

Active Learning and Discovery of Object Categories in the Presence of Unnameable Instances

Christoph Käding, Alexander Freytag, Erik Rodner, Paul Bodesheim, and Joachim Denzler
 Computer Vision Group, Friedrich Schiller University Jena, Germany, www.inf-cv.uni-jena.de

Motivation: Current visual recognition algorithms are “hungry” for data but massive annotation is extremely costly. Therefore, active learning algorithms are required that reduce labeling efforts to a minimum by selecting examples that are most valuable for labeling. In active learning, all categories occurring in collected data are usually assumed to be known in advance and experts should be able to label every requested instance. But do these assumptions really hold in practice? Could you name all categories in every image? A visual example is given in Figure 1.




Figure 1: The answers of an annotator are likely to be: (a) laptop (b) *what the heck is that?* (c) *no object!* (d) *some toy figure* (e) *no idea* (f) cat (g) cup (h) keyboard (i) cup. In an active learning and discovery scenario, an annotator might reject examples that do not show valid objects, that cannot be identified, or that are not part of the problem domain, which should be considered when requesting annotations.

Unnameable Instances: Existing active learning algorithms assume that an annotator always knows an answer, e.g., the requested label of a given image. From our experience, this assumption does *not* hold in practice and annotators are likely to reject labeling in two scenarios: either the sample does not belong to a valid category (e.g., lens artifacts, motion blur, or segmentations covering parts of multiple objects) or the sample is categorical, but unknown to the annotator or unrelated to the current labeling task. Ideally, active learning techniques should therefore be able to *discover new classes* and at the same time cope with queries an expert is *not able or willing to label*. To meet these observations, we present a variant of the expected model output change principle [2] for active learning and discovery in the presence of unnameable instances.

Expected model output changes: In this paper, we extend the active learning principle of expected model output changes (EMOC), which is an approximation to the expected decrease in classification error. With a given model f and an unlabeled example x' , EMOC estimates possible class labels $y' \in \mathcal{Y}$, updates the model under each possible label assumption to obtain f' , and compares differences in model outputs using a suitable loss function \mathcal{L} evaluated on all available data samples $x_j \in \mathcal{D}$:

$$\Delta f(x') = \sum_{y' \in \mathcal{Y}} \left(p(y'|f(x')) \cdot \frac{1}{|\mathcal{D}|} \sum_{x_j \in \mathcal{D}} \mathcal{L}(f(x_j), f'(x_j)) \right). \quad (1)$$

A sample that induces large changes should be preferred for being labeled, which is further supported by theoretical bounds on error reduction. When

we introduced the EMOC principle in our recent ECCV'14 paper [2], however, the presented technique was limited to binary classification scenarios and additionally required perfect oracles. We now extend the EMOC strategy in several ways:

1. We provide a suitable multi-class extension (see Eq. (5) in the full paper).
2. We additionally incorporate the data density to be unaffected by far-off non-categorical samples (see Eq. (15)).
3. We finally model possible rejections occurring from unnameable instances (see Eq. (16)), denoted by GP-EMOC_{PDE+R}.

Experimental results: We present experimental results in several application areas, including face identification [3], digit recognition [1], and object classification [4] – scenarios, where unnameable instance commonly occur. For comparison, we also include a variety of established techniques. Our experiments show that in these realistic scenarios, our approach substantially outperforms previous active learning methods, which are often not even able to improve with respect to the baseline of random query selection (detailed survey and description included in the full paper). An exemplary result is given in Figure 2.

Source code is available at <https://github.com/cvjena/>.

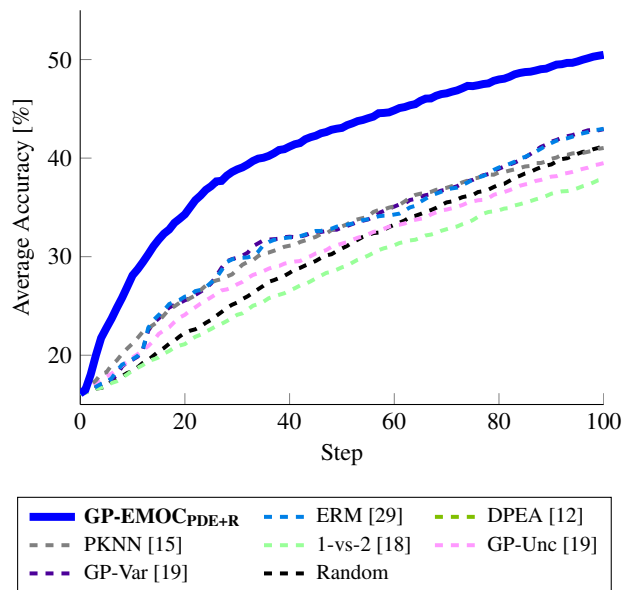


Figure 2: Improving recognition accuracy with active learning in an object classification scenario [4]. Baselines are indicated with dotted lines, whereas our technique is plotted solidly. See full paper for details.

[1] K. Bache and M. Lichman. UCI machine learning repository, 2013. URL <http://archive.ics.uci.edu/ml>.
 [2] A. Freytag, E. Rodner, and J. Denzler. Selecting influential examples: Active learning with expected model output changes. In *ECCV*, pages 562–577, 2014.
 [3] G. Huang, M. Ramesh, T. Berg, and E. Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst, 2007.
 [4] T. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, pages 740–755, 2014.