

Minimizing Supervision in Multi-label Categorization

Rajat* Munender Varshney* Pravendra Singh Vinay P. Namboodiri
Department of Computer Science and Engineering, IIT Kanpur, India
{rajatrk, munender, psingh, vinaypn}@iitk.ac.in

Abstract

Multiple categories of objects are present in most images. Treating this as a multi-class classification is not justified. We treat this as a multi-label classification problem. In this paper, we further aim to minimize the supervision required for providing supervision in multi-label classification. Specifically, we investigate an effective class of approaches that associate a weak localization with each category either in terms of the bounding box or segmentation mask. Doing so improves the accuracy of multi-label categorization. The approach we adopt is one of active learning, i.e., incrementally selecting a set of samples that need supervision based on the current model, obtaining supervision for these samples, retraining the model with the additional set of supervised samples and proceeding again to select the next set of samples. A crucial concern is the choice of the set of samples. In doing so, we provide a novel insight, and no specific measure succeeds in obtaining a consistently improved selection criterion. We, therefore, provide a selection criterion that consistently improves the overall baseline criterion by choosing the top k set of samples for a varied set of criteria. Using this criterion, we are able to show that we can retain more than 98% of the fully supervised performance with just 20% of samples (and more than 96% using 10%) of the dataset on PASCAL VOC 2007 and 2012. Also, our proposed approach consistently outperforms all other baseline metrics for all benchmark datasets and model combinations.

1. Introduction

Deep learning techniques have shown tremendous efficacy in various tasks such as Image classification [19, 31, 13], Object Detection [11] and image segmentation [3, 4, 26]. We are particularly concerned with the object categorization task that is commonly treated as a multi-class classification task. However, most images commonly contain multiple categories simultaneously being present.

Therefore, treating the categorization task as a multi-class classification would not be justified. In fact, it could be treated as such only if we consider an ensemble of binary classifiers, each considering the presence or absence of each individual category. A more natural approach would treat it as a *multi-label* classification task. A different way could be using a more structured recognition task such as object detection or segmentation and using full supervision for solving the task. However, a distinct approach is one that considers it as a multi-label categorization task, yet, associates a localization aspect associated with it, such as detection [8] or segmentation [7]. These methods have shown good performance for classification, indeed, better than most classification methods for obtaining multi-label classification with a smaller number of samples. In this paper, we solve for minimizing the supervision required for realizing these methods.

The localization based multi-label categorization models such as WELDON and WILDCAT use similar pre-trained backbone (ResNet-101 [31, 13]) to extract features. They use spatial pooling to select relevant regions in the final pooling layer to compute the final score for each class. They show that by localizing positive evidence for an object class, they are able to obtain better performance as compared to standard models that use the entire regions/activation for final prediction. The need for obtaining supervision for all the multiple labels for each instance remains.

The approach proposed by us is to incorporate an active learning strategy to minimize the amount of supervision required. In active learning, based on the current model, one chooses a set of samples that need to be labeled. The ground-truth labels for these samples are obtained by an annotator, and the model is further retrained to obtain the next of samples. The selection of samples thus forms the core element in any active learning strategy. Unfortunately, our experiments suggest that no particular strategy or measure for active learning consistently suffices to provide a set of samples. To address this drawback, we provide a simple strategy that is performant. We choose a varied set of active learning methods, such as those based on entropy and classification uncertainty. We augment these strategies with

*Equal contribution.

a couple of strategies that consider the localization aspects obtained through the specific localization based categorization methods. Given this set of strategies, if we need to choose n samples given k strategies, we obtain the ordered set of top n/k samples that are unique using the k strategies. We show that this simple method outperforms any individual strategy, costs similar to any single baseline strategy, and is also better than other variants for obtaining an ensemble. We term this strategy a ‘metric agnostic’ sampling strategy.

2. Related Work

Object recognition and detection have been extensively studied and achieved state-of-the-art results for the deep learning models such as CNN [7, 12, 11, 25, 8, 19]. The earlier method in object detection using deep learning is based on explicitly generating region proposals, following by feature extraction, and classification [11, 28]. The newer method outputs the location of the object and class directly without any intermediate proposal stage, such as SSD [25] and YOLO [27]. Due to these marvelous performances, CNN models have been extensively used to solve multi-label object recognition problem by localization using region proposal [25, 11, 12]. However, localization in CNN models needs bounding box annotation for large dataset [12, 25, 11], which is a laborious and costly job. Due to this laborious work, the weakly supervised learning models got huge attention [7, 8]. These models use spatial pooling with Fully convolutional networks, which helps in localizing discriminative regions [38, 7]. But, the annotation cost to label all objects in the dataset is huge, so active learning is used to label a smaller set of images.

Active learning has been wide used in various computer vision task where the labeled data is scarce and difficult to obtain such as image classification, segmentation, and detection [15], automated and manual video annotation [18], and CAPTCHA recognition [33], segmentation [15, 17]. Specifically active learning has also helped in object classification [15, 17, 37, 16] and object detection [2]. These tasks have been explored for various specific applications where annotated data is scarce, such as vehicle images [32] and satellite images [1]. A detailed survey of active learning algorithms, which are used to select fewer labeled training data for achieving high accuracy, are present in the settles’s survey [29]. The approaches developed so far doesn’t include any active learning work (multi-label classification) for deep learning settings. Active learning concepts relevant for this work include uncertainty sampling [33, 15, 37] and query-by-committee [30]. These active learning methods have helped in solving computer vision problems. These method also find application in different learning models including Gaussian processes [17], decision trees [20] and SVMs [34].

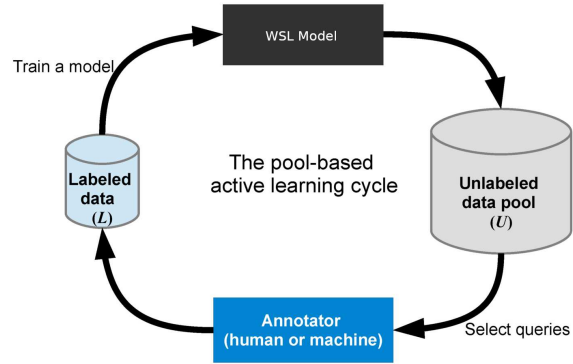


Figure 1: An iteration of active learning.

Even though active learning has been used in machine learning for decades, it has only recently been applied to deep learning architectures. Wang *et al.*, in [36] and [35] has proposed to use uncertainty based active learning for deep network model. Several metrics are offered to estimate the uncertainty in the model, such as margin, least confidence, or entropy sampling. Starket *et al.*, [33] proposed another variation of margin sampling. Another work in paper [24] presents uncertainty sampling. In addition, a query by committee strategy, as well as weighted incremental dictionary learning for active learning, is proposed. Gal *et al.* [10] used Monte Carlo dropout to measure Bayesian uncertainty for selecting samples in CNN models. To the best of our knowledge, this is the first work that addresses the issue of multi-label object recognition in a deep learning setting using active learning. However, there are some good work has been done non-deep setting for multi-label active learning which leverage than the correlation between labels [22, 14, 6].

3. Active Learning for WSL (Weakly Supervised Learning) models

A trained multi-label classifier takes an input image and outputs the probability of the presence of one or more objects of a class for each of the predefined categories. To train this model, we need images that are annotated on the global level, i.e., class-level annotations. Such annotations are available in public datasets like PASCAL VOC [9] and MS COCO [23]. We implement an active learning algorithm in the way shown in Fig. 1, where the baseline model is trained with a small amount of annotated samples. The baseline model is then used over the unlabelled dataset to give a pool of most informative samples using the given active learning metric. These pooled unlabelled samples are then annotated using human annotators and then added to the training set to train a new model. This cycle is an iteration of an active learning algorithm, which can be repeated

multiple times until the annotation capability or the desired level of performance is reached.

We first take a look at our proposed foreground-background active learning metric in Sec. 3.1, which utilizes the WSP module of the WELDON [8] and the multi-map WSL transfer layer of WILDCAT [7] model to select the most informative samples from the unlabelled data pool. After that, we look at the metric agnostic approach (Sec. 3.2), which utilizes existing active learning metrics to pool the best samples.

3.1. Foreground Background Separation

The main idea for this active learning metric utilizes the way in which these WSL models, both WELDON [8] and WILDCAT [7] are trained and how they perform classification.

In the WELDON architecture, the network consists of two parts, and the first part acts as a feature extractor having a CNN like ResNet [13] or VGG [31] trained on a large scale dataset such as ImageNet [5], and the second part which contains the WSP module which uses final convolution layer followed by a spatial pooling to make classification predictions. This final WELDON spatial pooling, which improves upon the MIL [21] and gets insights from it, pools both top instances as well as bottom instances, after which it averages these values to obtain the final score. The reason to perform such operation, being that the top instances correspond to localized positive evidence for the presence of an object, and the bottom instances correspond to negative evidence for its absence. Thus, the model’s ability to make a correct prediction for the presence/absence of an object is dependent on its ability to discern the foreground (localized positive evidence) from the background (negative evidence) in the given sample. Thus, we reason that samples in which the model can easily discern the foreground from the background, which here is being quantified by the difference between the value of maximum top instances and minimum negative instances, will not contain much new information for the model to learn from. We use this reasoning to propose an active learning metric which selects samples based on how well the model can discern foreground from the background instead of using the model’s final output like the other traditional active learning metrics.

The proposed Foreground-Background approach can be applied over WSL models or similar kinds of architecture only, i.e., WILDCAT architecture. The WILDCAT architecture, which similar to the WELDON architecture, also has two parts a feature extractor and a prediction network. The latter contains the multi-map WSL transfer layer. Unlike the WELDON network, WILDCAT uses two different kinds of pooling to make the final prediction. These pooling methods include a class-wise pooling, which is then fol-

lowed by a spatial pooling, latter of which is similar to that used in WELDON. In this network architecture, we again quantify the foreground and background separation as defined above, using the tensor/output obtained after performing the class-wise pooling on the WSL transfer layer. The proposed Foreground-Background approach is only limited by choice of architecture and not the task.

Using the metric as explained above, we select samples to annotate from the unlabelled dataset for which the model performance on the foreground and background separation is weak, i.e., the samples for which difference between the maximum of top and minimum of bottom instances are small. We also use different aggregation techniques for combining this foreground-background separation performance across different classes. For more detail, see Sec. 5.

Algorithm 1 Metric Agnostic Approach

```

1: Input:
2:  $S = [S_{M1}, S_{M2}, S_{M3}, \dots, S_{MT}]$  where  $S_{Mi}$  is list of
   selected samples sorted using metric  $Mi$ 
3: Number of samples to be selected  $n$ 
4: Output
5: List of samples given by metric agnostic approach,  $\mathcal{S}$ .
6: Algorithm:
7: Create an empty set,  $\mathcal{S} = \phi$ 
8: Create an index list,  $idxs = [1, 1, \dots, 1]$ 
9: while  $|\mathcal{S}| < n$  do
10:   for  $t = 1$  to  $T$  do
11:     if  $idxs[t] \leq len(S_{Mt})$  then
12:        $\mathcal{S} = \mathcal{S} \cup S_{Mt}[idxs[t]]$ 
13:        $idxs[t]++$ 
14:     if  $|\mathcal{S}| \geq n$  then
15:       break
16: return  $\mathcal{S}$ 

```

3.2. Metric Agnostic Approach (AG)

From the evaluation of various active learning techniques, we observed that the performance of these techniques (metrics) depends on various conditions (amount of data, dataset, and model). We, therefore, propose a new method that could combine the best set of samples obtained from each metric. This new method we term as being a “metric agnostic” approach as it can be thought of as an algorithm that is based on other active learning metrics.

In the proposed metric agnostic approach, the main idea of the proposed method is to pool samples from samples selected using other benchmark active learning metrics. Specifically, we take the top picks of each of the individual active learning metrics in equal proportion and then annotate these aggregated samples and then add them to the training dataset. The way we pool samples from different

metrics is we start with lists of samples selected by each of the metrics, sorted on some priority based on how that metric selects the samples. We then start with an empty set, and for each iteration, we select one sample from each of the list, which has not already been selected (by some other metric) until we have reached our predefined limit as shown in Algorithm 1. The final set of unlabelled samples obtained this way is considered as the final selection of the metric agnostic approach (see Alg. 1).

It should be noted that though this paper shows the usage of the metric agnostic approach with WSL models for multi-label classification. But, the metric agnostic approach is neither bound to the model type nor the task at hand. Hence it can be used with all kinds of models and tasks. Exploration of this capability of the metric agnostic approach will be done as part of future work.

4. Baseline Active Learning Metrics

We compare our proposed active learning metrics against various other widely used benchmark active learning metrics for the classification task, as well as with passive learning. We define $P(y_j | x_i)$ as the model’s confidence for the presence of class j in the i^{th} sample, these confidence values are obtained by taking the sigmoid of the model’s output.

4.1. Classification Uncertainty (UNC)

As defined $P(y_j | x_i)$ denotes the model confidence for the presence of class j in the i^{th} sample, we can measure the model’s certainty for class j as $abs(P(y_j | x_i) - 0.5)$, as a value of $P(y_j | x_i)$ close to 0 denotes that the model is certain the class is not present and a value of $P(y_j | x_i)$ close to 1 denotes that model is certain that the class is present. We then aggregate these certainties value across different classes using sum operator and select samples with following strategy $\min_i \sum_j abs(P(y_j | x_i) - 0.5)$.

4.2. Maximum Entropy (ENT)

In this active learning metric we select samples with maximum entropy values. The query strategy is defined as: $\max_i - \sum_j P(y_j | x_i) \log P(y_j | x_i)$

4.3. Min-Max (MM)

In this active learning metric, we select samples with minimum-maximum model confidence across all classes, defined as $\min_i \max_j P(y_j | x_i)$

4.4. Passive Learning (R)

In this strategy, we randomly select samples from an unlabelled dataset, annotate them, and add to the training set.

5. Experiments

We evaluate our proposed metrics on both WELDON [8] and WILDCAT [7] models. All our experiments were performed using the PyTorch framework. For all our experiments we used authors implementation of the model which are available online, WELDON (<https://github.com/durandtibo/weldon.resnet.pytorch/>) and WILDCAT (<https://github.com/durandtibo/wildcat.pytorch/>). The final results for each of the metrics are obtained by performing three trials using three different sets of randomly selected images to train the baseline model. Also, for the baseline and the first three of the five iterations, the model was trained for 35 epochs, and for the last two iterations, the model was trained for 25 epochs. For each active learning metric, we show relative(%) mAP of the model to that trained on the entire dataset in each graph.

For the foreground-background separation metric described in Sec. 3.1, we use three different aggregation methods for experimentation.

- **SEPSUM**, the performance of the model on a sample is represented using the sum of the foreground-background separation value of all classes.
- **SEPMAX**, the performance of the model on a sample is represented by the maximum value of foreground-background separation across all classes.
- **SEPMIN**, the performance of the model on a sample is represented by the minimum value of foreground-background separation across all classes.

For each of the variants, we select those samples from the unlabelled dataset on which the model’s performance is poor, i.e., the samples with a minimum value of **SEPSUM**, **SEPMAX** and **SEPMIN**. Finally, all the experiments with the metric agnostic approach are performed, as explained in Sec. 3.2.

Datasets: The performance of our proposed active learning approach has been evaluated on three benchmark datasets such as PASCAL VOC 2007, PASCAL VOC 2012, and MS COCO [9, 23]. For all classification experiments protocols used in [8] were followed. For each of these datasets, we start with a baseline model trained on a small number of randomly selected images from the training set, and the remaining training dataset is then used as the pool of unlabelled images from which we pool images in small increment for multiple iterations using active learning metrics. Since the annotations for the training set of each of the datasets are available, we provide the annotations for the selected images without using human annotators. For each of the experiments, we use up to 20% of the training dataset to train the model and then compare its performance with the model trained on the entire training set.




Figure 2: Results for PASCAL VOC 2007 using 500 images and WELDON. For each active learning metric we show relative(%) mAP of model to that trained on entire dataset.




Figure 3: Results for PASCAL VOC 2007 using 1000 images and WELDON.




Figure 4: Results for PASCAL VOC 2007 using 500 images and WILDCAT.

5.1. PASCAL VOC 2007

Experimental Setup: For Pascal VOC 2007 dataset, we begin with a baseline model trained on 100 randomly selected images from the training set (containing ~ 5000 images). Then for each active learning metric (excluding passive learning case), we add new images in five iterations adding respectively 100, 100, 100, 100, and 500 images in each iteration training the final model using 1000 images.

Results: For WELDON using 500 images in Fig. 2 and using 1000 images in Fig. 3. For WILDCAT using 500




Figure 5: Results for PASCAL VOC 2007 using 1000 images and WILDCAT.

images in Fig. 4 and using 1000 images in Fig. 5. The images show the performance of the model using each of the active learning metrics relative to the performance of the model trained on the entire training dataset.

We have three main observations from the results obtained on the PASCAL VOC 2007 dataset, first being that when using 500 images, our proposed foreground-background separation metric (in this case **SEPMIN**) outperforms all other metrics (except for metric agnostic) while using WELDON network. This metric helps the WSL model to perform much better while using a relatively small amount of data. Secondly, the performance of active learning metrics varies greatly with change in model, as evident with the performance of uncertainty(UNC) metric which goes from being the worst-performing metric while using WELDON (Fig. 3) to the best performing metric (except for metric agnostic) while using the WILDCAT model (Fig. 5). We also observe a significant difference in performance as a result of the change in the amount of data while using the same model. For the case of both WELDON (Fig 2 and Fig. 3) and WILDCAT (Fig. 4 and Fig. 5) comparative performance of all the metrics changes significantly as we go from 500 images to 1000 images. The third important observation is the consistent performance of our metric agnostic approach across all cases. The metric agnostic approach outperforms all other active learning metrics, especially for the case when 500 images are being used to train the model. This goes to show that apart from the consistent performance across different scenarios, the metric agnostic approach, when used in scenarios with limited annotation capability, can help get a non-trivial improvement over the performance obtained from using any single active learning metric.

5.2. PASCAL VOC 2012

Experimental Setup: For VOC 2012 dataset we begin with a baseline model trained on 95 randomly selected images from the training set (containing ~ 5700 images). Then for each active learning metric (excluding passive




Figure 6: Results for PASCAL VOC 2012 using 570 images and WELDON.




Figure 9: Results for PASCAL VOC 2012 using 1140 images and WILDCAT.




Figure 7: Results for PASCAL VOC 2012 using 1140 images and WELDON.




Figure 10: Results for MS COCO using 8200 images and WELDON.




Figure 8: Results for PASCAL VOC 2012 using 570 images and WILDCAT.

learning case) we add new images in six iterations adding respectively 95, 95, 95, 95, 95 and 570 images, training the final model using 1140 images. Final results is reported similar to VOC 2007 Sec. 5.1

Results: For WELDON using 570 images in Fig. 6 and using 1140 images in Fig. 7. For WILDCAT using 570 images in Fig. 8 and using 1140 images in Fig. 9. Similar to the case of the PASCAL VOC 2007 case, when using 570 images, our proposed foreground-background separation metric outperforms all other metrics (except for metric agnostic) not only while using WELDON network (**SEPSUM**) but also for the case for WILDCAT network

(**SEPMIN**). This further makes a case for the suitability of this metric is for WSL networks. The performance of active learning metrics again varies with the change in model and the amount of data being used to train the model. Also similar to VOC 2007 case, the metric agnostic approach consistently outperforms all other active learning metrics, especially for the case when using 570 images to train the model.

5.3. MS COCO

Experimental Setup: For the MS COCO dataset, we begin with a baseline model trained on 1640 randomly selected images from the training set (containing ~ 82000 images). Then for each of the active learning metrics (excluding passive learning case), we add new images in five iterations adding respectively 1640, 1640, 1640, 1640, and 8200 images, training the final model using 16400 images.

Results: For WELDON using 8200 images in Fig. 10 and using 16400 images in Fig. 11. For WILDCAT using 8200 images in Fig. 12 and using 16400 images in Fig. 13. When using 8200 images, our proposed foreground-background separation metric (**SEPMAX**) outperforms all other metrics (except for metric agnostic) not only while using WELDON network but also for the case for WILDCAT network. Also, when using 16400 images our proposed




Figure 11: Results for MS COCO using 16400 images and WELDON.




Figure 14: Vote distribution obtained using WELDON for both PASCAL VOC 2007 and 2012 dataset. Active learning metrics being used are **UNC**, **ENT**, **MM**, **SEPMAX**, **SEPMIN** and **SEPSUM**.




Figure 12: Results for MS COCO using 8200 images and WILDCAT.




Figure 15: Vote distribution obtained using WILDCAT for both PASCAL VOC 2007 and 2012 dataset. Active learning metrics being used are **UNC**, **ENT**, **MM**, **SEPMAX**, **SEPMIN** and **SEPSUM**.




Figure 13: Results for MS COCO using 16400 images and WILDCAT.

foreground-background separation metric outperforms all other metrics (except for metric agnostic) for both WELDON (**SEPMIN**) and WILDCAT (**SEPMAX**). This goes to show that our proposed foreground-background metric is able to scale well for large datasets. Similar to what we observed for VOC datasets, the performance of active learning metrics again varies with the change in model and amount of data being used to train the model. Also, similar to the VOC datasets case, the metric agnostic approach consistently outperforms all other active learning metrics.

5.4. Ablation Study

Various ablation studies have been conducted to illustrate various aspects of the proposed metric agnostic approach. First, We compare our aggregation strategy of the metric agnostic method explained in Sec. 3.2, to another voting based aggregation approach in which for each of the sample selected by one or more active learning metric, we treat the number of active learning metrics that selected that sample as votes for that sample, i.e., sample selected by only one metric will get one vote and so on. Finally, samples are selected based on the decreasing order of the number of votes. The vote distribution we obtained using WELDON is shown in Fig. 14 and using WILDCAT is shown in Fig. 15. From the results obtained, we see that more than half of the samples receive just one vote suggesting that, for the most part, there is little agreement between different active learning metrics on the importance of a sample. Furthermore, the classification performance obtained using this voting approach was sub-par to that obtained using our proposed approach.

The second ablation study is to show that though our

Dataset	VOC 2007	VOC 2012
UNC (sec)	40.888	50.415
ENT (sec)	40.868	50.335
MM (sec)	40.771	50.429
SEPMAX (sec)	41.557	51.328
SEPMIN (sec)	41.728	51.442
SEPSUM (sec)	41.490	51.496
AG (sec)	43.385	53.454
Δ_{AG} (%)	5.259	5.002

Table 1: Time analysis of active learning metrics for both PASCAL VOC 2007 and 2012 using WELDON. Δ_{AG} is the relative(%) time difference between **AG** and average of all other metrics.

Dataset	VOC 2007	VOC 2012
UNC (sec)	101.966	120.388
ENT (sec)	101.756	120.467
MM (sec)	102.035	120.338
SEPMAX (sec)	101.916	120.971
SEPMIN (sec)	102.220	120.957
SEPSUM (sec)	102.755	121.367
AG (sec)	106.557	124.166
Δ_{AG} (%)	4.357	2.831

Table 2: Time analysis of active learning metrics for both PASCAL VOC 2007 and 2012 using WILDCAT. Δ_{AG} is the relative(%) time difference between **AG** and average of all other metrics.

metric agnostic approach is an ensemble-based, the additional time cost for this ensemble is relatively quite low. The computation for each of the active learning metrics being used under it can be done with a single pass through the network. Table 1 shows the results of our time analysis using WELDON using an NVIDIA GeForce GTX 1080 graphics card. Table 2 shows result for WILDCAT model. For both models, we are pooling 100 images from the PASCAL VOC 2007 dataset(from ~ 4900 images), and for PASCAL VOC 2012, we are pooling 95 images (from ~ 5600 images).

The third ablation study is to show that effective selection of samples based on model behavior is critical and cannot be replaced by “naively” selecting samples that contain the most number of classes to maximize the number of examples of each class a model sees. To test this hypothesis, we created a sort of “adversarial”(ADV) dataset in which samples are added in decreasing order of the number of classes it contains. The performance comparison of this adversarial method with our metric agnostic method is shown in Fig. 16 and Fig. 17, for both cases, we see that the metric agnostic method outperforms the adversarial method.




Figure 16: Performance comparison of Metric Agnostic with Adversarial on both PASCAL VOC 2007 (using 1000 images) and 2012 (using 1140 images) using WELDON.




Figure 17: Performance comparison of Metric Agnostic with Adversarial on both PASCAL VOC 2007 (using 1000 images) and 2012 (using 1140 images) using WILDCAT. For both method we show relative(%) mAP of model to that trained on entire dataset.

6. Conclusions

The two active learning metrics are presented for multi-label object recognition in WSL (Weakly Supervised Learning) models. Our foreground-background separation metric pools images based on the model’s performance on being able to discern foreground from the background. Our second metric agnostic metric utilizes the existing active learning metrics to make its selection. Through extensive experimentation, we then show that the foreground-background separation metric outperforms existing active learning metrics, especially when the amount of annotations available is limited and that this metric scales well for the case of large datasets like MS COCO. The metric agnostic approach is shown to consistently outperform all other active learning metrics for various models and dataset combinations, indicating its robustness to change in datasets and models. Future works include adapting these proposed metrics to other domains of visual tasks like semantic segmentation.

References

- [1] Alberto Bietti. Active learning for object detection on satellite images. Technical report, Technical report, Caltech, 2012.
- [2] Clemens-Alexander Brust, Christoph Käding, and Joachim Denzler. Active learning for deep object detection. *arXiv preprint arXiv:1809.09875*, 2018.
- [3] L. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Semantic Image Segmentation with Deep Convolutional Nets and Fully Connected CRFs. In *The International Conference on Learning Representations (ICLR)*, 2015.
- [4] J. Dai, Y. Li, K. He, and J. Sun. R-FCN: Object detection via region-based fully convolutional networks. In *Neural Information Processing Systems (NIPS)*, 2016.
- [5] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [6] Bo Du, Zengmao Wang, Lefei Zhang, Liangpei Zhang, and Dacheng Tao. Robust and discriminative labeling for multi-label active learning based on maximum correntropy criterion. *IEEE Transactions on Image Processing*, 26(4):1694–1707, 2017.
- [7] Thibaut Durand, Taylor Mordan, Nicolas Thome, and Matthieu Cord. WILDCAT: Weakly Supervised Learning of Deep ConvNets for Image Classification, Pointwise Localization and Segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [8] Thibaut Durand, Nicolas Thome, and Matthieu Cord. WELDON: Weakly Supervised Learning of Deep Convolutional Neural Networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [9] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes (VOC) challenge. *International Journal of Computer Vision*, 88(2):303–338, 2010.
- [10] Yarin Gal, Riashat Islam, and Zoubin Ghahramani. Deep bayesian active learning with image data. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, pages 1183–1192. JMLR. org, 2017.
- [11] Ross Girshick. Fast r-cnn. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 1440–1448. IEEE, 2015.
- [12] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [13] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [14] Sheng-Jun Huang, Songcan Chen, and Zhi-Hua Zhou. Multi-label active learning: Query type matters. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.
- [15] Riashat Islam. *Active learning for high dimensional inputs using bayesian convolutional neural networks*. PhD thesis, Master’s thesis, Department of Engineering, University of Cambridge, 8 2016. 2, 2016.
- [16] Ajay J Joshi, Fatih Porikli, and Nikolaos Papanikolopoulos. Multi-class active learning for image classification. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2372–2379. IEEE, 2009.
- [17] Ashish Kapoor, Kristen Grauman, Raquel Urtasun, and Trevor Darrell. Gaussian processes for object categorization. *International journal of computer vision*, 88(2):169–188, 2010.
- [18] Vasilii Karasev, Avinash Ravichandran, and Stefano Soatto. Active frame, location, and detector selection for automated and manual video annotation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2123–2130, 2014.
- [19] A. Krizhevsky, I. Sutskever, and G. Hinton. Imagenet classification with deep convolutional neural networks. In *Neural Information Processing Systems (NIPS)*, 2012.
- [20] David D Lewis and Jason Catlett. Heterogeneous uncertainty sampling for supervised learning. In *Machine learning proceedings 1994*, pages 148–156. Elsevier, 1994.
- [21] Weixin Li and Nuno Vasconcelos. Multiple instance learning for soft bags via top instances. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.
- [22] Xin Li and Yuhong Guo. Active learning with multi-label svm classification. In *Twenty-Third International Joint Conference on Artificial Intelligence*, 2013.
- [23] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, Lubomir D. Bourdev, Ross B. Girshick, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. *CoRR*, abs/1405.0312, 2014.
- [24] Peng Liu, Hui Zhang, and Kie B Eom. Active deep learning for classification of hyperspectral images. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 10(2):712–724, 2017.
- [25] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016.
- [26] J. Long, E. Shelhamer, and T. Darrell. Fully Convolutional Networks for Semantic Segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [27] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.
- [28] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [29] Burr Settles. Active learning literature survey. Computer Sciences Technical Report 1648, University of Wisconsin-Madison, 2009.

- [30] H Sebastian Seung, Manfred Opper, and Haim Sompolinsky. Query by committee. In *Proceedings of the fifth annual workshop on Computational learning theory*, pages 287–294. ACM, 1992.
- [31] K. Simonyan and A. Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *The International Conference on Learning Representations (ICLR)*, 2015.
- [32] Sayanan Sivaraman and Mohan M Trivedi. Active learning for on-road vehicle detection: A comparative study. *Machine vision and applications*, 25(3):599–611, 2014.
- [33] Fabian Stark, Caner Hazırbaş, Rudolph Triebel, and Daniel Cremers. Captcha recognition with active deep learning. In *GCPR Workshop on New Challenges in Neural Computation*, volume 10, 2015.
- [34] Simon Tong and Daphne Koller. Support vector machine active learning with applications to text classification. *Journal of machine learning research*, 2(Nov):45–66, 2001.
- [35] Dan Wang and Yi Shang. A new active labeling method for deep learning. In *2014 International joint conference on neural networks (IJCNN)*, pages 112–119. IEEE, 2014.
- [36] Keze Wang, Dongyu Zhang, Ya Li, Ruimao Zhang, and Liang Lin. Cost-effective active learning for deep image classification. *IEEE Transactions on Circuits and Systems for Video Technology*, 27(12):2591–2600, 2017.
- [37] Yi Yang, Zhigang Ma, Feiping Nie, Xiaojun Chang, and Alexander G Hauptmann. Multi-class active learning by uncertainty sampling with diversity maximization. *International Journal of Computer Vision*, 113(2):113–127, 2015.
- [38] Xiaolin Zhang, Yunchao Wei, Jiashi Feng, Yi Yang, and Thomas S Huang. Adversarial complementary learning for weakly supervised object localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1325–1334, 2018.