Active Learning using Distribution Analysis for Image Classification

Dávid Papp, Gábor Szűcs

Nowadays, image classification is an indispensable tool for categorizing the huge amounts of visual contents both online and offline. The number of images available online increases with the rapidly growing internet usage. Besides, numerous electronic devices are capable to take a digital picture (e.g. cameras, telephones and so on), furthermore, smart devices are a click away to upload and share those pictures. This results massive data warehouses that are need to be structured, i.e. categorized. The classification of images requires labeled instances, but usually these contents are unlabeled, and labeling them is an expensive manual process. Active learning [1] is a way to address this problem since it selects a subset of the data by iteratively querying the most informative image(s) from the unlabeled ones, and then builds the classification model based on this subset instead of the whole data. In this way, the active learning algorithm aims to label as few instances as possible (i.e., minimizing the labeling cost) while it attempts to retain the same level of accuracy that could be achievable by using the total dataset. The most important question is how to estimate the informativeness of unlabeled instances, since different query strategies may lead to better or worse classification accuracy, compared to random sampling. There are many various proposed strategies in the literature, e.g. uncertainty sampling [2], query-by-committee [3], expected model change [4], expected error reduction [5]. Uncertainty sampling is a widely used, however the simplest query strategy framework, which happens to query the instances with least certainty about their labels. The focus of our work was to improve the accuracy through amplifying this technique by complementing it with a distribution analysis on the labeled dataset. For this purpose we defined a new penalty metric (see Eq. 1) which gave us an informativeness value for each unlabeled image:

$$Penalty_j = CTR_{j^*} \times \frac{1}{\#categories} \tag{1}$$

where j^* denotes the estimated category of the j^{th} image, and CTR_{j^*} increases with each iteration of queries where the received category is other then j^*. We merged these penalty values with the ones coming from uncertainty sampling, and this results the final decision scores. The advantage of this modification is a sort of balance between the classes of labeled instances. We demonstrated the efficiency of the proposed approach on two large datasets; on the PASCAL VOC2010 [6] training and validation data, and on the Caltech101 [7] image collection. We evaluated the Accuracy and the MAP (Mean Average Precision) metrics for this purpose. As can be seen in Fig. 1, our proposed approach (green line) outperforms both the general uncertainty sampling (red line) and the random sampling (blue line) methods.

Figure 1: Evaluation of the results on the Caltech101 and PASCAL VOC2010 datasets

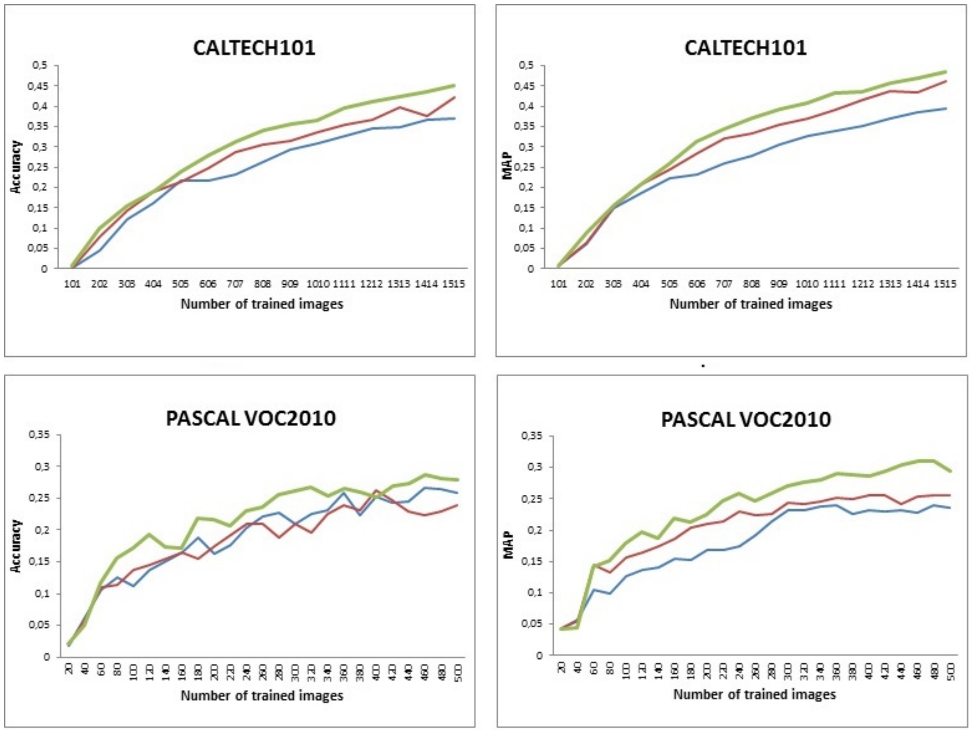

References

[1] Burr Settles, Active Learning Literature Survey, Computer Sciences Technical Report 1648, University of Wisconsin - Madison, 2009.

[2] Yang, Y., Ma, Z., Nie, F., Chang, X., & Hauptmann, A. G. (2015). Multi-class active learning by uncertainty sampling with diversity maximization. International Journal of Computer Vision, 113(2), 113-127.

[3] Tsai, Y. L., Tsai, R. T. H., Chueh, C. H., & Chang, S. C. (2014). Cross-Domain Opinion Word Identification with Query-By-Committee Active Learning. In Technologies and Applications of Artificial Intelligence (pp. 334-343). Springer International Publishing.

[4] Cai, W., Zhang, Y., & Zhou, J. (2013, December). Maximizing expected model change for active learning in regression. In Data Mining (ICDM), 2013 IEEE 13th International Conference on (pp. 51-60). IEEE.

[5] Mac Aodha, O., Campbell, N., Kautz, J., & Brostow, G. (2014). Hierarchical subquery evaluation for active learning on a graph. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (pp. 564-571).

[6] Everingham, M., Van Gool, L., Williams, C. K. I., Winn, J. and Zisserman, A. The PASCAL Visual Object Classes (VOC) Challenge, International Journal of Computer Vision, 88(2), 303 338, 2010.

[7] L. Fei-Fei, R. Fergus and P. Perona. Learning generative visual models from few training examples: an incremental Bayesian approach tested on 101 object categories. IEEE. CVPR 2004, Workshop on Generative-Model Based Vision. 2004.