

Formulating Semantic Image Annotation as a Supervised Learning Problem

Gustavo Carneiro^{†,*}

Department of Computer Science[†]
University of British Columbia
Vancouver, BC, Canada

Nuno Vasconcelos^{*}

Department of Electrical and Computer Engineering^{*}
University of California, San Diego
San Diego, CA, USA

Abstract

We introduce a new method to automatically annotate and retrieve images using a vocabulary of image semantics. The novel contributions include a discriminant formulation of the problem, a multiple instance learning solution that enables the estimation of concept probability distributions without prior image segmentation, and a hierarchical description of the density of each image class that enables very efficient training. Compared to current methods of image annotation and retrieval, the one now proposed has significantly smaller time complexity and better recognition performance. Specifically, its recognition complexity is $O(CxR)$, where C is the number of classes (or image annotations) and R is the number of image regions, while the best results in the literature have complexity $O(TxR)$, where T is the number of training images. Since the number of classes grows substantially slower than that of training images, the proposed method scales better during training, and processes test images faster. This is illustrated through comparisons in terms of complexity, time, and recognition performance with current state-of-the-art methods.

1. Introduction

Content-based image retrieval, the problem of searching large image repositories according to their content, has been the subject of a significant amount of computer vision research in the recent past [13]. While early retrieval architectures were based on the query-by-example paradigm, which formulates image retrieval as the search for the best database match to a user-provided query image, it was quickly realized that the design of fully functional retrieval systems would require support for semantic queries [12]. These are systems where the database images are annotated with semantic keywords, enabling the user to specify the query through a natural language description of the visual concepts of interest. This realization, combined with the cost of manual image labeling, generated significant interest in the problem of automatically extracting semantic descriptors from images.

The two goals associated with this operation are: a) the automatic annotation of previously unseen images, and b) the retrieval of database images based on semantic queries.

Current systems achieve these goals by training a classifier that automatically labels an image with semantic keywords. This can be posed as either a problem of supervised or unsupervised learning. The earliest efforts focused on the supervised learning of binary classifiers using a set of training images with and without the semantic of interest [6, 14]. The classifier was then applied to the image database, and each image annotated with respect to the presence or absence of the concept. Since each classifier is trained in the “one vs all” (OVA) mode, we refer to this framework as *supervised OVA*. More recent efforts have been based on unsupervised learning [1, 2, 4, 5, 7, 8, 9], and strive to solve the problem in its full generality. The basic idea is to introduce a set of latent variables that encode hidden states of the world, where each state defines a joint distribution on the space of semantic keywords and image appearance descriptors (in the form of local features computed over image neighborhoods). During training, a set of labels is assigned to each image, the image is segmented into a collection of regions, and an unsupervised learning algorithm is run over the entire database to estimate the joint density of words and visual features. Given a new image to annotate, visual feature vectors are extracted, the joint probability model is instantiated with those feature vectors, state variables are marginalized, and a search for the set of labels that maximize the joint density of text and appearance is carried out. We refer to this framework as “unsupervised”.

Both formulations have strong advantages and disadvantages. Generally, unsupervised labeling leads to significantly more scalable (in database size and number of concepts of interest) training procedures, places much weaker demands on the quality of the manual annotations required to bootstrap learning, and produces a natural ranking of keywords for each new image to annotate. On the other hand, it does not explicitly treat semantics as image classes and, therefore, provides little guarantees that the semantic annotations are optimal in a recognition or retrieval sense. That is, instead of annotations that achieve the smallest probability of retrieval error, it simply produces the ones that have largest joint likelihood under the assumed mixture model.

In this work we show that it is possible to combine the advantages of the two formulations through a slight reformulation of the supervised one. This consists of defining an

M -ary classification problem where each of the semantic concepts of interest defines an image class. At annotation time, these classes all directly compete for the image to annotate, which no longer faces a sequence of independent binary tests. This *supervised M -ary* formulation obviously retains the classification and retrieval optimality of supervised OVA, but 1) produces a natural ordering of keywords at annotation time, and 2) eliminates the need to compute a “non-class” model for each of the semantic concepts of interest. In result, it has learning complexity equivalent to that of the unsupervised formulation and, like the latter, places much weaker requirements on the quality of manual labels than supervised OVA. The method now proposed is compared to the state-of-the-art methods of [5, 8] using the experimental setup introduced in [4]. The results show that the approach now proposed has advantages not only in terms of annotation and retrieval accuracy, but also in terms of efficiency.

2. Semantic Labeling

The goal of semantic image labeling is to, given an image \mathcal{I} , extract, from a vocabulary \mathcal{L} of semantic descriptors, the set of keywords, or captions, \mathbf{w} that best describes \mathcal{I} . Learning is based on a training set $\mathcal{D} = \{(\mathcal{I}_1, \mathbf{w}_1), \dots, (\mathcal{I}_D, \mathbf{w}_D)\}$ of image-caption pairs. The training set is said to be weakly labeled if the absence of a keyword from \mathbf{w}_i does not necessarily mean that the associated concept is not present in \mathcal{I}_i . This is usually the case given the subjectivity of the labeling task.

2.1. Supervised OVA Labeling

Let $\mathcal{L} = \{w_1, \dots, w_L\}$ be the vocabulary of semantic labels, or keywords, w_i . Under the supervised OVA formulation, labeling is formulated as a collection of L detection problems that determine the presence/absence of the concepts w_i in the image \mathcal{I} . Consider the i^{th} such problem and the random variable Y_i such that

$$Y_i = \begin{cases} 1, & \text{if } \mathcal{I} \text{ contains concept } w_i \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

Given a collection of image features \mathbf{X} extracted from \mathcal{I} , the goal is to infer the state of Y_i with smallest probability of error, for all i . This can be solved by application of standard Bayesian decision theory, namely by declaring the concept as present if

$$P_{\mathbf{X}|Y_i}(\mathbf{x}|1)P_{Y_i}(1) \geq P_{\mathbf{X}|Y_i}(\mathbf{x}|0)P_{Y_i}(0) \quad (2)$$

where $P_{\mathbf{X}|Y_i}(\mathbf{x}|j)$ is the class-conditional density and $P_{Y_i}(j)$ the prior probability for class $j \in \{0, 1\}$.

Training consists of assembling a training set \mathcal{D}_1 containing all images labeled with the concept w_i , a training set \mathcal{D}_0 containing the remaining images, and using some density estimation procedure to estimate $P_{\mathbf{X}|Y_i}(\mathbf{x}|j)$ from \mathcal{D}_j , $j \in \{0, 1\}$. Note that any images containing concept w_i which are not explicitly annotated with this concept are incorrectly assigned to \mathcal{D}_0 and can compromise the classification accuracy. In this sense, the supervise OVA formulation is not amenable to weak labeling. Furthermore, the

set \mathcal{D}_0 is likely to be quite large when the vocabulary size L is large and the training complexity is dominated by the complexity of learning the conditional density for $Y = 0$.

Applying (2) to the query image \mathcal{I} , produces a sequence of labels $\hat{w}_i \in \{0, 1\}$, $i \in \{1, \dots, L\}$, and a set of posterior probabilities $P_{Y_i|\mathbf{x}}(1|\mathbf{x})$ that can be taken as degrees of confidence on the annotation. Notice, however, that these are posterior probabilities relative to different classification problems and do not establish a natural ordering of importance of the keywords w_i as descriptors of \mathcal{I} . Nevertheless, the binary decision regarding whether each concept is present in the image or not is a minimum probability of error decision.

2.2. Unsupervised Labeling

The basic idea underlying the unsupervised learning formulation [1, 4, 2, 5, 7] is to introduce a variable L that encodes hidden states of the world. Each of these states then defines a joint distribution for keywords and image features. The various methods differ in the definition of the states of the hidden variable: some associate a state to each image in the database [5, 8], others associate them with image clusters [1, 4, 2]. The overall model is of the form

$$P_{\mathbf{X}, \mathbf{W}}(\mathbf{x}, \mathbf{w}) = \sum_{l=1}^S P_{\mathbf{X}, \mathbf{W}|L}(\mathbf{x}, \mathbf{w}|l)P_L(l) \quad (3)$$

where S is the number of possible states of L , \mathbf{X} the set of feature vectors extracted from \mathcal{I} and \mathbf{W} the vector of keywords associated with this image. Since this is a mixture model, learning is usually based on the expectation-maximization (EM) [3] algorithm, but the details depend on the particular definition of hidden variable and probabilistic model adopted for $P_{\mathbf{X}, \mathbf{W}}(\mathbf{x}, \mathbf{w})$.

The simplest model in this family [5, 8], which has also achieved the best results in experimental trials, makes each image in the training database a state of the latent variable, and assumes conditional independence between image features and keywords, i.e.

$$P_{\mathbf{X}, \mathbf{W}}(\mathbf{x}, \mathbf{w}) = \sum_{l=1}^D P_{\mathbf{X}|L}(\mathbf{x}|l)P_{\mathbf{W}|L}(\mathbf{w}|l)P_L(l) \quad (4)$$

where D is the training set size. This enables individual estimation of $P_{\mathbf{X}|L}(\mathbf{x}|l)$ and $P_{\mathbf{W}|L}(\mathbf{w}|l)$, as is common in the probabilistic retrieval literature [13], therefore eliminating the need to iterate the EM algorithm over the entire database (a procedure of large computational complexity). In this way, the training complexity is equivalent to that of learning the conditional densities for $Y_i = 1$ in the supervised OVA formulation. This is significantly smaller than the learning complexity of that formulation (which, as discussed above, is dominated by the much more demanding task of learning the conditionals for $Y_i = 0$). The training of the $P_{\mathbf{W}|L}(\mathbf{w}|l)$, $l \in \{1, \dots, D\}$ consists of a maximum likelihood estimate based on the annotations associated with the l^{th} training image, and usually reduces to

counting [5, 8]. Note that, while the quality of the estimates improves when the image is annotated with all concepts that it includes, it is possible to compensate for missing labels by using standard Bayesian (regularized) estimates [5, 8]. Hence, the impact of weak labeling is not major under this formulation.

At annotation time, the feature vectors extracted from the query \mathcal{I} are used in (3) to obtain a function of \mathbf{w} that provides a natural ordering of the relevance of all possible captions for the query. This function can be the joint density of (3) or the posterior density

$$P_{\mathbf{W}|\mathbf{X}}(\mathbf{w}|\mathbf{x}) = \frac{P_{\mathbf{X},\mathbf{W}}(\mathbf{x}, \mathbf{w})}{P_{\mathbf{X}}(\mathbf{x})}. \quad (5)$$

Note that, while this can be interpreted as the Bayesian decision rule for a classification problem with the states of \mathbf{W} as classes, such class structure is not consistent with the generative model of (3) which enforces a causal relationship from L to \mathbf{W} . This leads to a very weak dependency between the observation \mathbf{X} and class \mathbf{W} variables, e.g. that they are independent given L in the model of (4). Therefore, in our view, this formulation imposes a mismatch between the class structure used for the purposes of designing the probabilistic models (where the states of the hidden variable are the dominant classes) and that used for labeling (which assume the states of \mathbf{W} to be the real classes). This implies that the annotation decisions are not optimal in a minimum probability of error sense.

3. Supervised M -ary Labeling

The supervised M -ary formulation now proposed explicitly makes the elements of the semantic vocabulary the classes of the M -ary classification problem. That is, by introducing 1) a random variable W , which takes values in $\{1, \dots, L\}$, so that $W = i$ if and only if \mathbf{x} is a sample from the concept w_i , and 2) a set of class-conditional distributions $P_{\mathbf{X}|W}(\mathbf{x}|i)$, $i \in \{1, \dots, L\}$ for the distribution visual features given the semantic class. Similarly to supervised OVA, the goal is to infer the state of W with smallest probability of error. Given a set of features \mathbf{x} from a query image \mathcal{I} this is accomplished by application of the Bayes decision rule

$$i^* = \arg \max_i P_{\mathbf{X}|W}(\mathbf{x}|i)P_W(i) \quad (6)$$

where $P_W(i)$ is a prior probability for the i^{th} semantic class. The difference with respect to the OVA formulation is that instead of a sequence of L binary detection problems, we now have a single M -ary problem with L classes.

This has several advantages. First, there is no longer a need to estimate L non-class distributions ($Y_i = 0$ in (1)), an operation which, as discussed above, is the computational bottleneck of the OVA formulation. On the contrary, as will be shown in Section 4, it is possible to estimate all semantic densities $P_{\mathbf{X}|W}(\mathbf{x}|i)$ with computation equivalent to that required to estimate one density per image. Hence, the supervised M -ary formulation has learning complexity equivalent to the simpler of the unsupervised labeling approaches (4).

Second, the i^{th} semantic class density is estimated from a training set \mathcal{D}_i containing all feature vectors extracted from images labeled with concept w_i . While this will be most accurate if all images that contain the concept include w_i in their captions, images for which this keyword is missing will simply not be considered. If the number of images correctly annotated is large, this is likely not to make any practical difference. If that number is small, missing labeled images can always be compensated for by adopting Bayesian (regularized) estimates. In this sense, the supervised M -ary formulation is equivalent to the unsupervised formulation and, unlike the supervised OVA formulation, not severely affected by weak labeling.

Finally, at annotation time, the supervised M -ary formulation provides a natural ordering of the semantic classes, by the posterior probability $P_{\mathbf{W}|\mathbf{X}}(w|\mathbf{x})$. Unlike the OVA case, under the M -ary formulation these posteriors are relative to the same classification problem, a problem where the semantic classes compete to explain the query. This ordering is, in fact, equivalent to that adopted by the unsupervised learning formulation (5), but now leads to a Bayesian decision rule that is matched to the class structure of the underlying generative model. Hence, this concept ordering is optimal in a minimum probability of error sense.

4. Estimation of Semantic Class Distributions

Given the collection of semantic class-conditional densities $P_{\mathbf{W}|\mathbf{X}}(w|\mathbf{x})$, supervised M -ary labeling is relatively trivial (it consists of a search for the solution of (6)). Two interesting questions arise, however, in the context of density estimation. Hereafter, assume that \mathbf{x} consists of a feature vector extracted from an image region of small spatial support.

4.1. Modeling Classes Without Segmentation

So far, we have assumed that all samples in the training set \mathcal{D}_i are from concept w_i . In practice, however, this would require careful segmentation and labeling of all training images. While concepts such as “Indoor”, “Outdoor”, “Coastline”, or “Landscape” tend to be holistic (i.e. the entire image is, or is not, in the class), most concepts refer to objects and other items that only cover a part of any image (e.g. “Bear”, “Flag”, etc.). Hence, most images contain a combination of various concepts. The creation of a training set \mathcal{D}_i of feature vectors exclusively drawn from the i^{th} class would require manual segmentation of all training images, followed by labeling of the individual segments.

Since this is unfeasible, an interesting question is whether it is possible to estimate the class-conditional density from a training set composed of images with a significant percentage of feature vectors drawn from other classes. The answer to this question is affirmative, it is the basis of so-called *multiple instance* learning [10], where each image is labeled positive with respect to a specific concept w if at least one of its regions \mathcal{R} is an exemplar of this concept,

and labeled negative otherwise. While no explicit correspondence between regions and concepts is included in the training set, it is still possible to learn a probability distribution for a concept by exploiting the consistent appearance of samples from this distribution in all images of the concept. To see this, let \mathcal{R}_w be the region of the feature space populated by the concept w that appears in all positive images. Assume, further, that the remaining samples are uniformly distributed. Since the probability of this uniform component must integrate to one, it must necessarily have small amplitude. Hence, the probability density of the image ensemble is dominated by the probability mass in \mathcal{R}_w . As the number of images goes to infinity, this property holds independently of how small the probability mass in \mathcal{R}_w is for each image.

4.2. Density Estimation

Given the training set \mathcal{D}_i of images containing concept w_i , the estimation of the density $P_{\mathbf{X}|W}(\mathbf{x}|i)$ can proceed in four different ways: *direct estimation*, *model averaging*, *naive averaging*, *hierarchical estimation*.

Direct Estimation: direct estimation consists of estimating the class density from a training set containing all feature vectors from all images in \mathcal{D} . The main disadvantage of this strategy is that, for classes with a sizable number of images, the training set is likely to be quite large. This creates a number of practical problems, e.g. the requirement for large amounts of memory, and makes sophisticated density estimation techniques unfeasible. One solution is to discard part of the data, but this is suboptimal in the sense that important training cases may be lost.

Model Averaging: model averaging performs the estimation of $P_{\mathbf{X}|W}(\mathbf{x}|i)$ in two steps. In the first step, a density estimate is produced for each image, originating a sequence $P_{\mathbf{X}|L,W}(\mathbf{x}|l,i)$, $l \in \{1, \dots, D\}$ where L is a hidden variable that indicates the image number. The class density is then obtained by averaging the densities in this sequence

$$P_{\mathbf{X}|W}(\mathbf{x}|i) = \frac{1}{D} \sum_l P_{\mathbf{X}|L,W}(\mathbf{x}|l,i). \quad (7)$$

Note that this is equivalent to the density estimate obtained under the unsupervised labeling framework, if the text component of the joint density of (3) is marginalized and the hidden states are images (as is the case of (4)). The main difference is that, while under M -ary supervised labeling the averaging is done only over the set of images that belong to the semantic class, under unsupervised labeling it is done over the entire database. This, once again, reflects the lack of classification optimality of the later formulation.

The direct application of (7) is feasible when the densities $P_{\mathbf{X}|L,W}(\mathbf{x}|l,i)$ are defined over a (common) partition of the feature space. For example, if all densities are histograms defined on a partition of the feature space \mathcal{X} into Q cells $\{\mathcal{X}_q\}$, $q = 1, \dots, Q$, and $h_{i,j}^q$ the number of feature vectors from class i that land on cell \mathcal{X}_q for image j , then the average class histogram is simply

$$\hat{h}_i^q = \frac{1}{D} \sum_j h_{i,j}^q.$$

However, when 1) the underlying partition is not the same for all histograms or 2) more sophisticated models (e.g. mixture or non-parametric density estimates) are used, model averaging is not as simple.

Naive Averaging: consider, for example, the Gauss mixture model

$$P_{\mathbf{X}|L,W}(\mathbf{x}|l,i) = \sum_k \pi_{i,l}^k \mathcal{G}(\mathbf{x}, \mu_{i,l}^k, \Sigma_{i,l}^k), \quad (8)$$

where $\pi_{i,l}^k$ is a probability mass function such that $\sum_k \pi_{i,l}^k = 1$. Direct application of (7) leads to

$$P_{\mathbf{X}|W}(\mathbf{x}|i) = \frac{1}{D} \sum_{k,l} \pi_{i,l}^k \mathcal{G}(\mathbf{x}, \mu_{i,l}^k, \Sigma_{i,l}^k) \quad (9)$$

i.e. a D -fold increase in the number of Gaussian components per mixture. Since, at annotation time, this probability has to be evaluated for each semantic class, it is clear that straightforward model averaging will lead to an extremely slow annotation process.

Mixture hierarchies: one efficient alternative to the complexity of model averaging is to adopt a hierarchical density estimation method first proposed in [15] for image indexing. This method is based on a mixture hierarchy where children densities consist of different combinations of subsets of the parents components. A formal definition is given in [15], we omit the details for brevity. The important point is that, when the densities conform to the mixture hierarchy model, it is possible to estimate the parameters of the class mixture directly from those available for the individual image mixtures, using a two-stage procedure. The first stage, is the naive averaging of (9). Assuming that each mixture has K components, this leads to an overall mixture with DK components of parameters

$$\{\pi_j^k, \mu_j^k, \Sigma_j^k\}, j = 1, \dots, D, k = 1, \dots, K. \quad (10)$$

The second is an extension of the EM algorithm, which clusters the Gaussian components into a T -component mixture, where T is the number of components at the class level. Denoting by $\{\pi_c^t, \mu_c^t, \Sigma_c^t\}$, $t = 1, \dots, T$ the parameters of the class mixture, this algorithm iterates between the following steps.

E-step: compute

$$h_{jk}^t = \frac{\left[\mathcal{G}(\mu_j^k, \mu_c^t, \Sigma_c^t) e^{-\frac{1}{2} \text{trace}\{(\Sigma_c^t)^{-1} \Sigma_j^k\}} \right]^{\pi_j^k N} \pi_c^t}{\sum_l \left[\mathcal{G}(\mu_j^k, \mu_c^l, \Sigma_c^l) e^{-\frac{1}{2} \text{trace}\{(\Sigma_c^l)^{-1} \Sigma_j^k\}} \right]^{\pi_j^k N} \pi_c^l}, \quad (11)$$

where N is a user-defined parameter (see [15] for details).

M-step: set

$$(\pi_c^t)^{new} = \frac{\sum_{jk} h_{jk}^t}{PK} \quad (12)$$

$$(\mu_c^t)^{new} = \sum_{jk} w_{jk}^t \mu_j^k, \text{ where } w_{jk}^t = \frac{h_{jk}^t \pi_j^k}{\sum_{jk} h_{jk}^t \pi_j^k} \quad (13)$$

$$(\Sigma_c^t)^{new} = \sum_{jk} w_{jk}^t [\Sigma_j^k + (\mu_j^k - \mu_c^t)(\mu_j^k - \mu_c^t)^T] \quad (14)$$

Notice that the number of parameters in each image mixture is orders of magnitude smaller than the number of feature vectors in the image itself. Hence the complexity of estimating the class mixture parameters is negligible when compared to that of estimating the individual mixture parameters for all images in the class. It follows that the overall training complexity is dominated by the latter task, i.e. only marginally superior to that of naive averaging and significantly smaller than that associated with direct estimation of class densities. On the other hand, the complexity of evaluating likelihoods is exactly the same as that achievable with direct estimation, and significantly smaller than that of naive averaging.

One final interesting property of the EM steps above is that they enforce a data-driven form of regularization which improves generalization. This regularization is visible in (14) where the variances on the left hand-side can never be smaller than those on the right-hand side. We have observed that, due to this property, hierarchical class density estimates are much more reliable than those obtained with direct learning.

5. Experimental Results

In this section, we present experimental results on a data set, Corel, that has been continuously adopted as a standard way to assess annotation and retrieval performance [4, 8, 5]. The Translation Model of [4] was the first milestone in the area of semantic annotation, in the sense of demonstrating results of practical interest. After various years of research, and several other contributions, the best existing results are, to our knowledge, those presented in [5]. We therefore adopt an evaluation strategy identical to that used in this work. The data set used in all experiments consists of 5,000 images from 50 Corel Stock Photo CDs, and was divided into two parts: a training set of 4,500 images and a test set of 500 images. Each CD includes 100 images of the same topic, and each image is associated with 1-5 keywords. Overall there are 371 keywords in the dataset. In all cases, the YBR color space was adopted, and the image features were coefficients of the 8×8 discrete cosine transform (DCT). Note that this is a feature set different than that used in [4, 8, 5], which consists of color, texture, and shape features.

5.1. Automatic Image Annotation

We start by assessing performance on the task of automatic image annotation. Given an un-annotated image, the task is to automatically generate a caption which is then compared to the annotation made by a human. Similarly to [8, 5] we define the automatic annotation to consist of the five classes under which the image has largest likelihood. We then compute the recall and precision of every word in the test set. Given a particular semantic descriptor w , if there are $|w_H|$ human annotated images with the descriptor w in the test set, and the system annotates $|w_{\text{auto}}|$ images

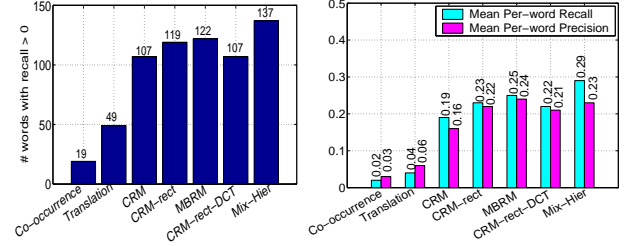


Figure 1. Performance comparison of automatic annotation on the Corel dataset.

with that descriptor, where $|w_C|$ are correct, recall and precision are given by $\text{recall} = \frac{|w_C|}{|w_H|}$, $\text{precision} = \frac{|w_C|}{|w_{\text{auto}}|}$.

Fig. 1 shows the results obtained on the complete set of 260 words that appear in the test set. The values of recall and precision were averaged over the set of testing words, as suggested by [8, 5]. Also presented are results (borrowed from [8, 5]) obtained with various other methods under this same experimental setting. Specifically, we consider: the Co-occurrence Model [11], the Translation Model [4], The Continuous-space Relevance Model (CRM-rect)[8, 5], and the Multiple-Bernoulli Relevance Model (MBRM) [5]. The method now proposed is denoted by 'Mix-Hier'. We also implemented the CRM-rect using the 8×8 DCT features, which is denoted as 'CRM-rect-DCT'.

Overall, the method now proposed achieves the best performance. When compared to the previous best results (MBRM) it exhibits a gain of 16% in recall for an equivalent level of precision. Similarly, the number of words with positive recall increases by 15%. It is also worth noting that the CRM-rect model with DCT features, performs slightly worse than the original CRM-rect. This indicates that the performance of Mix-Hier may improve with a better set of features. We intend to investigate this in the future.

Another important issue is the complexity of the annotation process. The complexity of CRM-rectangles and MBRM is $O(TR)$, where T is the number of training images and R is the number of image regions. Compared to those methods, Mix-Hier has a significantly smaller time complexity of $O(CR)$, where C is the number of classes (or image annotations). Assuming a fixed number of regions R , Fig. 2 shows how the annotation time of a test image grows for Mix-Hier and MBRM, as a function of the number of training images. In our experiments, over the set of 500 test images, the average annotation time was 268 seconds for Mix-Hier, and 371 seconds for CRM-rect-DCT.

5.2. Image Retrieval with Single Word Queries

In this section we analyze the performance of semantic retrieval. In this case, the precision and recall measures are computed as follows. If the n most similar images to a query are retrieved, recall is the percentage of all relevant images that are contained in that set and precision the percentage of the n which are relevant (where relevant means that the ground-truth annotation of the image contains the

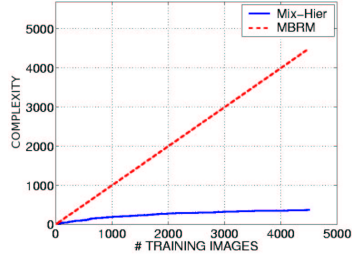


Figure 2. Time complexity for annotating a test image of the Corel data set.

Table 1. Retrieval results on Corel.

Mean Average Precision for the Corel Dataset		
Models	All 260 words	Words with recall > 0
Mix-Hier	0.31	0.49
MBRM	0.30	0.35

Image	Annotation
	<i>Human:</i> bear,polar, snow,tundra <i>Mix-Hier:</i> polar,tundra, bear,snow,ice
	<i>Human:</i> water,beach people,sunset <i>Mix-Hier:</i> sunset,sun, palm,clouds,sea
	<i>Human:</i> buildings,clothes, shops,streets <i>Mix-Hier:</i> buildings,street, shops,people,skyline

Figure 3. Comparisons of annotations made by our system and annotations made by a Human subject.

query descriptor). Once again, we adopted the experimental setup of [5]. Under this set-up, retrieval performance is evaluated by the mean average precision. As can be seen from Table 1, for ranked retrieval on Corel, Mix-Hier produces results superior to those of MBRM. In particular, it achieves a gain of 40% mean average precision on the set of words that have positive recall.

5.3. Results: Examples

In this section we present some examples of the annotations produced by our system. Fig. 3 illustrates the fact that, as reported in Table 1, Mix-Hier has a high level of recall. Frequently, when the system annotates an image with a descriptor not contained in the human-made caption, this annotation is not necessarily wrong. Finally, Fig. 4 illustrates the performance of the system on one word queries.




Figure 4. First five ranked results for the queries 'tiger' (first row) and 'mountain' (second row) in the Corel data set using our retrieval system.

Acknowledgments

We would like to thank Kobus Barnard for providing the Corel dataset used in [4].

References

- [1] K. Barnard et al. Matching words and pictures. *JMLR*, 3:1107–1135, 2003.
- [2] D. Blei and M. I. Jordan. Modeling annotated data. In *ACM SIGIR*, pages 127–134, 2003.
- [3] A. Dempster, N. Laird, and D. Rubin. Maximum-likelihood from Incomplete Data via the EM Algorithm. *J. of the Royal Statistical Society*, B-39, 1977.
- [4] P. Duygulu et al. Object recognition as machine translation: Learning a lexicon for a fixed image vocabulary. In *ECCV*, 2002.
- [5] S.L.Feng, R. Manmatha, and V. Lavrenko. Multiple bernoulli relevance models for image and video annotation. In *IEEE CVPR*, 2004.
- [6] D. Forsyth and M. Fleck. Body Plans. In *IEEE CVPR*, 1997.
- [7] H. Kuck, P. Carbonetto, and N. Freitas. A Constrained Semi-Supervised Learning Approach to Data Association. In *ECCV*, 2004.
- [8] V. Lavrenko, R. Manmatha, and J. Jeon. A model for learning the semantics of pictures. In *NIPS*, 2003.
- [9] J. Li and J.Z. Wang. Automatic linguistic indexing of pictures by a statistical modeling approach. *IEEE PAMI*, 25(10), 2003.
- [10] O. Maron and A. Ratan. Multiple instance learning for natural scene classification In *ICML*, 1998.
- [11] Y. Mori, H. Takahashi, and R. Oka. Image-to-word transformation based on dividing and vector quantizing images with words. In *First International Workshop on Multimedia Intelligent Storage and Retrieval Management*, 1999.
- [12] R. Picard. Digital Libraries: Meeting Place for High-Level and Low-Level Vision. In *ACCV*, 1995.
- [13] A. Smeulders et al. Content-based image retrieval: the end of the early years. *IEEE PAMI*, 22(12):1349–1380, 2000.
- [14] M. Szummer and R. Picard. Indoor-Outdoor Image Classification. In *Workshop in Content-based Access to Image and Video Databases*, 1998, Bombay, India.
- [15] N. Vasconcelos. Image Indexing with Mixture Hierarchies. In *IEEE CVPR*, 2001.