Beyond the closed-world assumption: The importance of novelty detection and open set recognition

Computer Vision Group Jena
Joachim Denzler, Erik Rodner, Paul Bodesheim, Alexander Freytag

Friedrich Schiller University Jena
{firstname.lastname}@uni-jena.de
http://www.inf-cv.uni-jena.de

Motivation: Current work on visual object recognition focuses on object classification and is implicitly based on the closed-world assumption, *i.e.,* a test sample is assigned to the most plausible class out of a fixed set of classes known during training. Knowledge about objects and classes is usually available in terms of representative training data and is used for model training. However, in real-world applications it is often not possible to obtain training data for all categories that can occur in the test phase beforehand. An example is quality control, where it is not only impossible to define all future defects – even worse, possible defects are in most cases not even known to the human expert who supervises the training step. In addition, even if one would know possible defects a priori, the small number of training images leads to ill-posed problems. A second application is life-long learning where a system needs to identify new, unknown objects classes and has to incrementally add them to its knowledge base. Finally, complex event detection in videos is also to impossible to tackle with a fixed set of classes. Although several solutions for the novelty detection problem have been proposed during the past years, they usually suffer from strong limitations (*e.g.,* model complexity), necessary assumptions (*e.g.,* Gaussian distribution), or heuristics (*e.g.,* separation from artificial negative data). On top of that, it is unknown so far whether or not such methods can successfully be applied in an open set scenario.

Problem specification: The actual problem is diverse and occurs on different levels of representation: (i) how can we determine whether or not a set of (local) features is novel with respect to the training set (*e.g.,* we've never seen a wheel before)? (ii) how can we determine the novelty of constellations of features (*e.g.,* a car upside down)? and (iii) how to measure the novelty of class constellations (*e.g.,* a car in a lake)? On top of that, the *level of novelty* for a specific application is not clear beforehand and needs to be well defined, *e.g.,* is a car within a lake novel, or abnormal, or still simply a car? What makes certain image regions to be considered as *novel* for a human inspector? And how can a system check for novelty and simultaneously classify objects which are not considered as novel?

In open set scenarios, novel objects are not only identified in comparison with the meta-class of all known objects, but known objects are additionally classified individually together with a label representing new classes or events. In terms of novelty detection, this situation has been investigated as *multi-class novelty detection* or under the more recently used term *open set recognition.* The main problems that arise in such scenarios are:

i) a reasonable definition of features to discriminate between known objects and to additionally separate them from unknown objects,
ii) the question, whether discriminative or generative models or a mixture of both are preferable or necessary,
iii) the incremental update of features and feature space as soon as new objects or events are identified,
iv) the question, how novelty can be measured, *i.e.,* whether or not specific distance measures or metrics need to be learned,
v) how novelty within class hierarchies can actually be defined, and
vi) whether context information (*e.g.,* class labels of surrounding regions) is beneficial, irrelevant, or even misleading for the open set recognition task.

The impact of results in this area is beyond pure computer vision applications, like object recognition. In modern disciplines, like biology, chemistry, etc. more and more image sensing devices are applied producing more and more image data. This mass of information can by no means be analyzed by hand. Consequently, methods that are able to identify known as well as unknown objects in such data would have an enormous, far-reaching impact on research in such areas. To summarize the problem in a single sentence: *how can we detect novel objects, events, or patterns and at the same time being able to distinguish known ones*?

Related work: In the literature, novelty detection is often handled as a one-class classification problem where all training samples share the same positive label indicating that they stem from a known category. The goal is then to model the empirical distribution of the training data such that it can be separated from the surrounding open space in each direction in the feature space. Popular approaches for one-class classification are one-class SVM [1], support vector data description (SVDD) [2], and methods based on Gaussian process regression [3, 4]. One-class SVM separates the training samples from the origin of the feature space with a maximum margin, support vector data description encloses the training data with a hypersphere of minimum volume, and the Gaussian process techniques are based on measures of the posterior distribution within a probabilistic regression framework. However, treating multiple known classes as a single one in the one-class setup disregards the information we have from the multi-class labels and leads to poor performance as shown in [5]. Even if we model each class by an individual one-class classifier as proposed in [6], novelty detection performance is rather low [5].

Promising approaches combine the idea of one-class classification with maximum margin separation between known classes. For binary classification, the small sphere and large margin approach of [7] encloses the positive class with a tight hypersphere as in SVDD while simultaneously maximizing the margin to the negative class. Recent work on open set recognition [8] uses the one-vs-rest multi-class SVM framework and estimates the separating hyperplane as well as a second hyperplane parallel to the first one in order to restrict the positive half space to a corridor between both hyperplanes. The authors propose using linear SVMs which lead to corridors that are still unbounded in some directions in the feature space. To avoid open class regions, the null space approach in [5] maps all training samples of the same class to a single point and maximizes distances between different classes. In contrast to unbounded (open) half spaces obtained from SVM models, representing classes by a single point directly encourages separation from unknown categories in every direction of the (transformed) feature space.

However, the novelty detection approaches currently available do not scale well in multi-class scenarios with an increasing number (hundreds to thousands) of categories. We believe that studying novelty detection and open set recognition in large-scale scenarios, where the feature space is occupied by more and more diverse known categories, is an important aspect for future research.

Conclusions We want to call computer vision researchers to tackle the task of open set recognition, where novelty detection and multi-class classification needs to be fused within a single recognition system. This leads to several challenges like balancing between novelty detection and recognition or the combination of generative and discriminative methods to achieve both exploration of new and separation of known categories. We think it is worth spending effort towards a closed problem formulation for tackling both tasks in a joint manner, which would allow for going in the direction of autonomous lifelong learning.

References

1. Schölkopf, B., Platt, J.C., Shawe-Taylor, J., Smola, A.J., Williamson, R.C.: Estimating the support of a high-dimensional distribution. Neural Computation **13**(7) (2001) 1443–1471
2. Tax, D.M.J., Duin, R.P.W.: Support vector data description. Machine Learning **54**(1) (2004) 45–66
3. Kemmler, M., Rodner, E., Denzler, J.: One-class classification with gaussian processes. In: Proceedings of the Asian Conference on Computer Vision (ACCV). (2010) 489–500
4. Bodesheim, P., Rodner, E., Freytag, A., Denzler, J.: Divergence-based one-class classification using gaussian processes. In: British Machine Vision Conference (BMVC). (2012) 50.1–50.11
5. Bodesheim, P., Freytag, A., Rodner, E., Kemmler, M., Denzler, J.: Kernel null space methods for novelty detection. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2013) (accepted for publication)
6. Tax, D.M.J., Duin, R.P.W.: Growing a multi-class classifier with a reject option. Pattern Recognition Letters **29**(10) (2008) 1565–1570
7. Wu, M., Ye, J.: A small sphere and large margin approach for novelty detection using training data with outliers. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) **31**(11) (2009) 2088–2092
8. Scheirer, W.J., Rocha, A., Sapkota, A., Boult, T.E.: Towards open set recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) (2013) (to appear)