

Text-based approaches for non-topical image categorization

Carl L. Sable, Vasileios Hatzivassiloglou

Department of Computer Science, 450 Computer Science Building, Columbia University, 1214 Amsterdam Avenue, New York, NY 10027, USA; E-mail: {sable,vh}@cs.columbia.edu

Published online: 22 September 2000 – © Springer-Verlag 2000

Abstract. The rapid expansion of multimedia digital collections brings to the fore the need for classifying not only text documents but their embedded non-textual parts as well. We propose a model for basing classification of multimedia on broad, non-topical features, and show how information on targeted nearby pieces of text can be used to effectively classify photographs on a first such feature, distinguishing between indoor and outdoor images. We examine several variations to a TF*IDF-based approach for this task, empirically analyze their effects, and evaluate our system on a large collection of images from current news newsgroups. In addition, we investigate alternative classification and evaluation methods, and the effects that secondary features have on indoor/outdoor classification. Using density estimation over the raw TF*IDF values, we obtain a classification accuracy of 82%, a number that outperforms baseline estimates and earlier, image-based approaches, at least in the domain of news articles, and that nears the accuracy of humans who perform the same task with access to comparable information.

Key words: Image categorization – High-level image features – Text similarity features – Probabilistic TF*IDF – Evaluation in the presence of uncertainty

1 Introduction

As digital collections on the World Wide Web, corporate intranets, and CD-ROMs increase vastly in size and availability, it is becoming increasingly important to find efficient methods of categorizing not only text documents but also images, video, sound files, and other multimedia embedded within a document. Work in information retrieval has focused primarily on text, and then on classifying an entire document as relevant to a particular query

or as a member of a specific class. Yet, much is to be gained by independently categorizing and indexing pieces of a document from different media; multimedia information arguably follows a different classification hierarchy than text, and more factors than topical relevance come into play when an image or other non-text data is included within a document. For example, a news article on the recent events in Kosovo may include a picture of an airplane at a U.S. base, even though that particular aircraft never participated in the operations described in the article. The same image can frequently be found in multiple related documents, and, conversely, an independent classifier of images could help select an image from a broad, separate collection for illustrating a summary of a text-only source. Undesirable images (e.g., advertisements) could also be detected and pruned before a document is displayed to the user.

In the present work, we explore such an independent classification for images, using information from associated text sources such as captions and the surrounding text in the document in which the image is embedded. We are informed and motivated in this endeavor by the parallel development of a multimedia, multiple-document summarizer [1], where appropriate images can enhance the text summary. Our approach centers on the development of a suitable class hierarchy of broadly applicable visual features that will facilitate the selection of appropriate images for such summaries, even when fine distinctions (such as the subject matter of the image) are not available. Such features include classifying the images as indoor or outdoor; as containing one or a few persons or a crowd or no people at all; and as depicting a natural landscape versus a city scene. If independent classifiers can be designed for these features, then we can infer the appropriateness of the image for a particular descriptive purpose with high likelihood given only a little domain guidance. For example, an outdoor image with no peo-

ple from the terrorism news domain is likely to show the scene of an event or its aftermath, while an indoor photograph with a crowd of people probably refers to a related press conference. Additional techniques can refine these inferences by using, for example, information extraction methods [44] to identify the location of an event or the names of specific participants in the images.

The more traditional technique of searching the text associated with an image for relevant keywords (such as “meeting” or “press conference”) may not always be successful, given the variability and richness of natural language. For example, an examination of press conference photographs in the collection of images we describe in Sect. 3 shows that often the captions do not include the words “press conference” or its synonyms, but other related words such as “speaks”, “answers questions”, “faces reporters”, etc. Sometimes the text cannot be interpreted without significant inferencing on the part of a word-indexing system, as is the case with the headline “Greenspan best prescription for economy”, announcing Greenspan’s re-appointment as chairman of the Federal Reserve Bank. On the other hand, an approach that bases classification on broader features that can be reliably extracted can offer an alternative means for classifying cases where synonymy, polysemy, and paraphrases hamper traditional text-based classification.

We report in this paper on our methods and results for classifying images as indoor versus outdoor. We chose this visual feature as a basis for a first division of the images because of its plausibility as an indicator of image content and because it is used as a high-level feature in image ontologies for image and digital signal processing [42]. It is also a feature for which purely visual classifiers can be built [22, 40]; in fact, we are developing such classifiers in parallel with the text-based ones described here and investigating ways to integrate them. Although we have focused on this category, the methods described in the paper are independent of the specific feature and can be applied to any of the broad categories identified earlier.¹

Our indoor/outdoor classifier for images is based on information retrieval measures of text similarity, such as term frequency and inverse document frequency (TF*IDF) [28, 30]. Unlike information retrieval, however, we have to work with small pieces of text (a caption or a portion of a caption). Hence, we examined and evaluated several potential improvements to standard IR techniques, such as using targeted parts of the available text, limiting ourselves to particular word classes, and partially disambiguating words according to their part of speech. We collected a large sample of 1675 images for training and evaluation, and had multiple human volunteers assign indoor or outdoor labels to them. We measured individual human performance on this task

against the standard implied by their agreement, and compared our system’s performance to the humans, a default baseline classifier, and image-based classifiers that operate on purely visual features (e.g., color, texture, and edge direction features). We optimized our classifier using three-fold cross-validation, varying several of the TF*IDF parameters and optional features and determining which of the features have a major effect on performance. Using probability density estimates for the output of the classifier, we are able to correct several potential misclassification errors. Our results show that the automatic system outperforms the baseline and image-based classifiers, approaching the accuracy of the human volunteers, at least in our collection of news-related images. We extend these results by considering an experiment on an evenly distributed data set, an alternative evaluation method, and the effects of lenient versus strict definitions of the indoor and outdoor categories. We also explore another method for identifying words that discriminate between the two categories, we measure the effect of additional high-level features (in this case, the number of people in each image) for indoor/outdoor classification, and we summarize parallel research we are involved in together with image processing researchers at Columbia, on the development of an image-based system for the same task.

2 Related work

Our classification approach draws on a long line of work for measuring text similarity, mostly in an information retrieval context. Most of the information retrieval approaches rely on single words (e.g., [28, 30]), although sometimes compounds and collocations have been used [35]. Some of the features we explore (e.g., ignoring capitalization) are also used by default in most IR systems. Other, more natural language-informed features have found mixed success in information retrieval (e.g., [Salton and Smith 1989; Gay and Croft 1990; Smeaton 1992]), although the usefulness of each feature needs to be evaluated separately for each application. (Classifying images based on captions is different than classifying entire documents.)

For topical image classification, keywords extracted from a document have been used to index an associated image [2, 38], and similarity between images has been measured on the basis of shared image features [19, 23, 37] and by a combination of textual and image feature matches [20, 38]. Rowe and Guglielmo [26] and Smeaton and Quigley [36] use information from captions for retrieving (rather than classifying) images given a query. Harmandas et al. [10] combined evidence from text content and the hypertext structure of the Web for retrieving images from the Web. Srihari [39] uses face detection techniques along with name extraction from the captions for linking faces in images to specific people. Satoh et al. [32] use face detection techniques along with name ex-

¹ The main cost for moving on to new categories involves the necessary manual labeling of a large set of images for training and evaluation.

Table 1. Number of images assigned each combination of labels by human evaluators for the indoor/outdoor classification. The table is presented in symmetric form, by combining cases where evaluator *A* uses label *X* and evaluator *B* label *Y* with the cases where the two evaluators are reversed

First evaluator	Second evaluator				
	Indoor	Likely indoor	Ambiguous	Likely outdoor	Outdoor
Indoor	401	62	28	17	39
Likely indoor		12	12	15	17
Ambiguous			14	11	21
Likely outdoor				12	76
Outdoor					938

traction from closed captions for linking faces in video to specific people.

Classification of images along broad-based, non-topical features such as those we are exploring has received less attention in the image processing literature, although this is beginning to change. Forsyth and Fleck [7] present an image-based detector for naked people. Szummer and Picard [40] describe an approach for separating consumer photographs into indoor and outdoor classes. Both of these approaches utilize as their input only low-level visual features, such as color and edge direction. Schneiderman and Kanade [33] use local appearance and spatial relationships to detect faces or profiles. Oliva et al. [21] and Torralba and Oliva [41] use spectral templates to classify real-world scenes along broad semantic axes. Fung and Loe [8] use a mix of primitive semantics and scene semantics to classify images into scene classes such as rural, farm, beach, etc. Lipson et al. [18] present an approach which uses qualitative spatial and photometric relationships within and across regions in low resolution images to classify natural scenes.

3 Data set

Our raw data set consists of 21,086 news postings from April 1997 to May 1998 from a variety of Usenet current news newsgroups. Of these, 1490 contain, in addition to a text article, one or two embedded images, each with an associated caption. Due to some articles having two embedded images with captions, our collection contains a total of 1675 image/caption pairs. Captions are generally two to four sentences long. The first sentence in the caption tends to describe the image, while the remainder usually gives background information and establishes the relevance of the image to the story. For example:

BANGKOK, THAILAND, 9-NOV-1997: New Thai Prime Minister Chuan Leekpai gives a traditional “wai” to thank members of his party applauding his entrance, November 9, during a ceremony appointing him as the country’s 23rd prime minister in Bangkok,

Thailand. Chuan was named prime minister for the second time, replacing Chavalit Yongchaiyudh at the helm of a country plagued by economic woes.

For training and testing, a web-based interface was set up allowing volunteers to label images according to two high-level features. The first feature corresponded to the indoor versus outdoor dichotomy, and the choices given were *Indoor*, *Outdoor*, *Likely Indoor*, *Likely Outdoor*, and *Ambiguous*. The second feature was the number of people, and the available choices were *No People*, *One Person*, *Two People*, *Three or More People*, *Crowd*, and *Ambiguous*. In both cases, the authors went over a sample of images in advance, identified potential problems, and supplied the evaluators with detailed instructions which can be viewed at <http://www.cs.columbia.edu/~sable/research/readme.html>.²

Using our interface, fourteen volunteers labeled the images under different access conditions: by viewing the image alone, the caption alone, both the image and the caption, or just the first sentence of the caption. Each image was categorized by two different people under the full access condition (when volunteers viewed both the image and caption), which we consider representative of normal use of the images in multimedia documents. We use the labels obtained for this condition as the basis for both our training and testing sets. A single label for each of the indoor/outdoor and number of people questions was obtained for each image under the other conditions, and these were used to estimate human performance and for comparison with our system (which uses only text information). Table 1 shows the distribution of label pairs for the indoor/outdoor classification task.

For the indoor versus outdoor distinction, analysis of the assigned labels reveals that in most cases (87.7%),

² For several classification tasks, a large portion of the data will be classified as “other”, in a category containing cases to which the classification does not apply (e.g., if the classification involves different types of vehicles, non-vehicle images will fall into the “other” category). We assumed that such cases would be infrequent in the two classification tasks examined here, a position that is supported by the low numbers of images classified as “Ambiguous” or in different categories by different evaluators.

a definite indoor or outdoor judgement was made, and only 3% of labels assigned were “Ambiguous”. Agreement between humans was also high: in 1377 of the 1675 images (82.2%), the two evaluators assigned exactly the same label to an image, and 90.4% of the images had compatible labels, although sometimes with different degrees of confidence. There was, however, some disagreement between human categorizers. One hundred and thirty-seven images had labels that differed by more than one step on the scale from “definite indoor” to “definite outdoor”, and 39 of them had in fact one “definite indoor” and one “definite outdoor” label. Our analysis of the labels for the number of people feature indicated a somewhat lower but still significant level of agreement (80.4%). Inspection of the images that received conflicting labels reveals that several of the disagreements are due to mistakes by the categorizers, but in some cases, even markedly different labels can be attributed to different opinions about how terms like “indoor” and “outdoor” should be defined. For example, close-ups of people within a vehicle such as a car or a plane, or pictures of people under the roof of a structure with no walls, were often labeled differently by different judges. Figure 1 shows one of the images that reasonable people could disagree on; more can be inspected at <http://www.cs.columbia.edu/~sable/unusual.html>.

We have compiled five different sets of images according to these manual categorizations. First, we consider the images for which both evaluators provide a definite judgement in the same direction on the indoor versus outdoor question. This set contains 1339 images (79.9% of the original 1675) and is the primary focus of our experiments. 401 (29.9%) of these images were classified as indoor while 938 (70.1%) were classified as outdoor.

Our second experimental data set relaxes the requirement of strong beliefs from each evaluator. It consists of those images that received two judgements in the same direction on the indoor versus outdoor question, regardless of the reported degrees of confidence. This set includes 1501 images (89.6% of the original 1675). Four hundred



Fig. 1. An image that is hard to classify as indoor or outdoor

and seventy-five (31.6%) of them are classified as indoor while 1026 (68.4%) are classified as outdoor.

We also created a third data set with an equal number of indoor and outdoor images. This is a subset of the primary set defined above. It consists of the same 401 indoor images and also 401 of the outdoor images. See Sect. 6 for a further discussion of how the outdoor images were selected for this data set.

Turning to the number of people question, we define a fourth set, consisting of the images that received identical (non-ambiguous) judgements from both evaluators on that question. This set includes 1346 images (80.4% of the total), further divided as 88 (6.5%) with no people, 304 (22.6%) with one person, 213 (15.8%) with two people, 609 (45.2%) with three or more people, and 132 (9.8%) with crowds. We also define a fifth experimental set for studying the interaction between the indoor/outdoor and number of people categories, as the intersection of the first and fourth sets described above. This last set contains 1081 images (64.5% of the total).

4 Measuring similarity for indoor/outdoor classification

We base our classification of images into indoor or outdoor classes on a measure of similarity between each *document* we examine and the two *category prototypes* that correspond to the two classes. The term *document* is used above with a general sense, standing for any piece of text that is associated with the image under consideration; in many of our experimental runs, this is much smaller than the entire article that contains the image.

For a single piece of text, a word’s TF, or *term frequency*, is the number of times that this word occurs in that text. For a category (such as all indoor images), the TF assigned to a word is the number of times that word occurs in all documents of that category. A word’s IDF, or *inverse document frequency*, is the logarithm of the ratio of the total number of documents to the word’s DF, or *document frequency*, which is the number of documents that contain that word; this measure remains constant independently of the particular document or category examined. Their product, TF*IDF,

$$TFIDF(word) = TF(word) \times \log \frac{\text{Total number of documents}}{DF(word)} \quad (1)$$

is therefore highest when a word contains a balance of high frequency within a document or category (signifying high importance to the document or category) and low overall dispersion within the collection (signifying high specificity). Note that (1) applies to both single documents and categories; in each case, only the TF quantity changes in the formula.

Every document and category is represented by a vector of TF*IDF values, with each element corresponding

to a word. By abstracting content in this manner, word vectors of documents and categories can be compared to determine how well a document fits in each category. We use the inner product between document and category vectors, i.e.,

$$\text{Score}(\text{document}, \text{category}) = \sum_i \text{TFIDF}_{\text{document}}[i] \times \text{TFIDF}_{\text{category}}[i] \quad (2)$$

as our measure of similarity. Each document is then assigned to the category for which the fit is best, i.e., for which (2) is maximized.

We varied this measure of similarity in different experimental runs by using different restrictions on what enters the TF*IDF formula (i.e., what a “word” is) and by modifying (2) with the introduction of normalizing factors. Our first set of parameters, corresponding to the definition of words, involves four choices:

- **Text span considered.** What is the text that should be associated with each image, becoming the “document” in the TF*IDF calculations above? We have experimented by using the entire article, the article without the image caption, just the caption, or only the first sentence of the caption. While the articles are longer and provide more information about the related story than the caption, they are less related to the specific image, and therefore may contain too much noise to be helpful for the type of categorization we are performing. Hence, we can trade some information of questionable quality for increased specificity by limiting ourselves to the caption only. Similarly, the first sentence of the caption tends to be more descriptive of the image than the rest, which often provides background information.
 - **Restriction to specific grammatical categories.** Should all the words in the selected text span be included in the TF*IDF computations? Open-class words (i.e., adjectives, nouns, verbs, and adverbs) carry in general most of the content information, while words such as numbers and pronouns do not usually affect an image’s classification. We used a statistical part-of-speech (POS) tagger [4] to automatically assign a grammatical category tag to each word, and then experimented with using all words, only open-class words and prepositions (because of the nature of the indoor/outdoor distinction), and open-class words and prepositions with proper nouns excluded.
 - **Disambiguation of words.** A word’s sense is frequently ambiguous, and sometimes knowing its grammatical part-of-speech can help disambiguate it. For example, *can* is most often an auxiliary verb, but sometimes a noun with a different meaning. We experimented with keeping the POS tag as part of the word (thus distinguishing between the two senses of *can/modal* and *can/noun* above), versus ignoring this information.
 - **Case sensitivity.** Should capitalization matter for treating words as different? Capitalization may indicate a proper noun, but may also be the result of sentence-initial placement. We experimented with collapsing words that differ only in capitalization to the same token versus treating words as different if they differ in case.
- Each combination of the above parameters results in a different set of TF*IDF vectors for each document. Three more parameters were varied when calculating the similarity between a document and a category:
- **Ignoring words with low TF*IDF during similarity computations.** We have experimented with optionally ignoring words whose TF*IDF values within a document fall below a given constant, for several alternative values of that constant. This eliminates relatively insignificant words, which have minimal impact on the classification, while potentially speeding up the necessary calculations and avoiding some rare words whose TF and IDF are hard to estimate accurately.
 - **Normalization of category vectors.** The size of each of the two classes does not enter (2) or the TF*IDF calculations. Yet, it is natural to expect that the a priori most frequent category will have higher TF values, simply because it contains more documents. This is a concern for our experiments, since the “outdoor” category contains more than two-thirds of the images in our collection. We therefore experiment with a modification to (2), where the TF*IDF value of each word in a category vector is divided by the total number of documents that fall into that category. This modification, which replaces total frequency with average per document frequency, makes the TF*IDF values directly comparable across categories.
 - **Density estimation.** The standard approach for assigning documents to categories is to select the category for which similarity is largest. This, however, implicitly assumes that the similarity scores are on the same scale for both categories, and makes it hard to tell when a difference between the similarity scores for the two categories is large enough for the system to be confident in its decision. We experimented with a modification of the category decision rule by transforming the difference of the raw similarity scores between the two categories into the corresponding probability that a document with the given score difference belongs in the indoor category. In other words, we empirically estimate the probability density of the composite random variable $\text{Score}(\text{document}, \text{indoor}) - \text{Score}(\text{document}, \text{outdoor})$. We calculate the histogram of this difference function from the training part of the data (see the next section), and then use a rectangular smoothing window on top of the histogram to estimate the probability density [34]. For a new image in the test set, we again compute the difference and

apply the conversion procedure that was fixed during training. The resulting probability is more directly interpretable than the difference of the raw similarity scores, automatically adjusts the cut-off point between the two categories (from the arbitrary 0 on the unrestricted difference scale to the now well-justified 0.5 on the 0–1 probability scale), and provides a measure of confidence in the system’s decision (values near 0 or 1 indicate higher confidence) that can be easily combined with information from other independent categorizers.

5 Results and evaluation

We randomly selected 894 (approximately two-thirds) of the 1339 images that had definite human agreement on the indoor versus outdoor classification question for training, and the remaining 445 images for testing. Two hundred and seventy-six (30.9%) of the training images were indoor while 618 (69.1%) were outdoor. One hundred and twenty-five (28.1%) of the testing images were indoor while 320 (71.9%) were outdoor. So, on that particular breakdown of our main experimental image set, a default classifier would achieve 71.9% accuracy on the test set by labeling every image with the more frequent category in the training set.

Using this training/testing partition, we calculated the TF*IDF vectors and similarity scores described in the previous section for each of the 768 possible combinations of parameters, performing a complete designed experiment [13]. The training set was randomly divided into three equal parts, and for each such experiment, we repeatedly trained on two parts and measured system performance on the third. This three-fold cross validation on the training set gives us the ability to compare the relative performance of the various settings for the experimental parameters. It also allows us to select the best combination of parameters, which is fixed for subsequent experiments and, in particular, for scoring against the completely unseen test set.

We found a wide variety in the obtained average accuracy score (percentage of correct categorizations) depending on the parameter settings. The parameters which had the most major effect were:

- **Text span.** Restricting analysis to the first sentences of captions accounted for the 37 top scoring experiments. First sentences clearly outperformed captions, while text spans that included the entire article (with or without the caption) were far behind. This provides support to our thesis that specifically selected and narrowly targeted pieces of text can be more useful for classifying embedded multimedia information than the document as a whole.
- **Restriction to specific grammatical categories.** Using only open-class words plus prepositions accounted for four of the top five experiments. The average accuracy over all experiments for this setting was also higher than that for using all parts of speech, which, in turn, was higher than that using open-class words plus prepositions but excluding proper nouns. So it appears that proper nouns help in this classification task, a somewhat counter-intuitive result, especially since we generally have a high number of low-frequency proper nouns.
- **Normalization of category vectors.** Normalizing category vectors accounted for 12 of the top 15 experiments, and had a higher average accuracy among all experiments, even more so for cases where density estimates were used. Interestingly, however, if we only look at experiments using the standard TF*IDF method (without transforming to probability estimates), normalization according to category size has, on average, a slight detrimental effect.
- **Density estimation.** Using probability densities instead of raw similarity scores improved performance in almost every case, including all combinations of parameters ranked near the top. This optional component had one of the most pronounced effects in overall system performance.

On the other hand, ignoring words with low TF*IDF, keeping the part of speech information for disambiguation, and ignoring case differences played much smaller roles. High thresholds for including words in the TF*IDF vector were clearly bad, but other than that, all setting of these parameters were used in some of the best experiments, and the average accuracy for each were similar. Table 2 summarizes the effect of each value of each parameter over all experiments, while Table 3 shows the top fifteen combinations of parameters (those which achieved over 82.5% accuracy) in terms of performance during the three-fold cross validation on the training set. The average cross-validated accuracy of all 384 experiments that directly use the TF*IDF scores was 71.74%, and of the 384 experiments that include the probability conversions, 74.26%. Note that these overall accuracies are close to the baseline of the default classifier (71.9%), while 31 of the 768 combinations of parameters performed better than 82% during cross validation. This indicates that an informed choice of the parameters is important for this classification task.

On the other hand, Table 2 shows the average accuracy for a given setting of each parameter, across all possible combinations of values for the other parameters. Several of these combinations lead to low accuracy scores, lowering the average accuracy scores listed in the Table, to the extent that many of them are close or even lower than the baseline accuracy of 71.9%. Because of this averaging, the scores on Table 2 should not be interpreted as a judgment on the relevance or usefulness of a particular parameter, but as a means for comparing different settings for a given parameter.

Table 2. Average overall accuracy during cross-validation of all experiments with the given value of each parameter

Parameter	Value	Average accuracy
Text span	First sentences of captions	79.45%
	Captions	76.06%
	Articles (including captions)	69.22%
	Articles (excluding captions)	67.26%
Part of speech restriction	Open-class and prepositions	73.54%
	All words	73.09%
	Open-class and prepositions, excluding proper nouns	72.36%
Keeping tags for disambiguation	Yes	73.08%
	No	72.91%
Case sensitivity	Yes	73.01%
	No	72.99%
Threshold on TF*IDF	Medium	73.63%
	Low	73.57%
	None	73.21%
	High	71.57%
Normalization according to category size	Yes	73.36%
	No	72.64%
Using probability density estimates	Yes	74.26%
	No	71.74%

Results for all combinations of parameters are available in a Microsoft Excel spreadsheet located at http://www.cs.columbia.edu/~sable/research/all_results.xls. The reader is invited to experiment with a pivot table we have created in this spreadsheet (in the “Play with Fields” ply) which allows the dynamic computation of average accuracy for any subset of parameter combinations.

On the basis of these cross-validation experiments, we selected the following combination of parameters for our system: using the first sentences of captions only; restricting words to those of an open class plus prepositions; treating words that differ only in part of speech as identical; keeping capitalization information; not applying any thresholds for including words in the TF*IDF vector; normalizing according to category size; and applying the density transformation. These are the parameters that were used in the experiment represented by the first line of Table 3, which was one of two that tied for the best results during cross-validation. With these parameters fixed, we retrained on the full training set and tested on the unseen test set. The corresponding categorizer achieved 82.02% accuracy on the test set, and 90.72% on the training set.³ If the density estimate transformations were not employed, the accuracy on the test set falls dramatically to 72.36%. Tables 4 and 5 are contingency tables further breaking down these accuracy scores on a per category basis, separately for the cases where the density adjustments are used or not. Without density

estimation, the F_1 measures for the indoor and outdoor categories, defined as

$$F_1 = \frac{2 \times \textit{Precision} \times \textit{Recall}}{\textit{Precision} + \textit{Recall}}$$

[43], were 54.80% and 77.84% respectively. With density estimation, the F_1 measures for the indoor and outdoor categories were 65.22% and 87.88%, respectively.

Note that the use of probability densities tends to shift the system’s categorizations from the smaller category to the larger category. Therefore, the smaller category winds up having a higher precision and lower recall, while the larger category ends up with a lower precision and higher recall. Both categories have significantly higher F_1 measures. Detailed results on our 445 individual test images can be observed at http://www.cs.columbia.edu/~sable/research/demo_results/demo_results.cgi.

Naturally, we want to compare these results with alternative classifiers, including humans. Our accuracy on the test set (82.02%) clearly surpasses that of the default classifier which always selects the “outdoor” label for every image (71.9%). We estimate human performance on this task by measuring the percentage of correct classifications achieved by a human volunteer who looked only at the captions of the images (i.e., who had access to the same kind of information that our system does). Of the 1339 images in our main set, 1172 (87.52%) were correctly categorized under this access condition.⁴ This figure can

³ An indoor output probability of more than 50% is translated to a decision in favor of the indoor category during this evaluation.

⁴ For this purpose, any categorization in the right direction (i.e., indoor or outdoor), regardless of the degree of confidence, was

Table 3. Top fifteen combinations of TF*IDF experiment parameters after three-fold cross validation on the training set. The “tags” column indicates whether tags were kept for disambiguating words; the “case” column indicates whether word comparisons were case sensitive; and the “norm.” column indicates whether the normalization for category size was applied during the similarity calculations

Text span	Part of speech restriction	Tags	Case	Threshold on TF*IDF	Norm.	Accuracy without densities	Accuracy with densities
first sentences of captions	open-class plus prepositions	no	yes	none	yes	75.06%	83.22%
first sentences of captions	open-class plus prepositions	no	yes	low	yes	75.06%	83.22%
first sentences of captions	all words	yes	no	medium	yes	78.08%	82.89%
first sentences of captions	open-class plus prepositions	no	no	low	yes	74.83%	82.89%
first sentences of captions	open-class plus prepositions	no	no	none	yes	74.61%	82.89%
first sentences of captions	all words	no	no	medium	yes	79.08%	82.77%
first sentences of captions	open-class plus prepositions	no	yes	none	no	78.75%	82.77%
first sentences of captions	all words	yes	no	medium	no	78.97%	82.66%
first sentences of captions	all words	no	yes	low	yes	77.29%	82.66%
first sentences of captions	all words	no	no	low	yes	76.73%	82.66%
first sentences of captions	open-class plus prepositions	yes	no	low	yes	75.17%	82.66%
first sentences of captions	open-class plus prepositions	no	yes	medium	no	81.99%	82.55%
first sentences of captions	all words	no	yes	none	yes	77.40%	82.55%
first sentences of captions	all words	no	no	none	yes	77.07%	82.55%
first sentences of captions	all words	yes	no	none	yes	76.96%	82.55%

Table 4. Contingency table showing the breakdown of the system’s categorizations on the test set with conversions to probability densities

	Actual indoor	Actual outdoor	Precision
System indoor	75	30	71.43%
System outdoor	50	290	85.29%
Recall	60.00%	90.63%	

Table 5. Contingency table showing the breakdown of the system’s categorizations on the test set using the raw similarity scores

	Actual indoor	Actual outdoor	Precision
System indoor	106	104	40.48%
System outdoor	19	216	91.91%
Recall	84.80%	67.50%	

serve as a reasonable, approximate upper bound for how well we might hope our system to perform given only text information. On the other hand, humans who looked only at the image matched our reference standard in 95.47% of the cases, indicating that for humans, the image is a more accurate source for answering the indoor/outdoor question than the accompanying text.

Recently, an image-based approach for classifying photographs as indoor or outdoor has been proposed [40]. This approach is based on a decomposition of the image by applying a 4×4 grid on it and taking measures of low-level image features such as color and texture on each of the 16 image regions. Then, similarities between blocks in a given image and blocks in known indoor and outdoor images are calculated, and the image is assigned to one of the two categories. In cooperation with image processing researchers at Columbia,⁵ we reimplemented this tech-

considered correct while assignments of the “Ambiguous” label received half credit.

⁵ Seungyup Paek, Alejandro Jaimes, and Shih-Fu Chang, of the Department of Electrical Engineering, Columbia University.

Table 6. System accuracy stratified according to high, medium, or low confidence

Confidence level	Number correct	Number incorrect	Accuracy
$p \geq 0.9$ or $p \leq 0.1$	234	21	91.76%
$0.7 \leq p < 0.9$ or $0.1 < p \leq 0.3$	89	32	73.55%
$0.3 < p < 0.7$	42	27	60.87%
Total	365	80	82.02%

nique and measured its performance on our collection of photographs. We found that its accuracy on our test set was 74%, significantly less than what we obtain with our text-based methods. We also added supplemental low-level features, such as edge direction histograms, to those used by Szummer and Picard, and a machine learning component for estimating classification thresholds. With these improvements, the resulting image-based classifier achieves 76% performance, still less than the methods described in this paper.

For each of the above comparisons, we calculated a level of significance by applying Pearson’s chi-square test [6] on the contingency table that represents the cross-classification of the answers of the two compared methods.⁶ We observe that the difference between the performance of our system and either the default baseline, Szummer’s and Picard’s image-based classifier, or (regrettably) the human judges, is strongly significant at the 1% level or less; the probability that similar or more pronounced differences in the observed accuracy rates between the compared methods would be observed by chance is 0.046%, 0.464%, and 0.460%, respectively. When comparing our system to our extended implementation of the image-based model, the difference is still significant at the 5% level (P -value of 3.24%).

We should stress here that without an extensive evaluation over multiple domains and classification tasks, we cannot draw definite conclusions about which of the methods (ours, the one by Szummer and Picard, or our extension of their method) would perform better on a new domain or task. Szummer and Picard worked with vacation images and report significantly higher accuracy values for indoor/outdoor classification in their domain. It is possible that their method works better in the vacation image domain rather than the news articles domain.⁷ On the other hand, the comparison of the techniques in our data collection and domain indicates that at least for news images, the text-based approach leads to a more accurate indoor/outdoor classification.

⁶ The large-sample assumption of the chi-square test is satisfied for these contingency tables. Because we test on several hundreds of images, the exact Fisher’s test [5] is computationally impractical.

⁷ Szummer and Picard’s images do not come with captions, as they are photographs taken by vacationers, so we cannot apply our technique to their data.

A final evaluation question is how reliable the confidence estimates provided by our system’s output probabilities are. Preferably, decisions with a high degree of confidence should be more likely to be accurate than decisions given a low degree of confidence. We have therefore broken down the test set into three subsets according to the probability assigned by our system, p , that a given image is indoor. These three ranges of p were defined as high confidence ($p \geq 0.9$ or $p \leq 0.1$), medium confidence ($0.7 \leq p < 0.9$ or $0.1 < p \leq 0.3$), and low confidence ($0.3 < p < 0.7$). Note that the indoor probability equals 1 minus the outdoor probability, with the classifier selecting the indoor category when $p > 0.5$ and the outdoor category otherwise; hence, probabilities of p and $1 - p$ are equivalent in terms of the expressed confidence. Table 6 shows the accuracy of our system within each confidence category, and verifies that decisions given a higher level of confidence are more likely to be correct, thus validating our confidence estimates. In particular, 255 (57.3%) of the 445 test images were labeled with over 90% confidence, and 91.76% of these categorizations were correct.

6 Experimenting with an evenly distributed data set

In order to ensure that our system does not require a skewed distribution between the target classes to achieve good results, we tested it on the third data set defined in Sect. 3 (that with an equal number of indoor and outdoor images). To create this data set, we started with the training set and testing set from our primary data set. For each, we kept all of the indoor images and randomly selected an equal number of outdoor images. Therefore, the new training set is a subset of the old training set and the new testing set is a subset of the old testing set. The new training set contains 552 images (276 indoor and 276 outdoor) and the new testing set contains 250 images (125 indoor and 125 outdoor).

We re-trained and tested our system on this balanced data set both with and without the use of density estimation. Without density estimation the overall accuracy was 76.00% and with density estimation the overall accuracy was 76.40%. These accuracy scores are lower than those obtained for our main data set, perhaps because the uneven prior probabilities in the main data set aided our system or perhaps because this data set has a smaller

Table 7. Contingency table showing the breakdown of the system’s categorizations on the balanced test set with conversions to probability densities

	Actual indoor	Actual outdoor	Precision
System indoor	97	31	75.78%
System outdoor	28	94	77.05%
Recall	77.60%	75.20%	

Table 8. Contingency table showing the breakdown of the system’s categorizations on the balanced test set using the raw similarity scores

	Actual indoor	Actual outdoor	Precision
System indoor	103	38	73.05%
System outdoor	22	87	79.82%
Recall	82.40%	69.60%	

training set, but they are still quite reasonable and well above the new baseline of 50%. Tables 7 and 8 are contingency tables further breaking down these accuracy scores on a per category basis, separately for the cases where the density adjustments are used or not. Without density estimation, the F_1 measures for the indoor and outdoor categories were 77.44% and 74.36%, respectively. With density estimation, the F_1 measures for the indoor and outdoor categories were 76.68% and 76.11%, respectively. In general, density estimation tended to push the precision and recall values closer together for each category, but it did not make a significant overall improvement when the collection consisted of equally sized categories. This is of course consistent with the main motivation that led to the introduction of density estimation in our system in the first place: density estimation offers a means for normalizing the TF*IDF classification scores across categories, but when categories are equiprobable, these scores are on the same scale, leaving little for the density estimation step to correct.

7 Identifying words with high discriminating power

Methodology. A second approach to the classification problem is to automatically locate words (or multi-word phrases) whose presence strongly indicates one of the competing classes. We explore this technique by first extracting all open-class words plus prepositions from the first sentences of captions. We exclude proper nouns from this analysis since they are unlikely to be general indicators of one of the categories, and only consider words occurring five times or more in our training set. This last

step is done to ensure that the words we keep will be frequent enough to be general discriminators, and to avoid cases where a particular word occurs in a few captions of images from a particular class simply by chance.

We construct a *log-linear regression model* [31] using binary variables corresponding to the occurrence of each of these words as predictors and the output feature (e.g., indoor or outdoor image) as the response. The model is fitted with *iterative reweighted least squares* [3], and the fit assigns a weight to each of the candidate discriminators. Words with higher weights are those that actually help discriminate between the two classes.

As an alternative machine learning technique, we also consider *decision trees* [24]. The prediction model remains the same, but now the tree is constructed with *recursive partitioning*, with the most discriminating variable being selected first. The resulting tree is *shrunk* [11] (node probabilities are optimally regressed to their parents) to reduce the possibility of overfitting; we select the shrinking parameter α through cross-validation within the training set.

Results. Using the same training/test set division as with the TF*IDF experiments reported in the previous section, our list of candidate discriminators contains 665 words. Both the log-linear regression model and the tree select a subset of these words as classification features; in the case of the constructed tree, 80 words are used during classification.

It is interesting to note which these words are, especially since the results of this procedure are likely to generalize to other sets of images. The five words most favoring an indoor classification are **conference**, **meeting**, **meets**, **hands** (plural noun), and **L**, while the five words most strongly indicating an outdoor image are **of**, **from**, **soldiers**, **police**, and **demonstration**. Some of them are expected (e.g., *demonstration* or *police* for an outdoor image, or *conference* for an indoor one), but some come as a surprise, for example, the “words” *C*, *L*, and *R* (indicating an indoor image) used in parentheses to identify people in images by position (i.e., center, left, or right).

Overall performance of the word discriminant method was 93.62% over the training set and 78.65% over the test set.

Integrating the two classifiers. The two classifiers discussed in the present section and in Sect. 5 utilize different methods to arrive at similar classification performance. Hence, it is natural to investigate how correlated their answers are, and whether a combined classifier might improve overall performance on the indoor/outdoor classification task.

We have built such combined classifiers using both general machine learning techniques discussed above (log-linear models and decision trees). However, the overall performance of the composite classifiers was in both cases slightly less than that obtained by the best individual classifier (82.02%). We attribute this to overtraining

during the construction of the composite classifiers, especially since the same training set was used for building each of them and for combining them.⁸ Nevertheless, our implementation of two classification methods provides us with two different general tools that can be easily ported to other high-level classification tasks; and the ability to identify key discriminating words may prove helpful in future exploration of what makes images in distinct categories different.

8 An alternative evaluation metric

So far, all reported accuracies considered the system to be completely right if the category with the higher probability was correct and completely wrong otherwise. An alternative evaluation method is to take the probability assigned by the system to the correct category and consider that to be the system's score for that document, in a manner similar to the partial credits proposed in [12]. For example, let's say that the system analyzes an image and says the probability that the image is indoor is 65% (meaning that the probability that the image is outdoor is 35%). If the image is actually indoor, the system is given a score of 0.65 for this image, while if the image is actually outdoor, the system is given a score of 0.35. The overall accuracy of the system is then the sum of the system's scores for all images divided by the total number of images. In the ideal case, the system would assign all indoor images a probability of 1 of being indoor, and all outdoor images a probability of 0 of being indoor. Thus its total overall accuracy would be 1, or 100%. Indeed, if the system always has complete confidence in its decisions (whether correct or not), the revised evaluation method becomes equivalent to the standard one.

In this way, the system receives partial credit for each answer, more if the system leans in the correct direction and directly increasing as the system's confidence in a correct decision increases. In general, when a system already classifies most images correctly under the original 0/1 scoring method, it will tend to be penalized for its uncertainty on correct decisions more than it is credited for uncertain wrong answers. This is the case in our task when our classifier is evaluated on our main set of images (those with definite agreement by the human volunteers); the system achieves 82.02% accuracy under the original evaluation method, and 76.71% under the revised one. However, we consider this modified method as more revealing, as it offers a way to evaluate the system's confidence in its decisions; the accuracy computed using this method is equal to the average confidence given to the correct category for each image. The alternative evaluation technique is not a substitute for the original evaluation technique, at least in cases where there is certainty

about the labels in the test set. It relates, however, to two evaluation dimensions not captured by the standard evaluation techniques:

- It provides an internal quality measure useful during system development. While ultimately what counts for applications is whether the system produces the correct answer or not, during development it is important to rank better a system that comes closer to the correct answer, even if it still rates higher a wrong answer.
- It provides a measurement of the confidence of the system in the correct answers. Hence, it can be used to predict how well the system will perform for different thresholds of confidence (high values of the measure justify the use of a high confidence threshold, with the expectation of increased accuracy and a small loss in coverage).

To further illustrate how this alternative evaluation technique can be used to reveal aspects of the evaluation not obvious from the traditional scores, and also the generality of our parameter selection mechanism, we repeated our training of the indoor/outdoor classifier on our second set of images, those that had any kind of agreement from the human judges (not necessarily with strong beliefs; see Sect. 3). We randomly selected 1000 (approximately two-thirds) of the 1501 images in that set for training and the remaining 501 images for testing, and retrained the classifier using the optimal combination of parameters determined in Sect. 5. Three hundred and eight (30.8%) of the training images were indoor while 692 (69.2%) were outdoor; within the test set, 167 (33.3%) images were indoor, while 334 (66.7%) were outdoor. Using our first, standard evaluation model, the system achieved on the test set 77.05% accuracy using the raw TF*IDF similarities and 80.04% after converting those to probability estimates. The latter of these results is the most important, and it is 1.98% lower than the result from the main set with definite agreement. This makes sense, since manual categorizations with a lower degree of confidence are less likely to be accurate, and also may indicate images that are inherently harder to classify. This is in fact reflected in the system's confidence measure, which tends to be lower on these problematic cases; applying the alternative evaluation method to this second test set, we obtain overall accuracy of 76.56%, almost as high as that measure is for the first test set.

9 Using information about number of people

Earlier on, we noted that our goal in this line of research is to develop multiple classifiers for a number of broadly applicable classification features. It is natural to consider interactions between such classifiers, as information about one feature may well help the categorization according to another feature. In this section, we report on investigations regarding the effect knowledge about the number of

⁸ A further subdivision of our image data in two separate training sets and a test set would leave us with too few images in each set.

Table 9. Breakdown of the set of images with definite agreement on indoor/outdoor and number of people features into indoor and outdoor images for each value of the number of people feature

Number of people	Indoor images	Outdoor images	Percentage of indoor images
No people	2	75	2.6%
One person	122	108	53.0%
Two people	75	85	46.9%
Three or more people	155	332	31.8%
Crowd	8	119	6.3%
Total	362	719	33.5%

people in a photograph has on our ability to classify the image as indoor or outdoor.

We have not yet built a text-based classifier for this second feature,⁹ so we use instead ideal knowledge, provided by the human categorization of images according to this feature. We analyze the set of images that has both definite agreement between the human judges in the indoor/outdoor question and agreement in the number of people question (excluding ambiguous labels). This set contains 1081 images, 362 (33.5%) of which are indoor and 719 (66.5%) are outdoor, a similar distribution as in the larger set which we used for our main experiments. However, if we take the number of people as given, the distribution of indoor versus outdoor images within each category of the secondary feature changes, sometimes dramatically, as Table 9 shows.

To utilize this information, we need a formula that connects $f(I|c, d)$, the probability density of an image being indoor given that it belongs to category c according to the number of people feature and that it receives a similarity difference of d , to our old probability density estimates, $f(I|d)$. Unfortunately, a Bayesian expansion of $f(I|c, d)$ involves the joint density $f(c, d)$, which we cannot estimate without a classifier that predicts the number of people c from the difference d (or vice versa). Therefore, we intuitively derive an approximate formula for $f(I|c, d)$ as follows: given N images with similarity difference in a small neighborhood Δd around d , approximately $P(I|\Delta d) \cdot N$ of them will be indoor. Now, for any image that has a specific number of people c , its odds for being indoor will change (for better or worse) from the global proportion of indoor images $P(I)$ by the ratio $P(I|c)/P(I)$. If $P(c)$ is the global proportion of images with c people in them, the overall number of indoor images with c people among the initial N images with similarity difference close to d can be estimated as

$$N(I|c, \Delta d) \approx P(c) \cdot \frac{P(I|c)}{P(I)} \cdot P(I|\Delta d) \cdot N \quad (3)$$

Similarly, the overall number of outdoor images with c people among the same N images can be estimated as

$$N(O|c, \Delta d) \approx P(c) \cdot \frac{1 - P(I|c)}{1 - P(I)} \cdot (1 - P(I|\Delta d)) \cdot N \quad (4)$$

⁹ Although work is under way for building one based on face detection combined with name extraction from captions.

By combining (3) and (4), we get our formula for updating $P(I|\Delta d)$:

$$P(I|c, \Delta d) \approx \frac{N(I|c, \Delta d)}{N(I|c, \Delta d) + N(O|c, \Delta d)} \quad (5)$$

We applied this update formula to the images in the set with definite agreement on both the indoor/outdoor and number of people questions. Since that set is a subset of our main experimental image set, we took those images that were in the training set for the main set (see Sect. 5) as our training images, and the remaining as test images. The resulting training set had 732 images, of which 249 (34.0%) were indoor, and the testing set contained 349 images, of which 113 (32.4%) were indoor. If the methods of Sect. 5 are applied to this training/test set partition while ignoring the number of people information, we obtain 79.94% accuracy on the test set. If instead we assume perfect knowledge of the number of people variable and update the probability estimates by applying (5) (estimating quantities such as $P(I)$ and $P(I|c)$ from the training set), we obtain 80.23% accuracy on the test set. This is only a minor improvement, not statistically significant. However, if the alternative evaluation metric of the previous section is employed, accuracy improves from 74.96% to 77.19%. So, while few categorizations actually changed from wrong to right or vice versa, the system's confidence values in its decisions were more appropriate when the number of people was taken into account. In other words, on average, correct decisions were given higher confidence while the reverse happened to incorrect decisions.

10 Integration with an image-based approach developed at Columbia University

Recently, image processing researchers at Columbia, working in cooperation with the authors of the present paper, have implemented a novel classifier [22] using a vector-based approach relying on visual objects' frequency and dispersion across images, a method that is analogous to TF*IDF. Images are divided into 64 (8×8) regions of equal size, and a set of texture- and color-related features are computed for each region. The feature vectors associated with all the blocks from the training set are clustered, and the centroid of each cluster that is generated defines an object. For each image in the test

set, feature vectors defined on the image's regions are compared to the cluster centroids from the training set, and the closest such centroid (or object) is considered matched. For a single image, the OF, or object frequency, is the number of times an object occurs in that image; for a category, the OF is the number of times that object occurs over all images in that category. An object's IIF, or inverse image frequency, is the log of the ratio of the total number of images to the number of images that contain that object. Once $OF \cdot IIF$ scores have been computed for each image and category, similarity scores between images and categories can be computed with a dot product, similar to the text-based $TF \cdot IDF$ scheme (1). As with our $TF \cdot IDF$ approach, our colleagues used density estimation to estimate probabilities for each category.

Our colleagues started with the same collection of images as we did but divided it into a training set and testing set differently. They randomly selected approximately half of the 1339 images for training and used the rest for testing. Using this data set, their image-based classifier achieved an 82.4% accuracy rate, and our text-based classifier achieved an 83.3% accuracy rate. The image-based approach developed by Szummer and Picard only achieved a 74.7% accuracy rate on this set. Note that the 82.4% accuracy rate obtained by the image-based classifier was on a different training/testing division of the data set used in the experiments of Sect. 5 (where we reported an 82.02% accuracy rate); on common data, the text-based approach still performs (slightly) better.

Although our text-based approach still beat the image-based approach developed by our colleagues, the two systems did not necessarily do well on the same types of images, and it was natural to integrate the two approaches together. We did this by computing for each image/category pair a single similarity score which was a weighted sum of the similarity scores that the $TF \cdot IDF$ and $OF \cdot IIF$ systems computed. This final system achieved a classification accuracy of 86.2% on the test set of Paek et al., which was significantly better than how either approach did on its own.

11 Conclusions and future work

We have shown that our methods for categorization of images as indoor or outdoor strongly beat baseline performance and previous image-based techniques, and even begin to approach human performance, at least for the task and domain tested in our experiments. In fact, our system provides 93.72% of the correct answers that a human judge with access to the same kind of information does (82.02% versus 87.52% overall accuracy). By staying within the $TF \cdot IDF$ paradigm but experimenting with several parameters and adding the use of probability density estimates, we have created a system that achieves 82% accuracy on unseen images, without using any visual information in the images. The output of our system is in

terms of a probability, which is readily interpretable and provides a level of confidence in the system's decision. We have explored additional techniques both for image classification and for evaluating the constructed classifiers. In addition, we investigated the possibility of using additional information about images that might change the a priori probabilities of an image being indoor or outdoor, and there is some promise that the system's results may be improved. Our methods are general, use no prior knowledge about the news domain, and could be applied to other high-level visual features or non-visual categories, although currently our model of probability densities assumes dichotomous classifications.

We have examined a classification approach that relates to the Rocchio paradigm [25, 27] and combines $TF \cdot IDF$ estimates with a probabilistic normalization. A future alternative is to compare our results with pure probabilistic approaches such as naive Bayes [15] and connectionist models [16]. Certainly, we have not exhausted the space of possible features and transformations of the input data; we plan to examine additional such options, including morphological transformations/stemming, semantic information linking related words, normalization schemas that take into account document length, and different weighing of identified named entities.

We also summarized work done in conjunction with our image processing colleagues. They have developed an image-based classification system, using a method analogous to $TF \cdot IDF$, that performs almost as well as our system for indoor versus outdoor classification on their test set. Integrating the two approaches together led to improved performance over either method individually.

One future step is to expand the range of classification questions considered. We wish to explore high-level classifications such as indoor versus outdoor, number of people, and city versus landscape, and complement the general classifiers with specific image feature detectors (e.g., detectors of skies, handshakes, or faces). Our goal is to provide a hierarchy of such classifiers and analyze their interactions so that we can build a model that relates a combination of the high-level visual features to specific conditions under which an image is appropriate for inclusion in a multimedia document. It is also possible to use the text-based classifier as a bootstrapping component, classifying a subset of images for a task for which only few images have associated text, and passing on the results to the training stage of an image-based classifier.

We have also recently generalized our method of density estimation so that it is usable when there are more than two categories. We are currently applying this extended density estimation technique to the task of classifying each of our news documents (as opposed to just the images) into one of the semi-topical categories Struggle, Politics, Crime, Disaster, or Other. Initial results indicate that our system works quite well, achieving about 87% accuracy on our test set. We believe that these new categories may have interesting interactions with the indoor

versus outdoor categories examined in this paper. Highly confident decisions for one set of categories should aid predictions for the other set of categories. For example, our system achieves almost perfect precision and recall for the Disaster category, and almost all images which are part of these documents are outdoor images. Additionally, combining decisions concerning different sets of categories might tell us more about an image than a classification into a single set of categories. For example, indoor images which are part of political documents are almost always pictures of meetings or press conferences.

We are also in the process of generalizing our system to work for binary categorization in which a document can have multiple categories and a separate “yes”/“no” decision is made for every document/category pair. We have applied our system to the Reuters-21578 corpus [17], a collection of news articles frequently used in text categorization tasks [14, 45], and initial results indicate that our generalized density estimation significantly improves the accuracy of a TF*IDF based approach, although we are still not doing quite as well on that task as other advanced methods such as Support Vector Machines [14].

Eventually, we would like to train our system on many sets of categories representing both broad-based, non-topical visual features of images and also higher level, content-based properties of images or entire news documents. Our colleagues using image-based approaches will also continue building classifiers for additional sets of categories. It is likely that text-based classifiers will achieve better results for some categories and image-based classifiers will achieve better results for others. One interesting question might be to further investigate which properties of categories affect which type of system will have better results. Once we have various types of categorizers trained, we hope to further explore possible interactions between such systems. In the end, we want to automatically determine as much as possible about images with associated text using all the information at our disposal.

Acknowledgements. We would like to thank Seungyup Paek, Alejandro Jaimes, Shih-Fu Chang, Luis Gravano, Barry Schiffman, Hongyan Jing, and especially Kathleen McKeown for numerous suggestions and discussions during the development of this work. We are also grateful to the volunteers who categorized the images in our collection, and to the anonymous reviewers of EDL-99 and the IJODL special issue for additional suggestions that helped to improve this paper. The work reported here has been funded in part by a National Science Foundation STIMULATE grant, IRI-96-19124. Any opinions, findings, or recommendations are those of the authors, and do not necessarily reflect the views of the NSF.

Additional information about this work, including our collection of images, human response collection protocol, and detailed results, can be found at <http://www.cs.columbia.edu/~sable/inout.html>.

References

1. Aho, A.V., Chang, S.-F., McKeown, K.R., Radev, D., Smith, J.R., Zaman, K.: Columbia Digital News Project. *International Journal of Digital Libraries* 1(4):377–385, 1998
2. Bach, J.R., Fuller, C., Gupta, A., Hampapur, A., Horowitz, B., Humphrey, R., Jain, R.C., Shu, C.: The VIRAGE image search engine: An open framework for image management. In: *Proceedings of the Symposium on Electronic Imaging: Science and Technology—Storage and Retrieval for Image and Video Databases IV. IS&T/SPIE*, February 1996
3. Bates, D.M., Watts, D.G.: *Nonlinear Regression Analysis and its Applications*. Wiley, New York, 1988
4. Church, K.W.: A stochastic parts program and noun phrase parser for unrestricted text. In: *Proceedings of the Second Conference on Applied Natural Language Processing. ANLP. Austin, Texas. Association for Computational Linguistics*, February 1988, pp. 136–143
5. Fisher, R.A.: *Statistical Methods for Research Workers*. Oliver and Boyd, Edinburgh, United Kingdom, 5th edition, 1934
6. Fleiss, J.L.: *Statistical Methods for Rates and Proportions*. Wiley, New York, 2nd edition, 1981
7. Forsyth, D.A., Fleck, M.M.: Finding naked people. In: *Proceedings of the European Conference on Computer Vision*, Berlin, Germany, 1996
8. Fung, C.Y., Loe, K.F.: A new approach for image classification and retrieval. In *Proc. of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR-99)*, 1999
9. Gay, L.S., Croft, W.B.: Interpreting nominal compounds for information retrieval. *Information Processing and Management* 26(1):21–38, 1990
10. Harmandas, V., Sanderson, M., Dunlop, M.D.: Image retrieval by hypertext links. In: *Proc. 20th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR-97)*, 1997
11. Hastie, T., Pregibon, D.: Shrinking trees. Technical report, AT&T Bell Laboratories, 1990
12. Hatzivassiloglou, V., McKeown, K.: Towards the automatic identification of adjectival scales: Clustering adjectives according to meaning. In: *Proceedings of the 31st Annual Meeting of the ACL*, Columbus, Ohio. Association for Computational Linguistics, June 1993, pp. 172–182
13. Hicks, C.R.: *Fundamental Concepts in the Design of Experiments*. Holt, Rinehart, and Wilson, New York, 3rd edition, 1982
14. Joachims, T.: Text categorization with support vector machines: Learning with many relevant features. In: *Proceedings of the European Conference on Machine Learning*, 1998
15. Lewis, D.: Naive (bayes) at forty: The independence assumption in information retrieval. In: *Proceedings of the European Conference on Machine Learning*, 1998
16. Lewis, D., Schapire, R., Callan, J., Papka, R.: Training algorithms for linear text classifiers. In: *Proc. of the 19th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR-96)*, 1996
17. Lewis, D.D.: Reuters-21578 text categorization test collection, readme file (version 1.2), 1997
18. Lipson, P., Grimson, E., Sinha, P.: Configuration based scene classification and image indexing. In: *Proceedings of the 1997 Conference on Computer Vision and Pattern Recognition (CVPR-97)*, 1997
19. Niblack, W., Barber, R., Equitz, W., Flickner, M., Glasman, E., Petkovic, D., Yanker, P., Faloutsos, C., Taubin, G.: The QBIC project: Querying images by content using color, texture, and shape. In: *Proceedings of Symposium on Electronic Imaging: Science and Technology—Storage and Retrieval for Image and Video Databases. SPIE*, February 1993
20. Ogle, V.E., Stonebraker, M.: Chabot: Retrieval from a relational database of images. *IEEE Computer Magazine*, 28(9):40–48, September 1995
21. Oliva, A., Torralba, A.B., Guerin-Dugue, A., Herault, J.: Global semantic classification of scenes using power spectrum templates. In: *Proceedings of the Challenge of Image Retrieval (CIR-99)*, 1999
22. Paek, S., Sable, C.L., Hatzivassiloglou, V., Jaimes, A., Schiffman, B.H., Chang, S.-F., McKeown, K.R.: Integration of visual and text-based approaches for the content labeling and classification of photographs. In: *ACM SIGIR Workshop on Multimedia Indexing and Retrieval (SIGIR-99)*, 1999

23. Pentland, A., Picard, R.W., Sclaroff, S.: Photobook: Tools for content-based manipulation of image databases. In: Proceedings of the Symposium on Electronic Imagic: Science and Technology—Storage and Retrieval for Image and Video Databases II, Bellingham, Washington, 1994, SPIE, pp. 34–47
24. Quinlan, J.R.: Induction of decision trees. *Machine Learning*, 1(1):81–106, 1986
25. Rocchio, J.: Relevance feedback in information retrieval. In *The SMART Retrieval System: Experiments in Automatic Document Processing*, Chap. 14. Prentice-Hall Inc., 1971, pp. 974–979
26. Rowe, N.C., Guglielmo, E.J.: Exploiting captions in the retrieval of multimedia data. *Information Processing and Management* 29(4):453–561, 1993
27. Salton, G.: Developments in automatic text retrieval. *Science* 253:974–979, 1991
28. Salton, G., Buckley, C.: Term weighting approaches in automatic text retrieval. *Information Processing and Management* 24(5):513–523, 1988
29. Salton, G., Smith, M.: On the application of syntactic methodologies in automatic text analysis. In: Proceedings of the Twelfth Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR-89), 1989
30. Salton, G.: *Automatic Text Processing: The Transformation, Analysis, and Retrieval of Information by Computer*. Addison-Wesley, Reading, Massachusetts, 1989
31. Santner, T.J., Duffy, D.E.: *The Statistical Analysis of Discrete Data*. Springer-Verlag, New York, 1989
32. Satoh, S., Sato, T., Smith, M., Nakamura, Y., Kanade, T.: Name-it: Naming and detecting faces in news video. In: Proceedings of the 1997 Conference on Computer Vision and Pattern Recognition (CVPR-97), 1997
33. Schneiderman, H., Kanade, T.: Probabilistic modeling of local appearance and spatial relationships for object recognition. In: Proceedings of the 1998 Conference on Computer Vision and Pattern Recognition (CVPR-98), 1998
34. Scott, D.W.: *Multivariate Density Estimation: Theory, Practice, and Visualization*. John Wiley and Sons, New York, 1992
35. Smeaton, A.F.: Progress in the application of natural language processing to information retrieval tasks. *The Computer Journal* 35(3):268–278, 1992
36. Smeaton, A.F., Quigley, I.: Experiments on using semantic distances between words in image caption retrieval. In: Proceedings of the 19th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR-96), 1996
37. Smith, J.R., Chang, S.-F.: VisualSEEK: A fully automated content-based image query system. In: Proceedings of the ACM Multimedia Conference, Boston, Massachusetts, November 1996
38. Smith, J.R., Chang, S.-F.: Visually searching the Web for content. *IEEE Multimedia*, 4(3):12–20, July–September 1997
39. Srihari, R.K.: Automatic indexing and content-based retrieval of captioned images. *IEEE Computer Magazine*, 28(9):49–58, September 1995
40. Szummer, M., Picard, R.W.: Indoor-outdoor image classification. In: *IEEE Workshop on Content Based Access of Image and Video Databases (CAIVD-98)*, Bombay, India, January 1998, pp. 42–51
41. Torralba, A.B., Oliva, A.: Semantic organization of scenes using discriminant spectral templates. In: Proceedings of the International Conference in Computer Vision (ICCV-99), 1999
42. Vailaya, A., Figueiredo, M., Jain, A.K., Zhang, H.: Bayesian framework for semantic classification of outdoor vacation images. In: Proceedings of SPIE (The International Society for Optical Engineering)—Storage and Retrieval for Image and Video Databases VII, San Jose, California, January 1999
43. van Rijsbergen, C.J. *Information Retrieval*. Butterworths, London, 2nd edition, 1979
44. Wacholder, N., Ravin, Y., Choi, M.: Disambiguation of proper names in text. In: Proceedings of the 5th ACL Conference on Applied Natural Language Processing. ANLP. Washington, D.C., April 1997, pp. 202–208
45. Yang, Y., Liu, X.: A re-examination of text categorization methods. In: Proceedings of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR-99), 1999