probability above some threshold, but a subset of these positive

images will be “noisy”, i.e. will not be *in-class*. Table 4 (text) gives an idea of the noise from the proportion of outliers. It averages at 40% if $n_+ = 100$. However, we can assume that the *non-class* images are *not* visually consistent – an assumption verified to some extent by the results in section 4.2. The case of negative images is more favourable: we select n_- images at random from all downloaded images (i.e. from all 18 classes, tens of thousands of images) and the chances of any image being of a particular class is very low. We did not choose to select the n_- images from the low ranked images of the text-ranker output, because the probability of finding *in-class* images there is higher than finding them in the set of *all* downloaded images.

Given this situation we choose to use a SVM classifier, since it has the potential to train despite noise in the data. The SVM training minimizes the following sum [23]:

$$\min_{\mathbf{w}, b, \xi} \frac{1}{2} \mathbf{w}^T \mathbf{w} + C_+ \sum_{i: y_i=1} \xi_i + C_- \sum_{j: y_j=-1} \xi_j \quad (7)$$

$$\text{subject to} \quad y_l (\mathbf{w}^T \Phi(\mathbf{x}_l) + b) \geq 1 - \xi_l, \quad (8)$$

$$\xi_l \geq 0, l = 1, \dots, (n_+ + n_-). \quad (9)$$

Where \mathbf{x}_l are the training vectors and $y_l \in \{1, -1\}$ the class labels. C_+ and C_- are the false classification penalties for the positive and negative images, with ξ being the corresponding slack-variables.

To implement the SVM we use the publicly-available SVM^{light} software [17] (with the option to remove inconsistent training data enabled). Given two input images I_i and I_j and their corresponding normalized histograms of visual words and HOG, S_i and S_j , this implementation uses the following χ^2 radial basis function (RBF) kernel: $K(S_i, S_j) = \exp(-\gamma \cdot \chi^2(S_i, S_j))$ [35]. With γ the free kernel parameter.

Thus, γ , C_+ and C_- are the three parameters that can be varied. The optimal value for these parameters is obtained by training the SVM using ten-fold cross validation. Note, we do not use the groundtruth at any stage of training, but we split the *noisy* training images into ten training and validation sets. I.e. the validation set is a subset of the n_+ top text-ranked images as well as the n_- background images. We require a performance measure for the cross-validation and use precision at 15% recall, computed on the validation subset of the n_+ images (treated as positive) and the n_- images (treated as negative). This parameter selection is performed automatically for *every* new object-class.

Sometimes C_+ and C_- are used to correct unbalanced training data [23]. In our case, however, the SVM is very sensitive to these parameters, probably due to the huge amount of noise in the data, and the optimal value does not directly correspond to the ratio of positive to negative images.

Finally, the trained SVM is used to re-rank the filtered image set based on the SVM classification score. The entire image harvesting algorithm is summarised in figure 7.

4.2 Results for textual/visual image ranking

In this section we evaluate different combinations of training and testing. If not stated otherwise the text+vision system of

- 1) Download images and meta-data for new class (e.g. “lion”) using WebSearch & GoogleImages (section 2).
- 2) Filter images: remove drawings&symbolic images (section 2.1).
- 3) Rank images based on text-attributes using the Bayes classifier (section 3).
- 4) Train visual SVM classifier on text-ranked images (section 4).
- 5) Rank all images from 1. (or 2.) using the visual classifier.

Fig. 7. Overview of the text+vision (t+v) image harvesting algorithm.

figure 7 was used. Results are given in table 4 for various choices of n_+ and n_- . For each choice, five different random selections are made for the sets used in the ten-fold cross-validation, and mean and standard deviation are reported. The clear improvement brought by the visual classifier over the text based ranking for most classes is obvious. Figure 8 and 9 compare the top ranked images from the text and vision steps.

We first investigate how the classification performance is affected by the choice of n_+ and n_- . It can be seen that increasing n_- tends to improve performance. It is, however, difficult to select optimal values for n_+ and n_- since these numbers are very class dependent. Table 4 indicates that using more images in the background class n_- tends to improve performance but there is no real difference between using 150/1000 and 250/1000 (n_+/n_-), which perform at $68.4\% \pm 1.9$ and $68.0\% \pm 2.0$ and are thus not significantly different. All numbers in this section report precision at 15% recall.

It can be seen (table 4) that HOG alone performs significantly worse than the bag of visual words $57.9\% \pm 1.8$, but the combination of BOW and HOG improves the overall performance to $69.8\% \pm 2.1$, compared to BOW alone $68.0\% \pm 2.0$.

In order to select the appropriate parameter values we use cross-validation, where the validation set is part of the n_+ and n_- images as described in section 4.1, together with precision at 15% recall as selection criterion. There are two possible cases that can occur: (i) a parameter setting that overfits to the training data. This problem is detected on the validation set, due to a low precision at 15% recall. (ii) all images (training and validation set) are classified as background. This leads to bad, but detectable, performance as well.

Here we describe a slight adjustment to this method, that ignores “difficult” images. Parameter settings that classify (almost) all images as fore- or background are not useful, neither are those that overfit to the training data. We reject those parameter settings. We then use the “good” parameter settings to train and classify all images. By looking at the distribution of SVM responses (over all parameter settings) we are able to eliminate “intermediate” images, i.e. images that are not classified as positive or negative images in the majority of cases. We assume that those images are difficult to classify and we don’t use those in our model selection step, because we cannot rely on their class labels being correct due to the noise. Note that training is still performed on *all* images, and the “difficult” images are only discarded during the computation of the precision at 15% recall during the cross-validation.

prec. 15%	ap	bv	bk	bt	cm	cr	dp	ep	gf	gr	hs	kg	mb	pg	sk	tr	ww	zb	avg.	std
text	40.61	30.00	70.44	52.22	49.65	61.16	70.27	70.70	82.39	63.18	57.72	53.93	67.94	68.42	54.05	43.81	78.03	48.41	59.05	–
t+v (250/250)	54.0	36.4	64.8	60.9	53.5	78.5	61.9	77.5	86.5	65.4	60.6	53.5	78.2	55.5	56.8	45.7	81.6	96.8	64.9	9.3 (2.5)
t+v (150/500)	59.1	33.8	69.7	63.4	56.8	91.8	61.6	70.7	83.2	67.3	67.0	50.0	77.4	70.4	69.3	59.6	86.3	85.5	67.9	8.3 (2.0)
t+v (250/500)	58.6	35.0	64.6	65.1	47.8	79.0	60.1	75.4	86.9	51.7	68.4	50.2	77.7	69.7	64.3	39.8	90.4	87.9	65.2	8.8 (2.0)
t+v (250/1000)	63.5	32.3	65.9	62.4	51.6	93.0	61.7	80.2	87.8	62.6	71.2	45.5	84.6	69.6	64.9	53.0	86.6	95.8	68.4	7.9 (1.9)
t+v (150/1000)	52.3	39.3	68.6	66.2	57.3	87.9	66.5	78.7	83.8	63.0	57.3	48.6	64.0	72.5	82.9	62.6	93.2	79.8	68.0	7.2 (2.0)
t+v (150/1000) C	49.8	42.7	71.0	66.2	55.1	88.8	64.7	78.0	85.7	62.3	49.2	56.2	67.0	72.6	69.8	58.5	88.9	86.1	67.4	8.3 (2.1)
HOG t+v	45.1	21.1	53.2	62.0	49.1	80.4	51.6	71.6	78.7	62.9	64.1	43.0	68.2	77.8	47.7	29.8	87.6	48.0	57.9	6.6 (1.8)
HOG+BOW t+v	42.8	46.4	70.4	60.9	54.5	93.5	67.7	84.8	88.3	62.5	68.8	55.2	72.1	78.9	80.7	57.3	89.7	81.4	69.8	8.1 (2.1)
HOG+BOW t+v C	51.3	42.3	68.2	60.3	63.2	91.1	69.7	78.7	88.7	66.3	70.1	53.9	76.6	90.2	66.1	50.2	92.5	91.7	70.6	7.2 (1.7)
gt (150/1000)	83.1	90.8	75.8	76.1	78.0	98.6	78.2	96.0	91.4	88.8	90.2	69.0	95.5	82.7	91.8	94.3	96.1	93.3	87.2	8.4 (2.0)
(B) (150/1000)	52.4	12.9	55.4	63.6	54.1	94.9	42.6	47.9	83.7	61.4	52.8	29.8	65.2	53.7	42.9	28.5	82.7	78.3	55.7	6.7 (1.7)

TABLE 4

Comparison of precision at 15% recall: ‘text’ refers to text re-ranking alone; ‘t+v’ is text+vision re-ranking using different training ratios n_+/n_- ; ‘gt’ is groundtruth (only positive images) training of the visual classifier; and (B) is the baseline, where the visual classifier is trained on $n_+ = 150$ images uniformly sampled from the filtered images of one class, instead of the text re-ranked images, and $n_- = 1000$ background images as before. The second last column (avg.) gives the average over all classes. The last column states the mean of the classwise standard deviations over five runs of cross-validation, as well as the standard deviation of the means over all classes, in parentheses.

This method does not give a significant improvement over all classes, but improves the performance dramatically for some classes, e.g. penguin $90.2\% \pm 5.0$ from $78.9\% \pm 13.2$. This modified version of the cross-validation is denoted ‘‘C’’ in table 4.

We next determine how much the performance is affected by the noise in the training data by training the SVM on groundtruth positive data, i.e. instead of selecting n_+ images from the text ranked images, we select n_+ *in-class* images using the groundtruth labelling. We find that the text+vision system performs well for the classes where the text ranking performs sufficiently well (see figure 8 and 9 for the top images returned by the text-ranker). HOG+BOW vs. gt: e.g. car $93.5\% \pm 7.1$ vs. $98.6\% \pm 1.1$, or giraffe $88.3\% \pm 2.7$ vs. $91.4\% \pm 4.8$. If the text ranking fails the groundtruth performs, as to be expected, much better than the text-ranked based training, e.g. for airplane, camel, kangaroo. These experiments show that the SVM based classifier is relatively insensitive to noise in the training data, as long as a reasonable noise level is given.

As a baseline comparison, we investigate performance if no text re-ranking is used, but the n_+ images are sampled uniformly from the filtered images. If the text re-ranking works well, and hence provides good training data, then text+vision improves over the baseline, e.g. elephant $84.8\% \pm 3.3$ vs. $47.9\% \pm 4.0$, penguin $78.9\% \pm 13.2$ vs. $53.7\% \pm 3.4$, or shark $80.7\% \pm 4.2$ vs. $42.9\% \pm 14.4$. In cases where the text ranking does not perform well the baseline can even outperform text+vision. This is due to the fact, that bad text ranking can provide visually consistent training data, that does *not* show the expected class (e.g. for airplanes it contains many images showing: airplane food, inside airplanes/airports, taken out of the window of an airplane). However, the uniformly sampled images still consist of about 35% *in-class* images (table 2) and the n_- are very unlikely to contain *in-class* images.

In addition to re-ranking the filtered images we applied the text+vision system to all images downloaded for one specific

class, i.e. the drawings&symbolic images were included. It is interesting to note that the performance is comparable to the case of filtered images (table 5). This means that the learnt visual model is strong enough, to remove the drawings&symbolic images during the ranking process. Thus, the filtering is only necessary to train the visual classifier and is not required to rank new images as evident from row two and three ‘‘tf ru’’ and ‘‘tu ru’’ in table 5. ‘‘tf ru’’ performs almost as good as ‘‘tf rf’’, $68.8\% \pm 1.8$ vs. $69.8\% \pm 2.1$, i.e. using unfiltered images during testing does not affect the performance significantly. However, using unfiltered images during training ‘‘tu ru’’, decreases the performance significantly down to $62.4\% \pm 2.5$. The main exception here is the airplane class where training with filtered images is a lot worse than with unfiltered images. In the case of airplane the filtering removed 91 good images and the overall precision of the filtered images is quite low 38.67%, which makes the whole process relatively unstable and can therefore explain the difference.

Discussion: Training and evaluating the system with the goal to build natural image databases, i.e. treating *abstract* images as *non-class* ($\not\in$) in all stages of the algorithm, as opposed to treating *abstract* images as *in-class*, gives similar performance, however slightly worse. The slight drop in performance can be explained by the fact that the text-ranker inevitably returns *abstract* images which are then used as training images. That our method is applicable to both those cases is further supported by the results we retrieve on the Berg *et al.* dataset (see section 5.1).

We also investigated alternative visual classifiers, topic models (see section 5.2), and feature selection. For feature selection our intention was to find discriminative visual words and then use these in a Bayesian framework. The discriminative visual words were obtained based on the likelihood ratio of a visual word occurring in the foreground to background images [10]. However, probably due to the large variation in both the foreground and background images, together with the noisy training data, we weren’t able to match the performance of the SVM ranker.

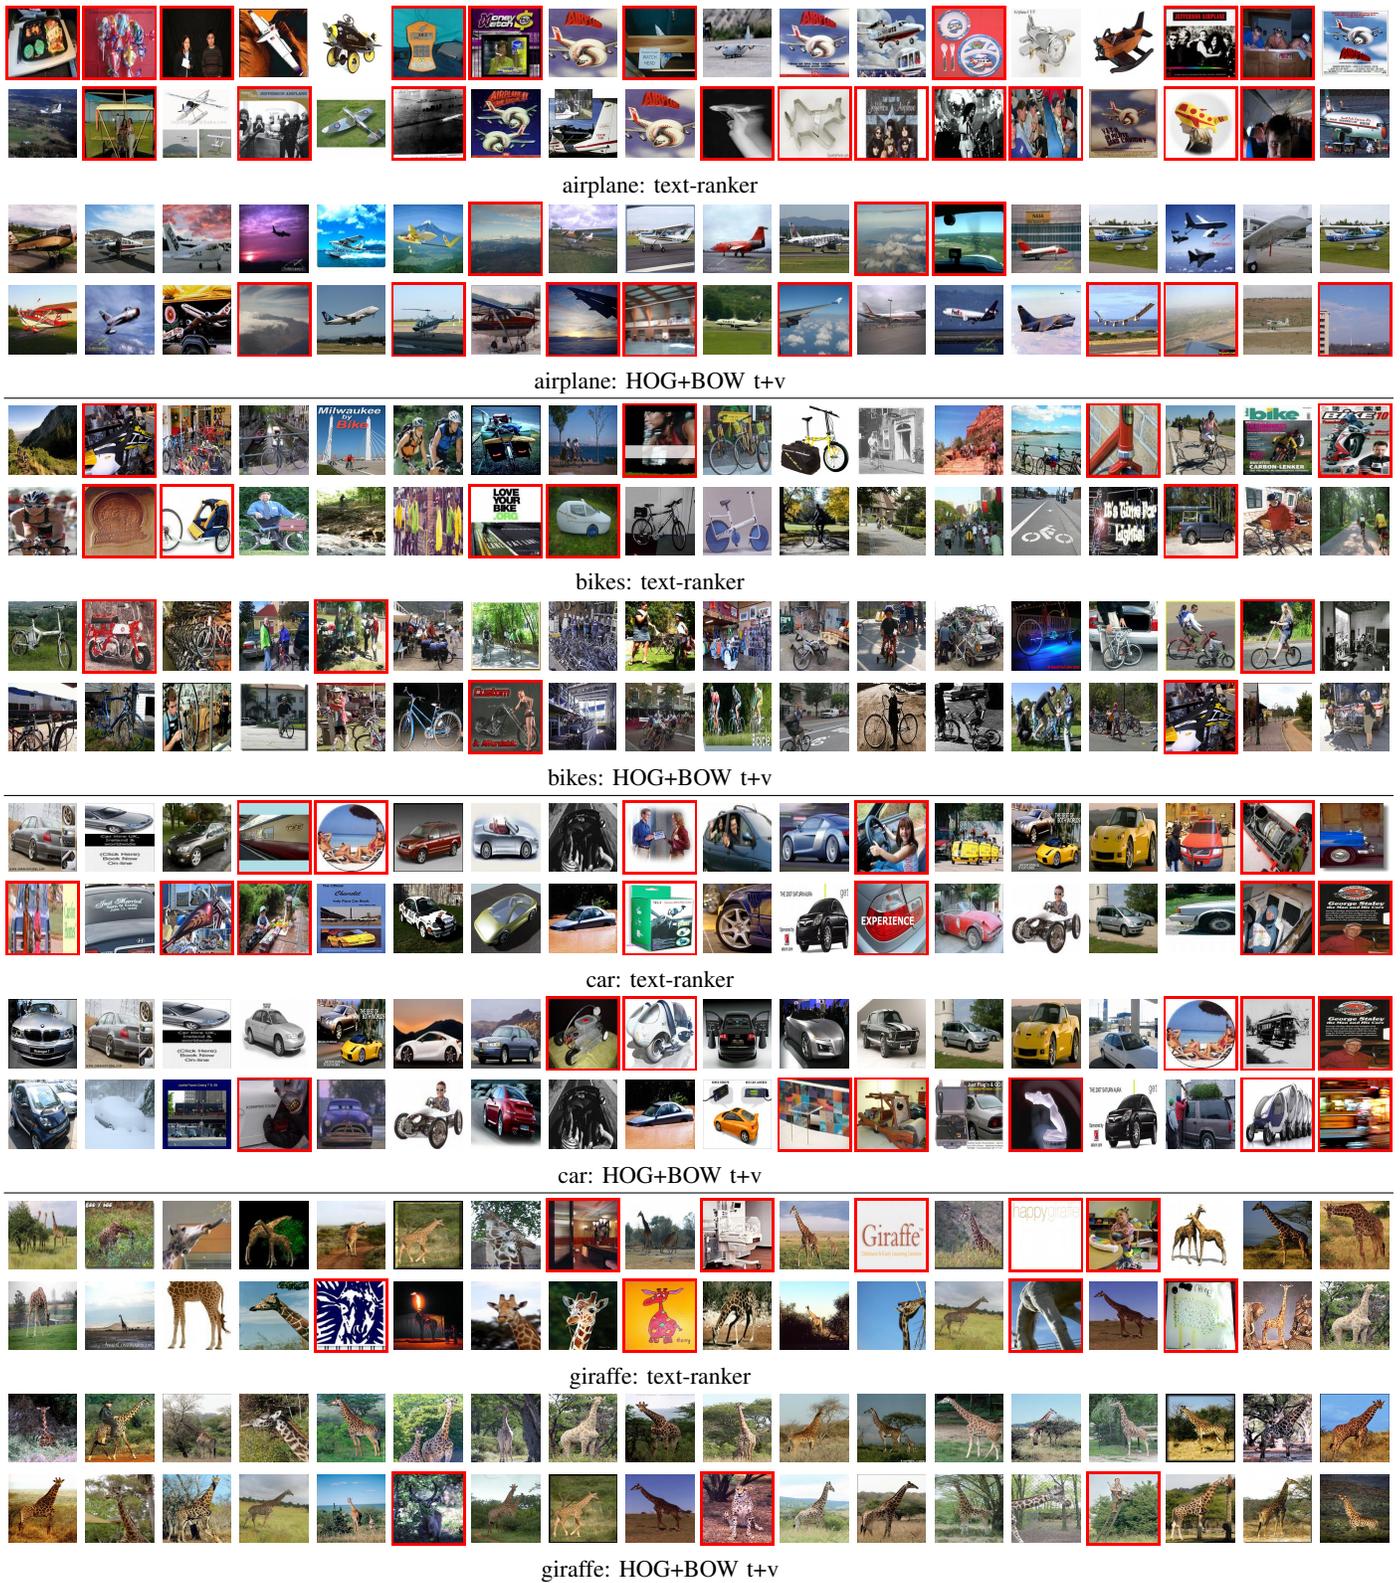


Fig. 8. Comparing top ranked 36 images using the text-ranker, and the full system. Red boxes indicate false positives.

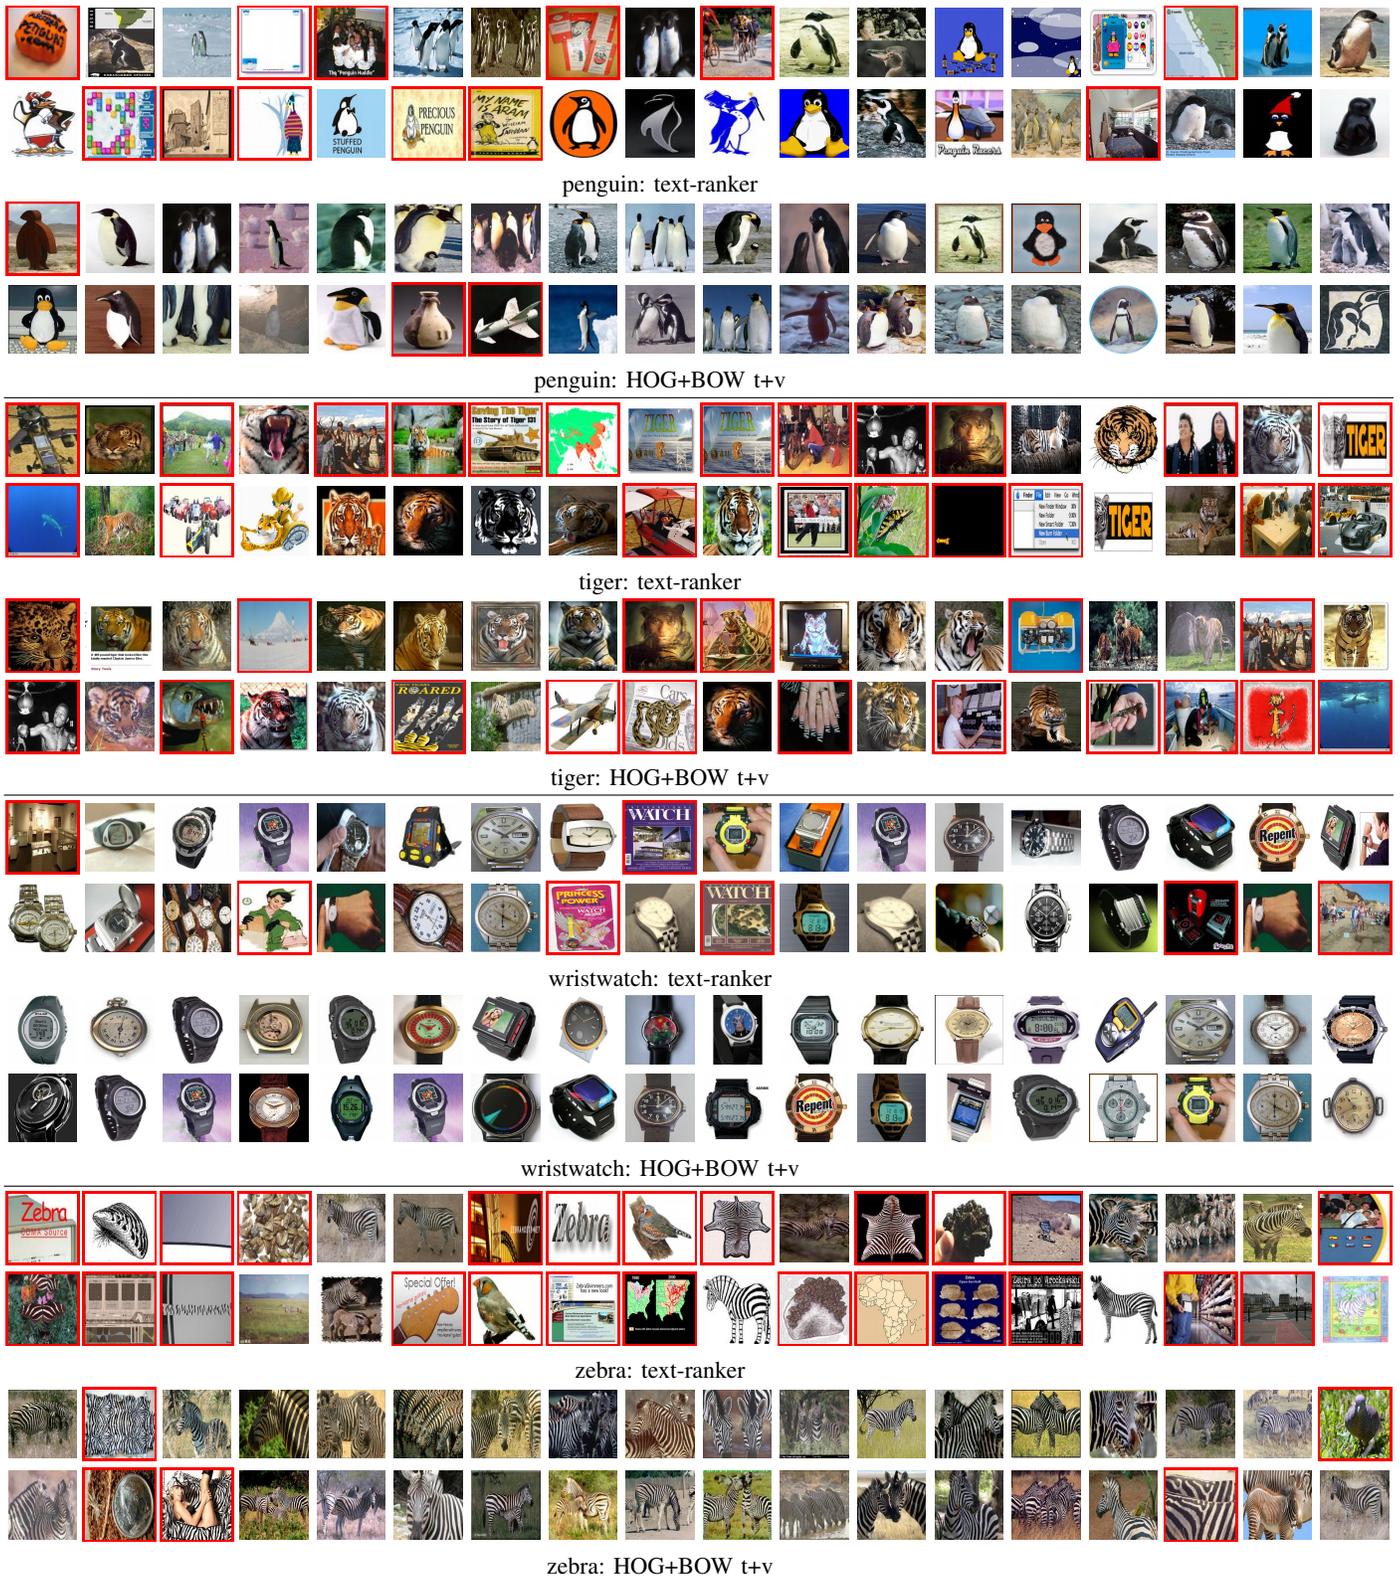


Fig. 9. Comparing top ranked 36 images using the text-ranker, and the full system. Red boxes indicate false positives.

prec. 15%	ap	bv	bk	bt	cm	cr	dp	ep	gf	gr	hs	kg	mb	pg	sk	tr	ww	zb	avg.	std.
tf rf	42.8	46.4	70.4	60.9	54.5	93.5	67.7	84.8	88.3	62.5	68.8	55.2	72.1	78.9	80.7	57.3	89.7	81.4	69.8	8.1 (2.1)
tf ru	38.7	39.0	68.7	62.3	55.5	93.4	63.6	81.8	84.4	64.3	68.9	53.0	74.7	80.0	79.1	58.2	92.5	80.2	68.8	6.6 (1.8)
tu ru	49.9	32.1	61.4	64.1	55.1	82.0	63.9	62.2	90.1	46.6	41.6	49.3	78.8	80.6	62.0	34.7	91.4	78.2	62.4	9.2 (2.5)
tf rf (B)	52.4	12.9	55.4	63.6	54.1	94.9	42.6	47.9	83.7	61.4	52.8	29.8	65.2	53.7	42.9	28.5	82.7	78.3	55.7	6.7 (1.7)

TABLE 5

Comparing filtered vs. non-filtered images and baseline: tf and (tu) denote training on filtered (unfiltered) images; rf and (ru) ranking on filtered (unfiltered) images; The first row gives the results for the whole system (HOG+BOW (t+v) in table 4). The second row shows the case where *all* images were re-ranked, including the ones that were filtered out in section 2.1. The third row shows training and ranking on unfiltered images. Last row gives results for the baseline (B) method where the visual classifier was trained on images selected from the initial database of that class, *i.e.* no text ranking applied, and tested on unfiltered images. The last column states the mean of the classwise standard deviations over five runs of cross-validation, as well as the standard deviation of the means over all classes, in parantheses.

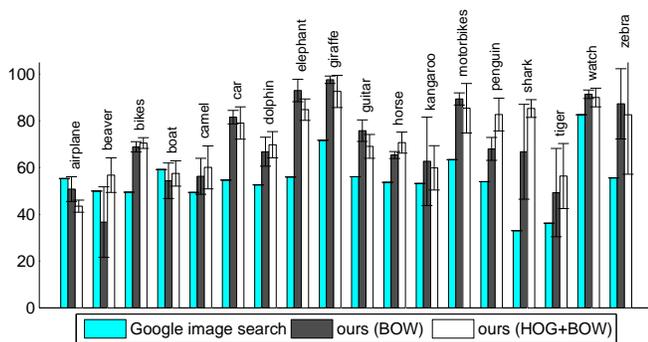


Fig. 10. **Comparison with Google image search.** Precision at 100 image recall. Compare to highlighted versions in first column of table 4 for *our* algorithm.

We found that combining the vision and text ranked lists using Borda count [1] or similar methods gave a slight improvement on average, but results were very class dependent. Combining the probabilistic outputs of text and SVM as in [34] remains an interesting addition for future work. The advantage of our system is that, once trained, we can rank images for which no meta-data is available, *e.g.* the Fergus and Berg images, see next section.

5 COMPARISON WITH OTHER WORK & METHODS

We compare *our* algorithm with three other approaches. We report the results for $n_+ = 150$, $n_- = 1000$ and HOG+BOW (see highlighted versions in first column of table 4). Again we report mean and standard deviation over five runs of ten-fold cross-validation.

5.1 Comparing with other results

Comparison with Google image search:

Here we re-rank the images downloaded with GoogleImages with our fully automatic system, *i.e.* text-based training of visual classifier followed by visual re-ranking of GoogleImages. Comparative results between our approach and GoogleImages are shown in figure 10.

	airplane	guitar	leopard	motorbike	wristwatch
our	58.5 ± 25.8	70.0 ± 3.9	49.6 ± 9.9	74.8 ± 7.3	98.1 ± 3.0
our (ϕ_k)	41.8 ± 18.4	31.7 ± 4.1	22.6 ± 10.2	50.5 ± 8.0	94.4 ± 3.2
[12] (ϕ_k)	57	50	59	71	88
Google (ϕ_k)	50	30	41	46	70
Fritz [14]	100	91	65	97	100

TABLE 6

Comparison with Fergus *et al.* [12] data: Average precision at 15% recall with one standard deviation. The images are from Google image search and were provided by the author. (ϕ_k) uses the same annotation as [12]. The first row of the table treats Fergus’ *ok* class as *in-class*, unlike [12].

As can be observed, our approach achieves higher average precision for 16 out of 18 classes with only airplane and boat being outperformed by the Google results. Those are cases where our text ranker doesn’t perform that well, which increases the noise in the training data and thus explains the decreased visual performance.

Images for the class penguin returned from both Google and our re-ranking (using the learnt visual classifier applied to the Google images) are shown in figure 11.

Comparison on Fergus *et al.* [12] data:

In this experiment we re-rank the Google images provided by [12]. In order to do so we train on the downloaded images and their corresponding meta-data. We downloaded “leopard” in addition to the previously described classes. It is difficult to directly compare results in [12] to our text+vision algorithm, as [12] treats *ok* images as non-class, whereas our system is not tuned to distinguish *good* from *ok* images. Due to this our system performs slightly worse than [12], when measured only on *good* images. However, it still outperforms Google image search on most classes even in this case. The same holds for the comparison with [14], who perform very well on this dataset. Table 6 also shows (first row) the results when *ok* images from the [12] data are treated as *in-class*. As expected the performance increases significantly for all classes. It needs to be noted that this dataset is relatively small and the distribution of images is not always diverse (*e.g.* many

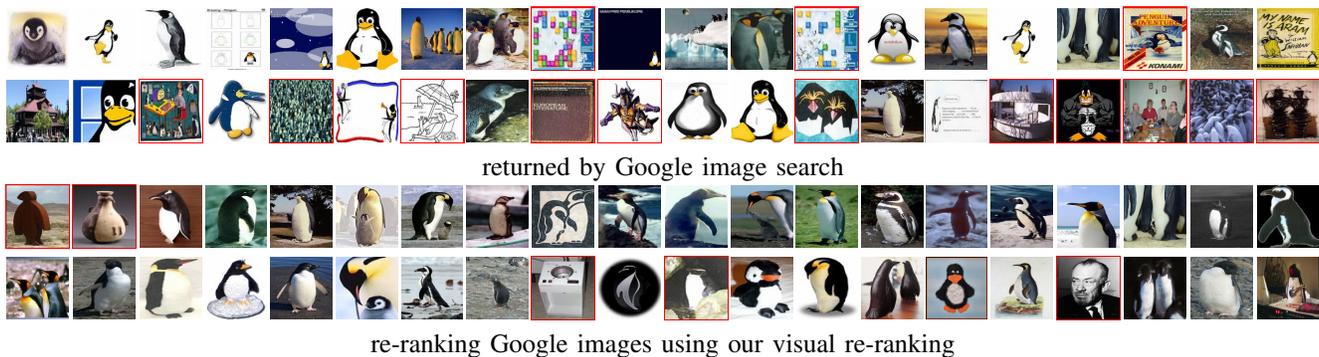


Fig. 11. **Re-ranking Google images.** The 40 top-ranked penguin images returned by Google image search (in their original order), as well as the 40 top-ranked images after re-ranking Google images using our visual (only) ranking. Red boxes indicate false positives. Google image search returns 14 *non-class* images (false positives), to be compared to the 6 returned by our system.

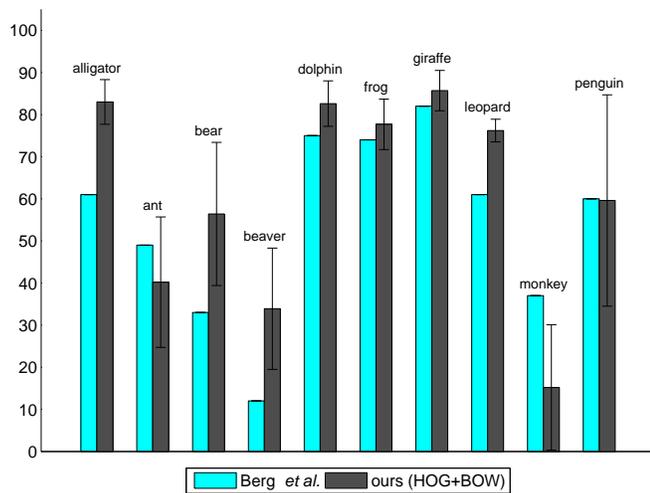


Fig. 12. **Comparison with Berg et al. [5].** Precision at 100-image recall level for the 10 animal classes made available by the authors of [5]. Note that our automatic algorithm is superior in many cases, even though the method of [5] involves manual intervention.

close duplicates in the airplane category).

Comparison with Berg et al. [5]: Here we run our visual ranking system on the dataset provided by [5]. In order to do so we downloaded an additional set of six classes (alligator, ant, bear, frog, leopard, monkey) for which no manual annotation was obtained. In our experiments we only use the test images from [5] for testing and training is performed on our downloaded images and the corresponding meta-data. Figure 12 compares the results reported in [5] to re-ranking of the test images available from [3] using our visual classifier. Note that we are training on our set of images which might stem from a different distribution than the Berg test set. We compare with the “classification on test data” category of [5], not to their “final dataset” which includes groundtruth from their manual step. Their provided groundtruth, which treats *abstract* images as *non-class*, was used. Note that our automatic algorithm produces results comparable or superior to those of Berg et

al., although their algorithm requires manual intervention.

5.2 Comparison with topic models

Here we compare our SVM based model to three widely used topic models (pLSA ([12],[16]), LDA [6], and HDP [30]). The underlying idea of all three topic models is to model each image as a distribution over topics, whereby each topic is a distribution over textons. Probabilistic Latent Semantic Analysis (pLSA) [16] models each entity (image) as a mixture of topics. Using the common notation for pLSA we have $P(w|d) = \sum_z P(w|z)P(z|d)$. In our case w are the visual words (described in section 4), d are the images and z are the classes (see figure 13). LDA [6] and HDP [30] are “probabilistic” extensions of pLSA, and HDPs model the data using hierarchical Dirichlet processes and avoid the need to explicitly specify the number of desired topics.

Given a set of topics and the test images ranked by their topic likelihood the question is how to select the “right” topic which induces the final ranking. This problem of selecting the right topic is difficult and there exist various approaches (see for example [11]). To give an idea of the performance of the topic models we avoid this topic selection stage and select the best topic based on groundtruth. The topic that gives the best 15% recall on the images evaluated on groundtruth is selected. This gives a huge advantage to the topic models in table 7, as groundtruth is used during testing. Each of these models is trained on all images that we want to rank later on. This results in a ranking of all images for each topic (number of topics specified for pLSA and LDA and learned for HDP).

Discussion: The best performing topic models are the pLSA with 500 topics $68.7\% \pm 1.0$ closely followed by the HDP $68.6\% \pm 0.7$ which uses an average of 257 topics. Our SVM based system performs similarly well with $68.0\% \pm 2.0$ *without* using groundtruth for model selection. There is also no single topic model that consistently outperforms the SVM on the majority of classes. As mentioned before the results reported for the topic models use groundtruth to select the best topic, and thus give an unfair advantage to topic models, but they also give evidence to the conclusion that the SVM is the “stronger” classifier. It could be possible to combine the high ranked

prec. 15%	ap	bv	bk	bt	cm	cr	dp	ep	gf	gr	hs	kg	mb	pg	sk	tr	ww	zb	avg.	std
t+v (150/1000)	52.3	39.3	68.6	66.2	57.3	87.9	66.5	78.7	83.8	63.0	57.3	48.6	64.0	72.5	82.9	62.6	93.2	79.8	68.0	7.2 (2.0)
pLSA(5)	71.2	23.4	70.2	61.9	49.6	85.7	60.0	59.9	78.2	56.7	56.3	38.2	64.5	53.4	55.7	47.0	67.2	63.6	59.0	4.4 (1.1)
pLSA(10)	70.3	21.4	69.0	59.3	53.2	84.2	66.2	61.3	78.0	63.9	65.0	45.4	61.5	62.3	65.6	52.3	74.8	74.6	62.7	5.5 (1.4)
pLSA(50)	74.8	36.7	74.2	62.2	56.5	87.3	65.5	65.6	81.0	66.4	71.5	59.9	71.8	73.2	62.4	54.5	82.9	81.5	68.2	3.8 (0.9)
pLSA(100)	73.7	38.2	73.8	63.1	56.1	88.9	66.4	63.7	80.2	67.0	72.3	57.5	71.7	72.7	65.7	55.5	85.4	78.6	68.4	3.9 (0.9)
pLSA(200)	76.1	34.6	70.9	60.2	55.8	88.7	65.8	62.6	81.8	67.6	67.8	55.2	71.6	73.9	61.9	57.5	86.1	81.0	68.4	3.7 (0.9)
pLSA(500)	76.0	37.5	68.3	61.6	54.4	88.3	66.9	63.5	81.5	65.2	68.2	55.0	73.0	69.6	63.6	64.6	86.2	88.4	68.7	3.7 (1.0)
LDA(5)	70.6	22.4	65.9	60.9	46.0	84.2	62.0	55.8	77.9	57.3	55.7	42.7	63.7	54.2	52.9	52.9	72.5	59.8	58.7	4.5 (1.2)
LDA(10)	68.9	22.7	70.0	59.0	47.2	82.6	63.9	60.5	80.7	65.5	59.9	54.7	62.3	61.1	54.2	62.5	78.3	66.5	62.2	4.6 (1.1)
LDA(50)	72.4	37.8	74.3	64.4	56.8	89.3	64.1	63.8	83.6	66.7	66.6	62.6	69.7	68.7	50.3	60.5	86.3	82.6	67.8	3.7 (0.9)
LDA(100)	74.5	36.6	74.5	63.2	56.2	85.9	62.9	62.0	83.6	66.1	65.8	62.5	70.5	71.3	53.9	62.0	86.1	80.9	67.7	3.5 (0.9)
LDA(200)	72.9	38.9	69.7	65.1	57.1	86.7	61.5	59.4	82.9	65.4	63.8	61.3	72.2	68.9	48.4	60.8	86.0	83.8	66.9	2.8 (0.7)
LDA(500)	74.3	42.2	65.3	62.7	51.6	85.3	61.7	59.5	79.1	66.2	64.2	60.4	72.0	65.5	45.9	58.8	83.0	84.6	65.7	3.0 (0.7)
HDP	71.2	33.6	76.2	61.7	53.5	86.1	71.7	76.6	81.8	64.8	70.2	60.8	72.6	74.3	55.0	60.2	83.6	80.5	68.6	2.9 (0.7)

TABLE 7

SVM vs. topic models: Given are the results as in table 4 (*i.e.* 15% recall) for all object classes. We compare the best result achieved, using BOW only, by the SVM model to the performance achieved by using topic models (pLSA, LDA, HDP) where the best topic is *selected using groundtruth*. This gives the topic models a huge advantage, despite this the topic models don't clearly outperform the *fully automatic* SVM model. HDP used an average of 257 topics. The last column states the mean of the classwise standard deviations over 10 runs of training the topic models, as well as the standard deviation of the means over all classes, in parantheses.

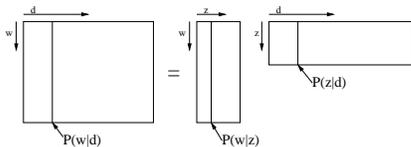


Fig. 13. pLSA models word frequencies in a document $P(w|d)$ as topics in a document $P(z|d)$ given the word distribution for a topic $P(w|z)$. pLSA then minimizes this equation with respect to KL-Divergence.

images from two topics, and thereby improve the ranking, although there is no straight forward method to do this.

6 CONCLUSION

This paper has proposed an *automatic* algorithm for harvesting the web and gathering hundreds of images of a given query class. Thorough quantitative evaluation has shown that the proposed algorithm performs similarly to state of the art systems such as [12], while outperforming both the widely used Google Image Search and recent techniques which rely on manual intervention [5].

Polysemy and diffuseness, are problems that are difficult to handle. This paper improves our understanding of the polysemy problem in its different forms. Interesting future direction could build on top of this understanding as well as the ideas in [26] and leverage multi-modal visual models to extract the different clusters of polysemous meanings, *i.e.* for tiger: Tiger Woods, the animal. It would also be interesting to divide diffuse categories described by the word airplane (airports, airplane interior, airplane food) into smaller visually distinctive categories. Recent work [26] addresses the polysemy problem directly and a combination with our work would be interesting.

Our algorithm does not rely on the high precision of top returned images *e.g.* from Google Image Search. Such images play a crucial role in [12, 19], and future work could take advantage of this precision by exploiting them as a validation set or by using them directly instead of the text based ranker to bootstrap the visual training.

There is a slight bias towards returning “simple” images, *i.e.* images where the objects constitute large parts of the image and are clearly recognizable. This is the case for object categories like car or wristwatch where an abundance of such images occurs in the top text-ranked images. For other object classes more difficult images are returned as well, *e.g.* elephant. The aim to return a more diverse set of images would require additional measures (see for example [19]). Although some classification methods might require difficult images, [27] gives an example of how a car model can be learnt from these images. This automatically learnt model is able to segment cars in unseen images.

ACKNOWLEDGMENTS

We are grateful for financial support from Microsoft Research through a European PhD scholarship, EU Project CLASS, ERC grant VisRec no. 228180, and the Royal Academy of Engineering.

REFERENCES

- [1] J. Aslam and M. Montague. Models for metasearch. In *Special Interest Group on Information Retrieval*, pages 276–284, New York, NY, USA, 2001. ACM Press. 11
- [2] K. Barnard, P. Duygulu, N. de Freitas, D. Forsyth, D. Blei, and M. Jordan. Matching words and pictures. *Journal of Machine Learning Research*, 3:1107–1135, February 2003. 1
- [3] T. Berg. Animals on the web dataset, 2006. URL <http://www.tamaraberg.com/animalDataset/index.html>. 12
- [4] T. Berg, A. Berg, J. Edwards, M. Mair, R. White, Y. Teh, E. Learned-Miller, and D. Forsyth. Names and Faces in the News. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Washington, DC, 2004*. 2

- [5] T. L. Berg and D. A. Forsyth. Animals on the web. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, New York*, 2006. 1, 2, 6, 12, 13
- [6] D. Blei, A. Ng, and M. Jordan. Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, January 2003. 1, 12
- [7] C. K. Chow and C. N. Liu. Approximating discrete probability distributions with dependence trees. *IEEE Information Theory*, 14, 1968. 4
- [8] B. Collins, J. Deng, K. Li, and Fei-Fei L. Towards scalable dataset construction: An active learning approach. In *Proceedings of the 10th European Conference on Computer Vision, Marseille, France*, 2008. 1
- [9] N. Dalal and B. Triggs. Histogram of Oriented Gradients for Human Detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, San Diego*, volume 2, pages 886–893, 2005. 6
- [10] G. Dorkó and C. Schmid. Selection of scale-invariant parts for object class recognition. In *Proceedings of the 9th International Conference on Computer Vision, Nice, France*, 2003. 8
- [11] R. Fergus, P. Perona, and A. Zisserman. A visual category filter for Google images. In *Proceedings of the 8th European Conference on Computer Vision, Prague, Czech Republic*. Springer-Verlag, May 2004. 1, 12
- [12] R. Fergus, L. Fei-Fei, P. Perona, and A. Zisserman. Learning object categories from Google’s image search. In *Proceedings of the 10th International Conference on Computer Vision, Beijing, China*, 2005. 1, 2, 6, 11, 12, 13
- [13] C. Frankel, M. J. Swain, and Athitsos V. Webseer: An image search engine for the world wide web. Technical report, University of Chicago, 1997. 4
- [14] M. Fritz and B. Schiele. Decomposition, discovery and detection of visual categories using topic models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Anchorage, Alaska*, 2008. 1, 6, 11
- [15] T. Hofmann. Probabilistic latent semantic analysis. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, 1999. 1
- [16] T. Hofmann. Unsupervised learning by probabilistic latent semantic analysis. *Machine Learning*, 43:177–196, 2001. 12
- [17] T. Joachims. SVM^{light}. URL <http://svmlight.joachims.org/>. 7
- [18] T. Kadir, A. Zisserman, and M. Brady. An affine invariant salient region detector. In *Proceedings of the 8th European Conference on Computer Vision, Prague, Czech Republic*. Springer-Verlag, May 2004. 6
- [19] J. Li, G. Wang, and L. Fei-Fei. OPTIMOL: automatic Object Picture collecTion via Incremental MOdel Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Minneapolis*, 2007. 1, 13
- [20] W.-H. Lin, R. Jin, and A. Hauptmann. Web Image Retrieval Re-Ranking with Relevance Model. In *Proceedings of the IADIS International Conference WWW/Internet*, 2003. 1
- [21] D. Lowe. Object recognition from local scale-invariant features. In *Proceedings of the 7th International Conference on Computer Vision, Kerkyra, Greece*, pages 1150–1157, September 1999. 6
- [22] K. Mikolajczyk, C. Schmid, and A. Zisserman. Human detection based on a probabilistic assembly of robust part detectors. In *Proceedings of the 8th European Conference on Computer Vision, Prague, Czech Republic*. Springer-Verlag, May 2004. 6
- [23] K. Morik, P. Brockhausen, and T. Joachims. Combining statistical learning with a knowledge-based approach - A case study in intensive care monitoring. In *Proceedings of the 16th International Conference on Machine Learning, Bled, Slovenia*, 1999. 7
- [24] Onix. ONIX Text Retrieval Toolkit. URL <http://www.lextek.com/manuals/onix/stopwords1.html>. 4
- [25] M. Porter, R. Boulton, and A. Macfarlane. The English (Porter2) stemming algorithm, 2002. URL <http://snowball.tartarus.org/>. 4
- [26] K. Saenko and T. Darrell. Unsupervised Learning of Visual Sense Models for Polysemous Words. In *Advances in Neural Information Processing Systems*, 2008. 13
- [27] F. Schroff. Semantic image segmentation and web-supervised visual learning. DPhil. Thesis, University of Oxford, 2009. 13
- [28] F. Schroff, A. Criminisi, and A. Zisserman. Harvesting Image Databases from the Web. In *Proceedings of the 11th International Conference on Computer Vision, Rio de Janeiro, Brazil*, 2007. 2
- [29] F. Schroff, A. Criminisi, and A. Zisserman. Harvesting Image Databases from the Web. <http://www.robots.ox.ac.uk/~vgg/data/mkdb>, 2007. 2
- [30] Y. Teh, M. Jordan, M. Beal, and D. Blei. Hierarchical dirichlet processes. Technical Report 653, 2003. 12
- [31] A. Torralba. Contextual priming for object detection. *International Journal of Computer Vision*, 53(2):153–167, 2003. 3
- [32] VGG. Affine Covariant Features. URL <http://www.robots.ox.ac.uk/~vgg/research/affine/index.html>. 6
- [33] S. Vijayanarasimhan and K. Grauman. Keywords to Visual Categories: Multiple-Instance Learning for Weakly Supervised Object Categorization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Anchorage, Alaska*, 2008. 1
- [34] G. Wang and D. Forsyth. Object image retrieval by exploiting online knowledge resources. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Anchorage, Alaska*, pages 1–8, 2008. 11
- [35] J. Zhang, M. Marszalek, S. Lazebnik, and Schmid C. Local features and kernels for classification of texture and object categories: a comprehensive study. *International Journal of Computer Vision*, 2007. 7