



Data Mining of Traffic Video Sequences

Final Report

Prepared by:

Ajay J. Joshi
Nikolaos P. Papanikolopoulos

Artificial Intelligence, Robotics and Vision Laboratory
Department of Computer Science and Engineering
University of Minnesota

CTS 09-25

Technical Report Documentation Page

1. Report No. CTS 09-25	2.	3. Recipients Accession No.	
4. Title and Subtitle Data Mining of Traffic Video Sequences		5. Report Date September 2009	
		6.	
7. Author(s) Ajay J. Joshi and Nikolaos Papanikolopoulos		8. Performing Organization Report No.	
9. Performing Organization Name and Address Artificial Intelligence, Robotics and Vision Laboratory Department of Computer Science and Engineering University of Minnesota 200 Union Street SE Minneapolis, Minnesota 55455		10. Project/Task/Work Unit No. CTS Project # 2008003	
		11. Contract (C) or Grant (G) No. (c) (wo)	
12. Sponsoring Organization Name and Address Intelligent Transportation Systems Institute University of Minnesota 200 Transportation and Safety Building 511 Washington Ave. SE Minneapolis, Minnesota 55455		13. Type of Report and Period Covered Final Report	
		14. Sponsoring Agency Code	
15. Supplementary Notes http://www.its.umn.edu/Publications/ResearchReports/			
16. Abstract (Limit: 250 words) Automatically analyzing video data is extremely important for applications such as monitoring and data collection in transportation scenarios. Machine learning techniques are often employed in order to achieve these goals of mining traffic video to find interesting events. Typically, learning-based methods require significant amount of training data provided via human annotation. For instance, in order to provide training, a user can give the system images of a certain vehicle along with its respective annotation. The system then learns how to identify vehicles in the future - however, such systems usually need large amounts of training data and thereby cumbersome human effort. In this research, we propose a method for active learning in which the system interactively queries the human for annotation on the most informative instances. In this way, learning can be accomplished with lesser user effort without compromising performance. Our system is also efficient computationally, thus being feasible in real data mining tasks for traffic video sequences.			
17. Document Analysis/Descriptors Surveillance, Machine learning, Learning, Active learning, Vehicle classification, Computer vision, Data mining		18. Availability Statement No restrictions. Document available from: National Technical Information Services, Springfield, Virginia 22161	
19. Security Class (this report) Unclassified	20. Security Class (this page) Unclassified	21. No. of Pages 29	22. Price

Data Mining of Traffic Video Sequences

Final Report

Prepared by
Ajay J. Joshi
Nikolaos P. Papanikolopoulos

Artificial Intelligence, Robotics and Vision Laboratory
Department of Computer Science and Engineering
University of Minnesota

September 2009

Published by
Intelligent Transportation Systems Institute
Center for Transportation Studies
University of Minnesota
200 Transportation and Safety Building
511 Washington Avenue SE
Minneapolis, MN 55455

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated under the sponsorship of the Department of Transportation University Transportation Centers Program, in the interest of information exchange. The U.S. Government assumes no liability for the contents or use thereof. This report does not necessarily reflect the official views or policy of the Intelligent Transportation Systems Institute or the University of Minnesota.

The authors, the Intelligent Transportation Systems Institute, the University of Minnesota and the U.S. Government do not endorse products or manufacturers. Trade or manufacturers' names appear herein solely because they are considered essential to this report.

ACKNOWLEDGEMENTS

The authors wish to acknowledge those who made this research possible. The study was funded by the Intelligent Transportation Systems (ITS) Institute, a program of the University of Minnesota's Center for Transportation Studies (CTS). Financial support was provided by the United States Department of Transportation's Research and Innovative Technologies Administration (RITA).

TABLE OF CONTENTS

CHAPTER 1: INTRODUCTION.....	1
POOL-BASED LEARNING SETUP	1
CHAPTER 2: MULTI-CLASS ACTIVE LEARNING	3
ENTROPY MEASURE (EP).....	3
BEST-VERSUS-SECOND BEST (BVSB)	4
COMPUTATIONAL COST	5
CHAPTER 3: PREVIOUS WORK.....	7
CHAPTER 4: EXPERIMENTAL RESULTS.....	9
STANDARD DATASETS	9
REDUCTION IN TRAINING REQUIRED	11
OBJECT RECOGNITION	11
TIME PERFORMANCE OF ACTIVE LEARNING	13
EXPLORING THE SPACE	13
CLASSIFYING VEHICLES IN IMAGES.....	14
CHAPTER 5: CONCLUSIONS	19
REFERENCES	21

LIST OF FIGURES

Figure 1: An illustration of why entropy can be a poor estimate of classification uncertainty.	4
Figure 2: Classification accuracy on (a) Pendigits, (b) Letter, and (c) USPS datasets from the UCI repository.	9
Figure 3: Active learning on the Caltech-101 dataset.....	12
Figure 4: Example selection time as a function of active pool size.	13
Figure 5: Space exploration of active selection - BvSB-based selection is almost as good as random exploration, while the former achieves much higher classification accuracy than random.....	14
Figure 6: Frame of highway video 1 for classification of vehicles.	15
Figure 7: Frame of highway video 2 for classification of vehicles.	15
Figure 8: Active learning for classification of vehicles in surveillance video.....	16
Figure 9: Changing illumination conditions over time due to changing cloud cover.	16

LIST OF TABLES

Table 1: Dataset properties and their corresponding sizes used for demonstrating active learning classification accuracy.	10
Table 2: Reduction in training set sizes for active learning.	11

EXECUTIVE SUMMARY

Video sequences and images from surveillance cameras have a large amount of information that needs to be automatically analyzed. Classification of vehicles in video sequences is one of the primary challenges underlying many transportation applications. The principal bottleneck in applying learning techniques to classification problems is the large amount of labeled training data required. Especially for images and video, providing training data is very expensive in terms of human time and effort. In this report we propose an active learning approach to tackle the problem. Instead of passively accepting random training examples, the active learning algorithm iteratively selects unlabeled examples for the user to label, so that human effort is focused on labeling the most useful examples. Our method relies on the idea of uncertainty sampling, in which the algorithm selects unlabeled examples that it finds hardest to classify. Specifically, we propose an uncertainty measure that generalizes margin-based uncertainty to the multi-class case and is easy to compute, so that active learning can handle a large number of classes and large data sizes efficiently. We demonstrate feasibility results on standard well-known datasets. We also show results on real traffic video for classifying vehicles in surveillance camera images. The proposed method gives large reductions in the number of training examples required over random selection to achieve similar classification accuracy, with little computational overhead. Hence the proposed methods offer a practical new way for analyzing large quantities of data with little human effort.

CHAPTER 1: INTRODUCTION

Most methods for image classification use statistical models that are learned from labeled training data. In the typical setting, a learning algorithm passively accepts randomly provided training examples. However, providing labeled examples is costly in terms of human time and effort. Further, small training sizes can lead to poor future classification performance. In this paper, we propose an active learning approach for minimizing the number of training examples required, and achieving good classification at the same time. In active learning, the learning algorithm selects “useful” examples for the user to label, instead of passively accepting data. Theoretical results show that active selection can significantly reduce the number of examples required compared to random selection for achieving similar classification accuracy (cf. [4] and references therein). Even though most of these results require strict assumptions and are applicable to binary classification, they serve as a motivation to develop active learning algorithms for multi-class problems.

The principal idea in active learning is that not all examples are of equal value to a classifier, especially for classifiers that have sparse solutions. For example, consider a Support Vector Machine trained on some training examples. The classification surface remains the same if all data except the support vectors are omitted from the training set. Thus, only a few examples define the separating surface and all the other examples are redundant to the classifier. We wish to exploit this aspect in order to actively select examples that are useful for classification. The primary contribution of this paper is an active learning method that i) can easily handle multi-class problems, ii) works without knowledge of the number of classes (so that this number may increase with time), and iii) is computationally and interactively efficient, allowing application to large datasets with little human time consumed. For better clarity, comparisons of our method to previous work are made in a later section after describing our approach in detail.

Pool-based learning setup

Here we describe pool-based learning, which is a very common setup for active learning. We consider that a classifier is trained using a small number of randomly selected labeled examples called the seed set. The active learning algorithm can then select examples to query the user (for labels) from a pool of unlabeled examples referred to as the active pool. The actively selected examples along with user-provided labels are then added to the training set. This querying process is iterative such that after each iteration of user feedback, the classifier is retrained. Finally, performance evaluation is done on a separate test set different from the seed set and the active learning pool. In this work, we use Support Vector Machines (SVM) as the primary classifier for evaluation; however, other classification techniques could potentially be employed.

CHAPTER 2: MULTI-CLASS ACTIVE LEARNING

Our approach follows the idea of uncertainty sampling [2, 6], wherein examples on which the current classifier is uncertain are selected to query the user. Distance from the hyperplane for margin-based classifiers has been used as a notion of uncertainty in previous work. However, this does not easily extend to multi-class classification due to the presence of multiple hyperplanes. We use a different notion of uncertainty that is easily applicable to a large number of classes. The uncertainty can be obtained from the class membership probability estimates for the unlabeled examples as output by the multi-class classifier. In the case of a probabilistic model, these values are directly available. For other classifiers such as SVM, we need to first estimate class membership probabilities of the unlabeled examples. In the following, we outline our approach for estimating the probability values for multi-class SVM. However, such an approach for estimating probabilities can be used with many other non-probabilistic classification techniques also.

Our uncertainty sampling method relies on probability estimates of class membership for all the examples in the active pool. In order to obtain these estimates, we follow the approach proposed by [12], which is a modified version of Platt's method to extract probabilistic outputs from SVM [15].

Entropy measure (EP)

Each labeled training example belongs to a certain class. However, we do not know true class labels for examples in the active pool. For each unlabeled example, we can consider the class membership variable to be a random variable denoted by Y . We have a distribution p for Y of estimated class membership probabilities computed in the way described above. Entropy is a measure of uncertainty of a random variable. Since we are looking for measures that indicate uncertainty in class membership Y , its discrete entropy is a natural choice.

Higher values of entropy imply more uncertainty in the distribution; this can be used as an indicator of uncertainty of an example. If an example has a distribution with high entropy, the classifier is uncertain about its class membership.

The algorithm proceeds in the following way. At each round of active learning, we compute class membership probabilities for all examples in the active pool. Examples with the highest estimated value of discrete entropy are selected to query the user. User labels are obtained and the corresponding examples are incorporated in the training set and the classifier is retrained.

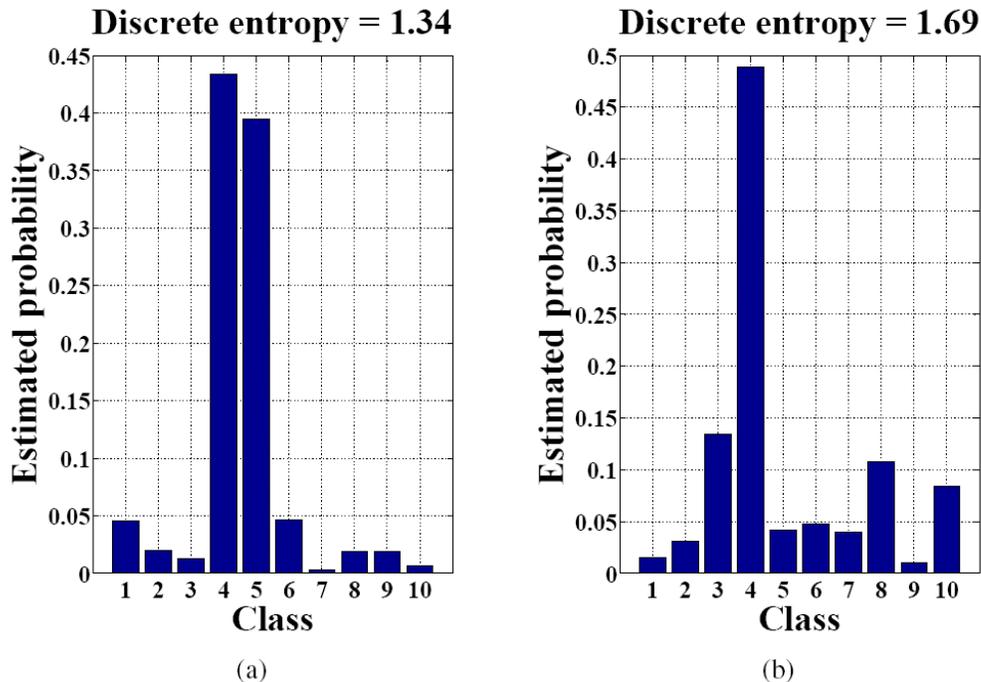


Figure 1: An illustration of why entropy can be a poor estimate of classification uncertainty. The plots show estimated probability distributions for two unlabeled examples in a 10 class problem. In (a), the classifier is highly confused between classes 4 and 5. In (b), the classifier is relatively more confident that the example belongs to class 4, but is assigned higher entropy. The entropy measure is influenced by probability values of unimportant classes.

Best-versus-second best (BvSB)

Even though EP-based active learning is often better than random selection, it has a drawback. A problem of the EP measure is that its value is heavily influenced by probability values of unimportant classes. See Figure 1 for a simple illustration. The figure shows estimated probability values for two examples on a 10-class problem. The example on the left has a smaller entropy than the one on the right. However, from a classification perspective, the classifier is more confused about the former since it assigns close probability values to two classes. For the example in Figure 1(b), small probability values of unimportant classes contribute to the high entropy score, even though the classifier is much more confident about the classification of the example. This problem becomes even more acute when a large number of classes are present. Although entropy is a true indicator of uncertainty of a random variable, we are interested in a more specific type of uncertainty relating only to classification amongst the most confused classes (the example is virtually guaranteed to not belong to classes having a small probability estimate).

Instead of relying on the entropy score, we take a more greedy approach to account for the problem mentioned. We consider the difference between the probability values of the two classes having the highest estimated probability value as a measure of uncertainty. Since it is a comparison of the best guess and the second best guess, we refer to it as the Best-versus-Second-Best (BvSB) approach. Such a measure is a more direct way of estimating confusion about class

membership from a classification standpoint. Using the BvSB measure, the example on the left in Figure 1 will be selected to query the user.

As mentioned previously, confidence estimates are reliable in the sense that classes assigned low probabilities are very rarely the true classes of the examples. However, this is only true if the initial training set size is large enough for good probability estimation. In our experiments, we start from as few as 2 examples for training in a 100 class problem. In such cases, initially the probability estimates are not very reliable, and random example selection gives similar results. As the number of examples in the training set grows, active learning through BvSB quickly dominates random selection by a significant margin.

Note that most methods that do not use margins rely on some kind of model on the data. Therefore, probability estimation with very few examples presents a problem to most active learning methods. Overcoming this issue is an important area for future work.

Computational cost

There are two aspects to the cost of active selection. One is the cost of training the SVM on the training set at each iteration. Second is probability estimation on the active pool, and selecting examples with the highest BvSB score. SVM training is by far the most computationally intensive component of the entire process. However, the essence of active learning is to minimize training set sizes through intelligent example selection. Therefore, it is more important to consider the cost of probability estimation and example selection on the relatively much larger active pool. The first cost comes from probability estimation in binary SVM classifiers. The estimation is efficient since it is performed using Newton's method with backtracking line search that guarantees quadratic rate of convergence. Given class probability values for binary SVMs, multi-class probability estimates can be obtained in $O(k)$ time per example [19], where k is the number of classes. Due to the linear relationship, the algorithm is scalable to problems having a large number of classes, unlike most previous methods. In the experiments, we also demonstrate empirical observations indicating linear time relationship with the active pool size. We were easily able to perform experiments with seed set sizes varying from 2 to 500 examples, active pool sizes of up to 10000 examples, and a up to 102-class classification problems.

A typical run with seed set of 50 examples, active pool of 5000 examples, and a 10-class problem took about 22 seconds for 20 active learning rounds with 5 examples added at each round. The machine used had a 1.87 Ghz single core processor with 2 Gb of memory. All the active selection code was written in Matlab, and SVM implementation was done using LIBSVM (written in C) interfaced with Matlab. The total time includes the time taken to train the SVM, to produce binary probability values, and to estimate multi-class probability distribution for each example in the active pool at each round.

CHAPTER 3: PREVIOUS WORK

Tong and Chang [17] propose active learning for SVM in a relevance feedback framework for image retrieval. Their approach relies on the margins for unlabeled examples for binary classification. Tong et al. [18] use an active learning method to minimize the version space¹ at each iteration. However, both these approaches target binary classification.

Gaussian processes (GP) have been used for object categorization by Kapoor and Grauman [10]. They demonstrate an active learning approach through uncertainty estimation based on GP regression, which requires $O(N^3)$ computations, cubic in the number of training examples. They use one-versus-all SVM formulation for multi-class classification, and select one example per classifier at each iteration of active learning. In our work, we use the one-versus-one SVM formulation, and allow the addition of a variable number of examples at each iteration.

Holub et al. [8] recently proposed a multi-class active learning method. Their method selects examples from the active pool, whose addition to the training set minimizes the expected entropy of the system. In essence, it is an information-based approach. Note that our method computes the uncertainty through probability estimates of class membership, which is an uncertainty sampling approach. The entropy-based approach proposed in [8] requires $O(k^3N^3)$ computations, where N is the number of examples in the active pool and k is the number of classes.

Qi et al. [16] demonstrate a multi-label active learning approach. Their method employs active selection along two dimensions – examples and their labels. Label correlations are exploited for selecting the examples and labels to query the user.

For handling multiple image selection at each iteration, Hoi et al. [7] introduced batch mode active learning with SVMs. Since their method is targeted towards image retrieval, the primary classification task is binary; to determine whether an image belongs to the class of the query image. Active learning with uncertainty sampling has been demonstrated by Li and Sethi [11], in which they use conditional error as a metric of uncertainty, and work with binary classification.

In summary, compared to previous work, our active learning method handles the multi-class case efficiently, allowing application to huge datasets with a large number of categories.

CHAPTER 4: EXPERIMENTAL RESULTS

This section reports experimental results of our active selection algorithm compared to random example selection. We demonstrate results on standard image datasets available from the UCI repository [1], the Caltech-101 dataset of object categories, and a dataset of 13 natural scene categories. All the results show significant improvement owing to active example selection.

Standard datasets

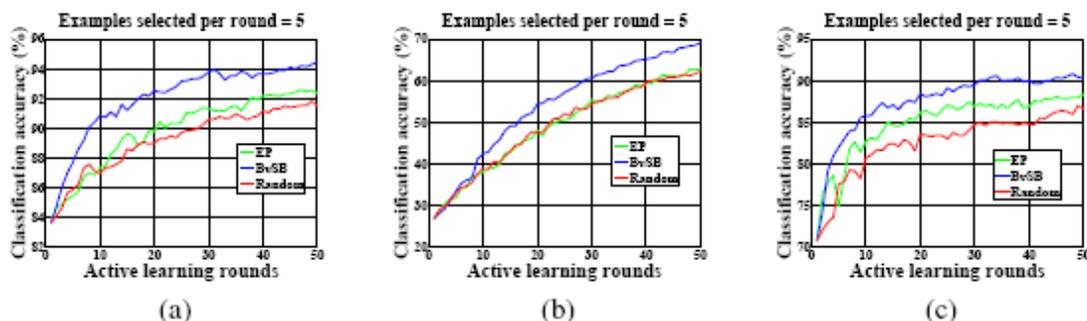


Figure 2: Classification accuracy on (a) Pendigits, (b) Letter, and (c) USPS datasets. Note the improvement in accuracy obtained by BvSB approach over random selection. For similar accuracy, the active learning method requires far fewer training examples. In (b), EP-based selection performs poorly due to the larger number of classes.

We choose three datasets that are relevant to image classification tasks. The chosen datasets and their properties are summarized in Table 1 along with seed set, active pool, and test set sizes used in our experiments. We also report the kernel chosen for the SVM classifier. For choosing the kernel, we ran experiments using the linear, polynomial, and Radial Basis Function (RBF) kernels on a randomly chosen training set, and picked the kernel that gave the best classification accuracy averaging over multiple runs.

Figure 2(a) shows classification results on the Pendigits dataset. The three methods compared are EP-based selection, BvSB-based selection, and random example selection. All three methods start with the same seed set of 100 examples. At each round of active learning, we select $n = 5$ examples to query the user for labels. BvSB selects useful examples for learning, and gradually dominates both the other approaches. Given the same size of training data, as indicated by the same point on the x-axis, BvSB gives significantly improved classification accuracy. From another perspective, for achieving the same value of classification accuracy on the test data (same point on the y-axis), our active learning method needs far fewer training examples than random selection. The result indicates that the method selects useful examples at each iteration, so that user input can be effectively utilized on the most relevant examples. Note that EP-based

Dataset	# classes	Feature dimension	Seed set size	Active pool size	Test set size	Kernel
Pendigits	10	16	100	5000	2000	Linear
USPS	10	256	100	7000	2000	RBF ($\sigma = 256$)
Letter	26	16	100	7000	5000	RBF ($\sigma = 16$)

Table 1: Dataset properties and the corresponding sizes used. Pendigits: Pen-based Recognition of handwritten digits. USPS: Optical recognition of handwritten digits originally from the US Postal Service. Letter: Letter recognition dataset. All obtained from UCI [1].

selection does marginally better than random. The difference can be attributed to the fact that entropy is a somewhat indicative measure of classification uncertainty. However, as pointed out before, the entropy value has problems of high dependence on unlikely classes. The BvSB measure performs better by greedily focusing on the confusion in class membership between the most likely classes instead. This difference between the two active selection methods becomes clearer when we look at the results on a 26 class problem. Figure 2(b) shows classification accuracy plots on the Letter dataset, which has 26 classes. EP-based selection performs even worse on this problem due to the larger number of classes, i.e., the entropy value is skewed due to the presence of more unlikely classes. Entropy is a bad indicator of classification uncertainty in this case, and it gives close to random performance. Even with a larger number of classes, the figure shows that BvSB-based selection outperforms random selection. After 50 rounds of active learning, the improvement in classification accuracy is about 7%, which is significant for data having 26 classes. In Figure 2(c), we show results on the USPS dataset, a dataset consisting of handwritten digits from the US Postal Service. The performance of all methods is similar to that obtained on the Pendigits dataset shown in Figure 2(a). Active selection needs far fewer training examples compared to random selection to achieve similar accuracy.

Reduction in training required

BvSB selection rounds	Random selection rounds	% Reduction in # training examples
3	6	11.53
4	10	20
5	13	24.24
6	19	33.33
7	28	43.75
8	29	42.85
9	43	53.96
10	44	53.12
11	43	50.79
12	48	52.94
13	50	52.85

Table 2: Percentage reduction in the number of training examples provided to the active learning algorithm to achieve classification accuracy equal to or more than random example selection on the USPS dataset.

In this section, we perform experiments to quantify the reduction in the number of training examples required for BvSB to obtain similar classification accuracy as random example selection. Consider a plot like Figure 2(c) above. For each round of active learning, we find the number of rounds of random selection to achieve the same classification accuracy. In other words, fixing a value on the y-axis, we measure the difference in the training set size of both methods and report the corresponding training rounds in Table 2. The table shows that active learning achieves a reduction of about 50% in the number of training examples required, i.e., it can reach near optimal performance with 50% fewer training examples. Table 2 reports results for the USPS dataset, however, similar results were obtained for the Pendigits dataset and the Letter dataset. The results show that even for problems having up to 26 classes, active learning achieves significant reduction in the amount of training required.

An important point to note from Table 2 is that active learning does not provide a large benefit in the initial rounds. One reason for this is that all methods start with the same seed set initially. In the first few rounds, the number of examples actively selected is far fewer compared to the seed set size (100 examples). Actively selected examples thus form a small fraction of the total training examples, explaining the small difference in classification accuracy of both methods in the initial rounds. As the number of rounds increase, the importance of active selection becomes clear, explained by the reduction in the amount of training required to reach near-optimal performance.

Object recognition

In Figure 3, we demonstrate results on the Caltech-101 dataset of object categories [5]. As image features, we use the precomputed kernel matrices obtained from the Visual Geometry group at

Oxford (<http://www.robots.ox.ac.uk/vgg/research/caltech/index.html>). These features give state-of-the-art performance on the Caltech dataset. The data is divided into 15 training and 15 test images per class, forming a total of 1530 images in the training and test sets each (102 classes including the ‘background’ class). We start with a seed set of only 2 images randomly selected out of the 1530 training images. We start with an extremely small seed set to simulate real-world scenarios. The remaining 1528 images in the training set form the active pool. After each round of active learning, classification accuracy values are computed on the separate test set of 1530 images. Note that in our experiments, the training set at each round of active learning is not necessarily balanced across classes, since the images are chosen by the algorithm itself. Such an experiment is closer to a realistic setting in which balanced training sets are usually not available (indeed, since providing balanced training sets needs human annotation, defeating our purpose). From Figure 3, we can see that active learning through BvSB-based selection outperforms random example selection in this 102 class problem. Interestingly, the difference in classification accuracy between active selection and random selection starts decreasing after about 70 rounds of learning. This can be attributed to the relatively limited size of the active pool; after 70 learning rounds, about half the active pool has been exhausted. Intuitively, the larger the active pool size, the higher the benefit of using active learning, since it is more unlikely for random selection to query useful images. In real-world image classification problems, the size of the active pool is usually extremely large, often including thousands of images available on the web. Therefore, the dependence on active pool sizes is not a limitation in most cases.

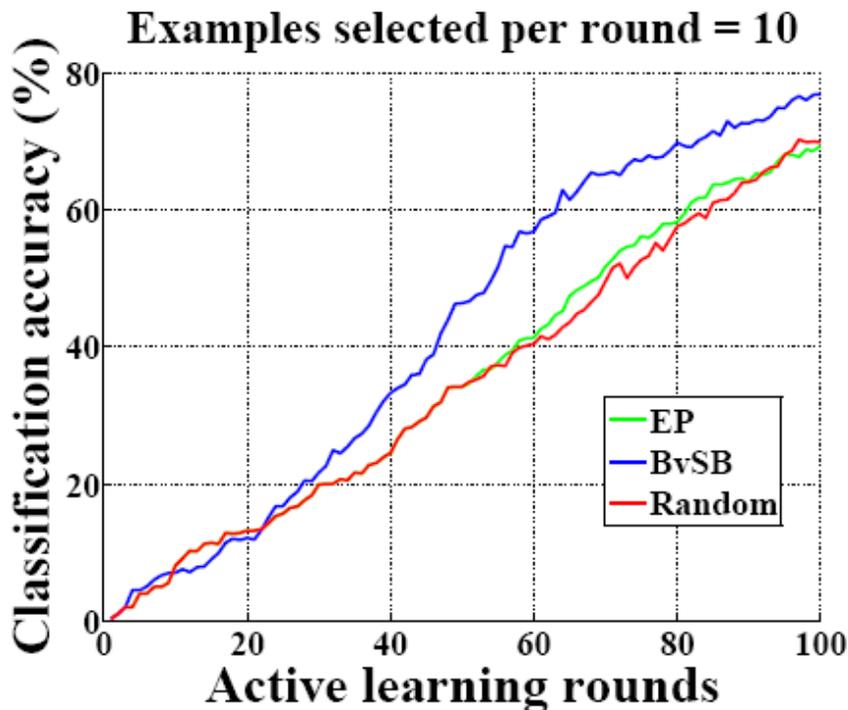


Figure 3: Active learning on the Caltech-101 dataset.

Time performance of active learning

From another perspective, the necessity of large active pool sizes points to the importance of computational efficiency in real-world learning scenarios. In order for the methods to be practical, the learning algorithm must be able to select useful images from a huge pool in reasonable time. Empirical data reported in Figure 4 suggests that our method requires time varying linearly with active pool size. The method is therefore scalable to huge active pool sizes common in real applications.

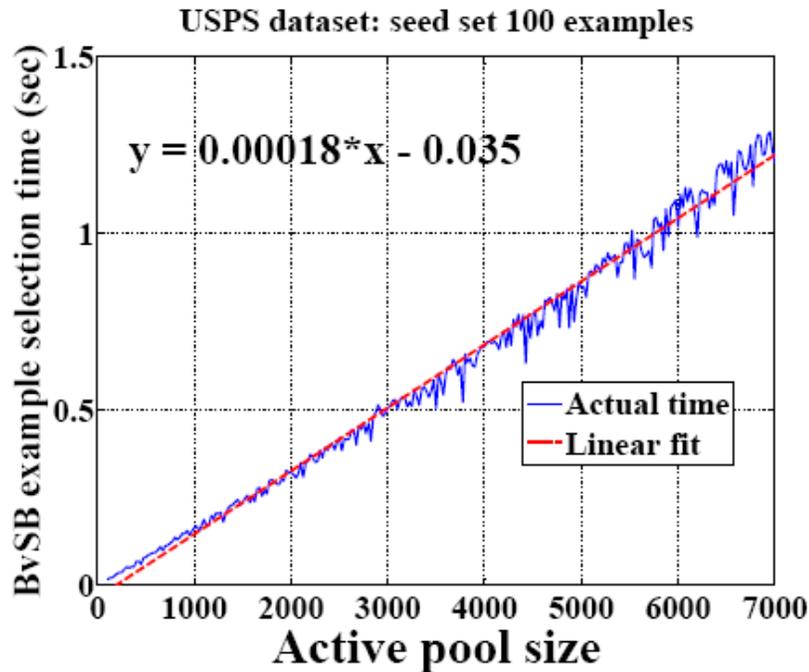


Figure 4: Example selection time as a function of active pool size. The relationship is linear over a large range with the equation shown in the figure. This demonstrates that the method is scalable to large active pool sizes that are common in real applications.

Exploring the space

In many applications, the number of categories to be classified is extremely large, and we start with only a few labeled images. In such scenarios, active learning has to balance two often conflicting objectives – exploration and exploitation. Exploration in this context means the ability to obtain labeled images from classes not seen before. Exploitation refers to classification accuracy on the classes seen so far. Exploitation can conflict with exploration, since in order to achieve high classification accuracy on the seen classes, more training images from those classes might be required, while sacrificing labeled images from new classes. In the results so far, we show classification accuracy on the entire test data consisting of all classes – thus good performance requires a good balance between exploration and exploitation. Here we explicitly

demonstrate how the different example selection mechanisms explore the space for the Caltech-101 dataset that has 102 categories. Figure 5 shows that the BvSB measure finds newer classes almost as fast as random selection, while achieving significantly higher classification accuracy than random selection. Fast exploration of BvSB implies that learning can be started with labeled images from very few classes and the selection mechanism will soon obtain images from the unseen classes. Interestingly, EP-based selection explores the space poorly.

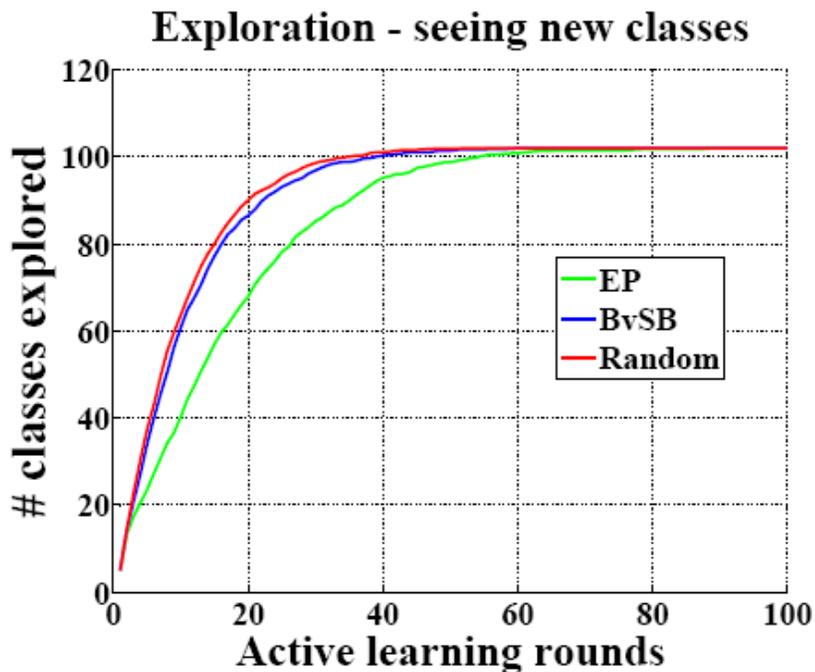


Figure 5: Space exploration of active selection – BvSB-based selection is almost as good as random exploration, while the former achieves much higher classification accuracy than random.

Classifying vehicles in images

The following images show frames from some video sequences on which classification of vehicles into multiple categories was done. The videos are taken in real traffic from a camera mounted near a highway.



Figure 6: Frame of highway video 1 for classification of vehicles.



Figure 7: Frame of highway video 2 for classification of vehicles.

All experiments were performed on video that represents realistic conditions for surveillance.

Figure 8 shows a comparison between active learning using BvSB- and EP-based example selection and random selection. We can see that a significant advantage can be achieved using active learning. The amount of training data is reduced by about 50% and hence more training in varied conditions can be easily obtained.

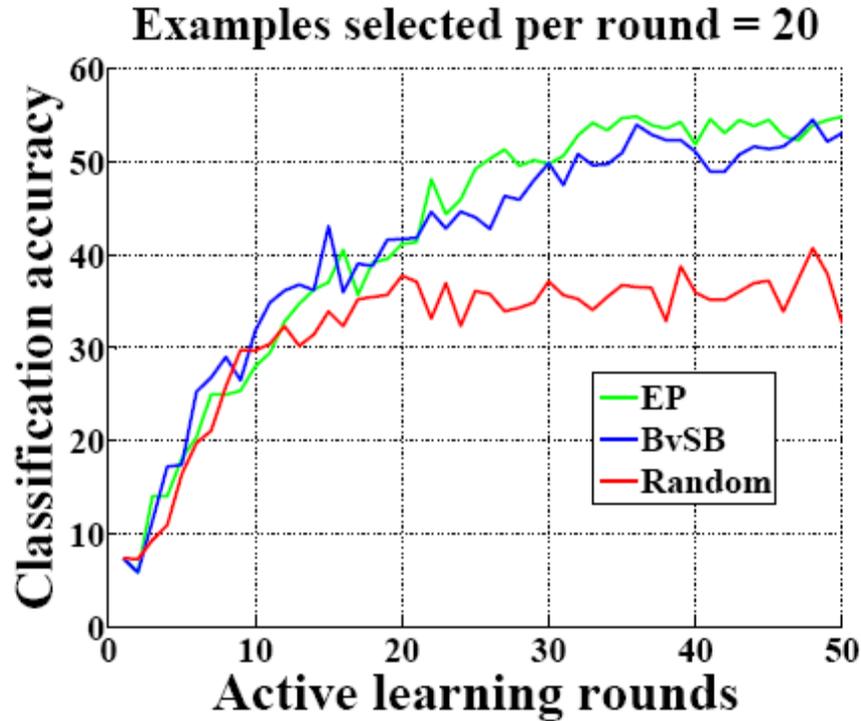


Figure 8: Active learning for classification of vehicles in surveillance video.

Figure 9 shows a sequence of images obtained from a camera monitoring a long stretch of a highway that shows significant illumination changes due to shadows and cloud cover. This presents one of the biggest challenges for current vision systems. Our active learning method can effectively tackle such scenarios by utilizing human input in an optimal way to perform good classification while limiting the amount of input necessary at the same time.



Figure 9: Changing illumination conditions over time due to changing cloud cover.

Such active learning schemes can be extremely important in real scenarios such as surveillance in transportation corridors. The challenges arise primarily from the fact that providing training for detecting/classifying certain tasks is difficult, in part due to changes in scene conditions that take place over time. Also, a system trained in one location is likely to fail in another due to the scene-specific training.

We have shown that active learning can substantially reduce the amount of training required, and thus make learning systems more practical in challenging scenarios like data mining for transportation applications.

CHAPTER 5: CONCLUSIONS

In this paper, we have proposed a simple active learning method for multi-class image classification. The proposed method achieves significant reduction in training required, along with efficient scaling to a large number of categories and huge data sizes.

In data mining for traffic video, there are many challenges as regards to the large amount of data to be analyzed, classification accuracy to be achieved, time limitations, and human effort for training the systems. In this work, based on our research on active learning, we make significant progress in multiple directions. We propose practical active learning schemes that can work with large data sizes in realistic conditions with little human input. Further, they also give state-of-the-art classification performance.

REFERENCES

- [1] A. Asuncion and D.J. Newman. *UCI machine learning repository*. University of California, Irvine, School of Information and Computer Sciences, 2007.
- [2] C. Campbell, N. Cristianini, and A.J. Smola. "Query learning with large margin classifiers." In *ICML '00: Proceedings of the International Conference on Machine Learning, 2000*, San Francisco, CA.
- [3] C.C. Chang and C.J. Lin. LIBSVM: A library for support vector machines, 2001. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- [4] S. Dasgupta. "Coarse sample complexity bounds for active learning." *Advances in Neural Information Processing Systems*, Vancouver, Canada. MIT Press, 2006.
- [5] L. Fei-Fei and P. Perona. "A bayesian hierarchical model for learning natural scene categories." In *CVPR '05: Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition, 2005*, San Diego, CA.
- [6] Y. Freund, H.S. Seung, E. Shamir, and N. Tishby. "Selective sampling using the query by committee algorithm." *Machine Learning*, 28:133–168, 1997.
- [7] S.C. Hoi, R. Jin, J. Zhu, and M.R. Lyu. "Semi-supervised SVM batch mode active learning for image retrieval." In *CVPR '08: Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition, 2008*, Anchorage, AK.
- [8] A. Holub, P. Perona, and M. Burl. "Entropy-based active learning for object recognition." In *CVPR '08: Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition, Workshop on Online Learning for Classification, 2008*, Anchorage, AK.
- [9] C.W. Hsu and C.J. Lin. "A comparison of methods for multi-class support vector machines." *IEEE Transactions on Neural Networks*, 13:415–425, 2002.
- [10] A. Kapoor, K. Grauman, R. Urtasun, and T. Darrell. "Active learning with Gaussian Processes for object categorization." In *ICCV '07: Proceedings of the IEEE International Conference on Computer Vision, 2007*, Rio de Janeiro, Brazil.
- [11] M. Li and I. Sethi. "Confidence-based active learning." *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28:1251–1261, 2006.
- [12] H.T. Lin, C.J. Lin, and R.C. Weng. "A note on Platt's probabilistic outputs for support vector machines." *Machine Learning*, 68:267–276, 2007.
- [13] T. Mitchell. *Machine Learning*. Boston, MA. McGraw-Hill, 1997.

- [14] A. Oliva and A. Torralba. "Modeling the shape of the scene: A holistic representation of the spatial envelope." *International Journal of Computer Vision*, 42(3):145–175, 2001.
- [15] J. Platt. "Probabilistic outputs for support vector machines and comparison to regularized likelihood methods." *Advances in Large Margin Classifiers*, Vancouver, Canada, MIT Press, 2000.
- [16] G.J. Qi, X.S. Hua, Y. Rui, J. Tang, and H.J. Zhang. "Two-dimensional active learning for image classification." In *CVPR '08: Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition, 2008*, Anchorage, AK.
- [17] S. Tong and E. Chang. "Support vector machine active learning for image retrieval." In *MULTIMEDIA '01: Proceedings of the ninth ACM international conference on Multimedia, 2001*, Ottawa, Canada.
- [18] S. Tong, D. Koller, and P. Kaelbling. "Support vector machine active learning with applications to text classification." *Journal of Machine Learning Research*, 2:45–66, 2001.
- [19] T.F. Wu, C.J. Lin, and R.C. Weng. "Probability estimates for multi-class classification by pairwise coupling." *Journal of Machine Learning Research*, 5:975–1005, 2004.