

The Influence of Cross-Validation on Video Classification Performance

Jan C. van Gemert
ISLA, Informatics Institute
University of Amsterdam
Kruislaan 403, 1098 SJ
Amsterdam, The Netherlands
jvgemert@science.uva.nl

Cees G.M. Snoek
ISLA, Informatics Institute
University of Amsterdam
Kruislaan 403, 1098 SJ
Amsterdam, The Netherlands
cgmsnoek@science.uva.nl

Cor J. Veenman
ISLA, Informatics Institute
University of Amsterdam
Kruislaan 403, 1098 SJ
Amsterdam, The Netherlands
cveenman@science.uva.nl

ABSTRACT

Digital video is sequential in nature. When video data is used in a semantic concept classification task, the episodes are usually summarized with shots. The shots are annotated as containing, or not containing, a certain concept resulting in a labeled dataset. These labeled shots can subsequently be used by supervised learning methods (classifiers) where they are trained to predict the absence or presence of the concept in unseen shots and episodes. The performance of such automatic classification systems is usually estimated with cross-validation. By taking random samples from the dataset for training and testing as such, part of the shots from an episode are in the training set and another part from the same episode is in the test set. Accordingly, data dependence between training and test set is introduced, resulting in too optimistic performance estimates. In this paper, we experimentally show this bias, and propose how this bias can be prevented using *episode-constrained* cross-validation. Moreover, we show that a 17% higher classifier performance can be achieved by using episode constrained cross-validation for classifier parameter tuning.

Categories and Subject Descriptors

H.3.1 [Information Storage and Retrieval]: Content Analysis and Indexing—*Indexing methods*; I.5.2 [Pattern Recognition]: Design Methodology—*Classifier design and evaluation*

General Terms

Experimentation, Performance, Reliability

Keywords

Multimedia performance evaluation, cross-validation, parameter tuning, semantic concept detection

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MM'06, October 23–27, 2006, Santa Barbara, California, USA.
Copyright 2006 ACM 1-59593-447-2/06/0010 ...\$5.00.

1. INTRODUCTION

Machine learning techniques have proven to be a valuable addition to the repertoire of a multimedia researcher. Applications of machine learning techniques in multimedia are found in semantic video labeling [10], video shot detection [9], audio classification [5], scene recognition [12], sports analysis [2], emotion recognition [1], Meeting analyse [6], and in many other areas. Moreover, multimedia researchers have contributed to specifically designed classifiers for multimedia analysis [7].

A prerequisite to good classification results is an accurate estimation of classifier performance [3, 4]. The estimated classification performance may be used in finding the best parameters of the classifier and helps deciding between different features. Hence, the estimated classification performance determines the quality of the classification results.

This paper shows that commonly used classifier performance evaluation techniques need special care when applied to multimedia classification. Much multimedia data is sequential in nature. For example, popular music has a verse and a chorus, in a multimedia presentation the slides are designed with a story in mind and in video data there is the temporal ordering of shots. This paper will show that sequential data requires extra effort to accurately estimate classification performance. As an instantiation of sequential multimedia data, we will focus on classifier estimation of semantic concept detectors in video. However, the described techniques readily apply to other types of data that share a sequential ordering.

The organization of this paper is as follows. The next section revisits standard classifier evaluation techniques. Then, section 3 introduces our cross-validation method for video classification. In section 4 the experimental setup is discussed followed by the results in section 5, and the conclusions in section 6.

2. CLASSIFIER PERFORMANCE EVALUATION

The error estimation of a classifier not only provides a qualitative assessment of the classifier, it also influences the classifier's performance. The estimation of the performance of the classifier over different features determines which features are used, and which features might be left out. Furthermore, several classifiers require parameters, which are tuned by maximizing the estimated performance over various settings. For example in a video classification task,

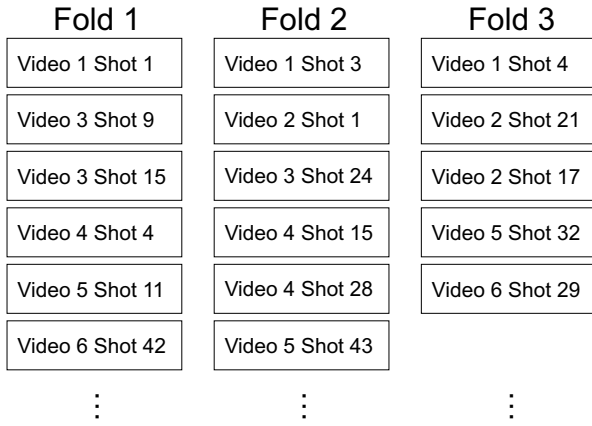


Figure 1: An example of partitioning a video set by using shot based 3-fold cross-validation.

Snoek et al. [10] use the estimated classifier performance to select the best low level features. Furthermore, they find the best parameters for a Support Vector Machine (SVM) [13] by maximizing the estimated classifier error. In their framework, inaccurate classifier performance estimation might result in choosing the wrong features, or in sub-optimal parameter settings. Hence, classifier error estimation affects the quality of the classifier.

Estimating the classification error is done by training the classifier on one set, and testing the classifier on a different set. Thus, a straightforward approach to classifier performance estimation is by keeping part of the available data in an unseen hold-out set. This hold-out set should be as large as possible, to accurately represent the class variation that may be expected. However, keeping a large part of the data from the training set hinders the classification performance. The advantage of using a hold-out set is that the test-data is completely separate from the training data. However, keeping a large training set to obtain a good classifier, counteracts a large hold-out set for accurate error estimation.

In contrast to having a single hold-out set, the cross-validation method rotates the hold-out set over all available data. Cross-validation randomly splits the available data in X folds, where each of these X folds is once used as a hold-out set. The error estimations of all folds are averaged, yielding the classifier error. The cross-validation procedure may be repeated R times, to minimize the effect of the random partitioning. An example of cross-validation for a video set is shown in figure 1. The advantage of using cross-validation is the combination of a large training set with several independent test sets. Therefore, cross-validation is the standard procedure for classification error estimation [3, 4].

3. CROSS-VALIDATION IN VIDEO CLASSIFICATION

Machine learning is heavily used in semantic video indexing. The aim of semantic video indexing is detecting all relevant shots in a dataset to a given semantic category. Some examples of semantic concepts are *Airplane*, *Car*, *Computer Screen*, *Hu Jintao*, *Military Vehicle*, *Sports*. Indexed semantic concepts provide a user with tools to browse, explore, and find relevant shots in a large collection of video. With growing digital video collections, there is a need for automatic



Figure 2: An example of narrative structure in video: five consecutive shots showing an interview with Lebanese President Émile Lahoud.

concept detection systems, providing instant access to digital collections. Therefore, machine learning techniques are vital to automatic video indexing.

In a video classification task, a shot is often the granularity of interest [8, 10]. However, a video document is the end result of an authoring process [10], where shots are used to convey a message. For example, a topic in news video, may consist of several shots, as shown in figure 2. Hence, a semantic concept might span several shots, while a video classification task is oriented towards single shots.

The mismatch between the granularity of the classification task and the granularity of the semantic concept requires special care in estimating classifier performance. Consider figure 2, and note the high similarity between shot 250 and shot 252. The similarity between these two shots can be expected, since they are part of the same narrative structure. However, the classification task focuses on single shots, and does not take this semantic relation between shots into account. Therefore, the common practice [8, 10] of estimating classifier performance by cross-validation on shots is biased. Cross-validation on shots will mix shots in a single topic to different folds while randomly partitioning the data. Thus, nearly identical shots will leak through to the rotating hold-out set. This leaking of near identical information creates a dependency between the training set and the hold-out set, which will manifest in too optimistic estimates for classifier performance. Moreover, if cross-validation is used for parameter tuning, the parameters will be biased towards near-duplicate data and might consequently fail to find the best classifier parameters for true independent hold-out data. Therefore, the sequential nature of video data should be taken into account when estimating classifier performance.

In order to preserve the semantic relation between shots

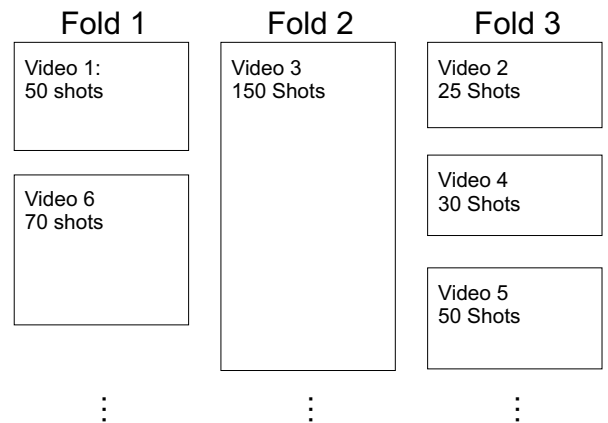


Figure 3: An example of a partitioning a video set by using episode-constrained 3-fold cross-validation.

	Shot Based	Episode-Constrained
Training set	0.525	0.309
Hold-out set	0.188	0.221

Table 1: The mean performance over all concepts, using the estimated parameters as selected by each method.

method selects $k = 3$ for 97 out of 100 concepts. Moreover, the estimated average precision scores for most concepts are disproportional high compared to the scores that are obtained on the hold-out set. Consider for example the concept *Aircraft*. The shot based cross-validation predicts a score of 0.393 where the best parameter is $k = 3$ neighbors. In contrast, the episode-constrained cross-validation predicts a score of 0.203 where the best parameter is $k = 2$ neighbors. Verifying the classification performance of *Aircraft* on hold-out data shows an average precision of 0.096 for $k = 3$ and a score of 0.132 for $k = 2$. Note that the performance estimate of episode-constrained cross-validation is not only closer to the hold-out performance, it also selects the best classifier parameter.

In table 1, we present the mean performance over all concepts, for both cross-validation methods. We show the estimated results on training data, and the results on hold-out data. The classifier parameter, the number of k -nearest neighbors, is tuned by selecting the maximum performance according to the cross-validation method at hand.

In analyzing table 1, we focus on two points: 1) the accuracy in estimating classifier performance and 2) the final classification performance. Starting with point 1, we consider the difference between the estimated performance on training data, and the reported performance on hold-out data. For shot based cross-validation there is considerable variation between the estimated performance on training data and the performance on hold-out data. Specifically, the performance estimate is 0.337 too optimistic. In contrast, for episode-constrained cross-validation the difference between training data and hold-out data is only 0.088. This clearly shows that the estimated performance of the episode-constrained cross-validation is more accurate than the performance estimate based on shots. Continuing with point 2, we compare the performance on hold-out data for both methods. It turns out that the episode-constrained method outperforms the shot based method by 17%. Analyzing the hold-out results per-concept, shows that episode-constrained cross-validation yields equal or better results for 93 out of 100 concepts, and it gives better results for 67 concepts. The shot based method gives the best results for 7 concepts. These results on hold-out data show that parameter tuning is considerably more accurate when using episode-constrained cross-validation.

6. CONCLUSIONS

In this paper, we compare two methods of cross-validation for estimating classification performance for semantic concept detection in video. The traditional method of cross-validation is based on shots, where we propose a method based on whole videos. This episode-constrained method for cross-validation prevents the leaking of similar shots to the rotating hold-out set. Experimental results show that the episode-constrained method yields a more accurate es-

timate of the classifier performance than the shot based method. Moreover, when cross-validation is used for parameter optimization, the episode-constrained method is able to better estimate the optimal classifier parameters, resulting in higher performance on validation data compared to shot based (traditional) cross-validation.

7. ACKNOWLEDGMENTS

This research is sponsored by the BSIK MultimediaN project.

8. ADDITIONAL AUTHORS

Additional authors: Arnold W.M. Smeulders (ISLA, Informatics Institute, University of Amsterdam, email: smeulders@science.uva.nl)

9. REFERENCES

- [1] R. Cowie, E. Douglas-Cowie, N. Tsapatsoulis, G. Votsis, S. Kollias, W. Fellenz, and J. G. Taylor. Emotion recognition in human-computer interaction. *Signal Processing Magazine*, 18(1):32–80, 2001.
- [2] L.-Y. Duan, M. Xu, X.-D. Yu, and Q. Tian. A unified framework for semantic shot classification in sports video. *Transactions on Multimedia*, 7(6):1066–1083, 2005.
- [3] R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern Classification*. Wiley-Interscience Publication, 2000.
- [4] A. K. Jain, R. P. W. Duin, and J. Mao. Statistical pattern recognition: A review. *TPAMI*, 22(1):4–37, 2000.
- [5] L. Lu, H. Jiang, and H. Zhang. A robust audio classification and segmentation method. In *ACM Multimedia*, pages 203–211, New York, NY, USA, 2001. ACM Press.
- [6] I. McCowan, D. Gatica-Perez, S. Bengio, G. Lathoud, and M. Barnard. Automatic analysis of multimodal group actions in meetings. *TPAMI*, 27(3):305–317, 2005.
- [7] M. R. Naphade and T. S. Huang. A probabilistic framework for semantic video indexing, filtering, and retrieval. *Transactions on Multimedia*, 3(1):141–151, 2001.
- [8] NIST. TRECVID Video Retrieval Evaluation, 2001–2005. <http://www-nlpir.nist.gov/projects/trecvid/>.
- [9] Y. Qi, A. Hauptmann, and T. Liu. Supervised classification for video shot segmentation. In *ICME*, July 2003.
- [10] C. G. M. Snoek, M. Worring, J.-M. Geusebroek, D. C. Koelma, F. J. Seinstra, and A. W. M. Smeulders. The semantic pathfinder: Using an authoring metaphor for generic multimedia indexing. *TPAMI*, 2006, in press.
- [11] C. G. M. Snoek, M. Worring, J. C. van Gemert, J.-M. Geusebroek, and A. W. M. Smeulders. The challenge problem for automated detection of 101 semantic concepts in multimedia. In *ACM Multimedia*, 2006.
- [12] J. C. van Gemert, J. M. Geusebroek, C. J. Veenman, C. G. M. Snoek, and A. W. M. Smeulders. Robust scene categorization by learning image statistics in context. In *CVPR-SLAM*, 2006.
- [13] V. Vapnik. *The Nature of Statistical Learning Theory*. Springer Verlag, New York, 1995.