

Deep Active Learning with Noisy Oracle in Object Detection

Marius Schubert^{1,*}^a, Tobias Riedlinger^{2,*}^b, Karsten Kahl¹^c and Matthias Rottmann¹^d

¹*School of Mathematics and Natural Sciences, IZMD, University of Wuppertal, Germany*

²*Institute of Mathematics, Technical University Berlin, Germany*

Keywords: Active Learning, Label Noise, Robustness, Label Error Detection, Object Detection.

Abstract: Obtaining annotations for complex computer vision tasks such as object detection is an expensive and time-intensive endeavor involving numerous human workers or expert opinions. Reducing the amount of annotations required while maintaining algorithm performance is, therefore, desirable for machine learning practitioners and has been successfully achieved by active learning. However, it is not merely the amount of annotations which influences model performance but also the annotation quality. In practice, oracles that are queried for new annotations frequently produce significant amounts of noise. Therefore, cleansing procedures are oftentimes necessary to review and correct given labels. This process is subject to the same budget as the initial annotation itself since it requires human workers or even domain experts. Here, we propose a composite active learning framework including a label review module for deep object detection. We show that utilizing part of the annotation budget to correct the noisy annotations partially in the active dataset leads to early improvements in model performance, especially when coupled with uncertainty-based query strategies. The precision of the label error proposals significantly influences the measured effect of the label review. In our experiments we achieve improvements of up to 4.5*mAP* points by incorporating label reviews at equal annotation budget.

1 INTRODUCTION

In the previous decade, deep learning has revolutionized computer vision models across many different tasks like supervised object detection (Ren et al., 2015; Redmon and Farhadi, 2018; Carion et al., 2020). Object detection has various potential real-world applications, many of which have not yet been developed practically in a sense that public datasets are oftentimes not available. When such a new field is to be developed, there are many practical challenges during dataset curation and creation. Oftentimes, data can be recorded with, e.g., cameras in large amount at acceptable cost, while acquiring corresponding labels might be comparatively costly and might require expert knowledge. Active learning aims at maintaining model performance while reducing the amount of training data by leveraging data informativeness for the label acquisition. The model is utilized in turn to find the data, in our case from a large pool of unlabeled

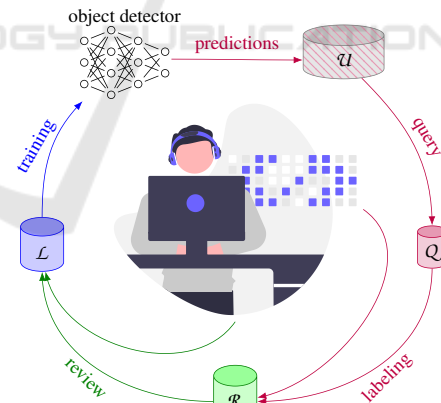






Figure 1: Our active learning cycle consists of **training** on labeled data \mathcal{L} , **querying and labeling** informative data points Q out of a pool of unlabeled data \mathcal{U} by an oracle and a **review** of acquired data $\mathcal{R} = \mathcal{L} \cup Q$.

beled data, for which new labels would improve the model performance most efficiently (Settles, 2009; Brust et al., 2018; Riedlinger et al., 2022). The goal is to request as few annotations with human labor as possible while still receiving a well-performing model that makes accurate predictions. When developing and simulating active learning models in a lab-

^a <https://orcid.org/0000-0002-9410-8949>

^b <https://orcid.org/0000-0002-1953-8607>

^c <https://orcid.org/0000-0002-3510-3320>

^d <https://orcid.org/0000-0003-3840-0184>

* Equal contribution.

oratory setup, one typically assumes an oracle that provides correct labels for queried data points (Brust et al., 2018; Riedlinger et al., 2022). However, in practice, such an oracle does not exist and any person that labels data produces errors with some frequency (Yan et al., 2014). Especially in complex domains such as medical applications where expert opinions are required for the annotation process, there exists variability between different oracles (Schilling et al., 2022). Some authors have considered active learning with noisy oracles in image classification before (Younesian et al., 2020; Younesian et al., 2021; Yan et al., 2016; Gupta et al., 2019). In the present work, we consider active learning with a noisy oracle, to the best of our knowledge for the first time, in object detection. More precisely, we utilize recent findings on label errors in state-of-the-art object detection benchmarks (Schubert et al., 2023) to simulate two types of predominantly occurring label errors in object detection oracles. On the one hand we treat missed bounding box labels which do not appear in the ground truth at all. On the other hand, we consider bounding box labels with incorrect class assignment which are likely to induce undesired training feedback. We do so for the EMNIST-Det dataset (Riedlinger et al., 2022) which is an extension of EMNIST (Cohen et al., 2017) to the object detection and instance segmentation setting. We complement this with the BDD100k dataset (Yu et al., 2020) which has mostly clean bounding box annotations of variable size. Both datasets have high quality bounding box labels such that we can simulate label errors without greater influence of naturally occurring label noise. We introduce independent and identically distributed errors into the labels which have been queried during the data-acquisition process. We simulate a label reviewer as a second human in the loop who has access to a label error detection module (Schubert et al., 2023) which is integrated into the active learning cycle, see fig. 1. We compare different sources for label error proposals which are to be considered after data acquisition. The efficiency of the proposal method controls the frequency of justified review cases, i.e., the efficiency of the budget utilization for label reviewing. The review oracle is assumed to contain smaller amounts of noise since labels do not have to be generated from scratch. Instead, only individual proposals have to be reviewed.

In our experiments, we observe that *a label error detection method applied to active learning with a noisy oracle clearly outperforms active learning with random label review and active learning without any label review*. We compare different query strategies with and without review in terms of performance as a

function of annotation budget (*split into labeling and reviewing cost*). Improvement of performance is observed consistently for random queries as well as for an uncertainty query based on the entropy of the object detector’s softmax output. Furthermore, our findings are consistent over two datasets, i.e., an artificial one and a real world one, as well as across two different object detectors. We find that the success of our method can likely be attributed to a strong performance of the label error detection method.

Our contribution can be summarized as follows:

- We contribute the first method that performs partially automated label review and active label selection for object detection.
- We provide an environment for performing rapid prototyping of methods for active learning with noisy oracles.
- Our method outperforms manual and review-free active learning for different queries, datasets and underlying object detectors.

Our method can be used with humans in the loop for labeling and label review to maximize model performance at minimal annotation budget, thus aiding data acquisition pipelines with partial automation.

2 RELATED WORK

Our contribution is located at the intersection of two disciplines which both aim at reducing the tiresome workload of repetitive image annotation by human workers. Active learning aims at reducing the overall amount of annotations given while the goal of label reviewing is to control or improve the quality of present annotations.

Label Error Detection in Object Detection. Previous work on the detection of label errors for object detectors tend to make use of a model which was trained on given, potentially error-prone data. (Hu et al., 2022) compare a softmax probability-based measure per prediction with the given annotations to obtain proposals for label errors. (Schubert et al., 2023) used an instance-wise loss computation to identify different types of label errors. (Koksal et al., 2020), in contrast, use a template matching scheme to find annotation errors in frame sequences for UAV detection.

Active Learning in Object Detection. Training in the context of more refined computer vision tasks such as object detection requires significant computational resources and the learning task itself comes with an elevated degree of complexity. All the more important is efficient handling when it comes to expensive data annotations which can be approached by

active learning. The following accounts for the pioneering successes that were accomplished in fully supervised active learning for deep object detection. (Brust et al., 2018) used margin sampling by aggregating probability margin scores over predicted bounding boxes in different ways where a class-weighting scheme addresses class imbalances. (Roy et al., 2018) similarly utilize softmax entropy and committee-based scoring making use of the different detection scales of object detection architectures. (Schmidt et al., 2020) compare different deep ensemble consensus-based selection strategies leveraging model uncertainty. (Choi et al., 2021) utilize Monte-Carlo dropout in conjunction with a Gaussian mixture model to estimate epistemic and aleatoric uncertainty, respectively. The utilized scoring function for querying takes both kinds of uncertainty into account. (Haussmann et al., 2020) compare methodically diverse scoring functions in the objectness entropy, mutual information, estimated learning gradient and confidence coupled with different diversity selection mechanisms.

Active Learning with Noisy Oracle in Classification. The intersection between active learning and training under label noise has been addressed in the context of classification tasks before. While (Kim, 2022) used an active query mechanism for cleaning up labels, the proposed training algorithm itself is not active. (Younesian et al., 2020) consider noisy binary and categorical oracles by assigning different label costs to both in an online, stream-based active learning setting. (Yan et al., 2016) treat the query complexity of noisy oracles with a reject option in a theoretical manner. (Gupta et al., 2019) consider batch-based active learning with noisy oracles under the introduction of the QActor framework by (Younesian et al., 2021) which has a label cleaning module in its active learning cycle is most related to our approach. One of the proposed quality models chooses examples to clean via the cross-entropy loss which are then re-labeled by the oracle.

3 ACTIVE LEARNING WITH NOISY ORACLE

In this section, we describe the task of active learning in object detection as well as the addition of a review module to the generic active learning cycle. While in the active learning setting, new labels are queried on the basis of an informativeness measure, the *presence of erroneous or misleading oracle responses can counteract the benefit of the informed data selection*. In order to account for new data containing incorrect

labels, we introduce a review module that generates proposals for label errors to review and to potentially correct.

3.1 Active Learning with Review in Object Detection

Most of the commonly used datasets in object detection, e.g., MS-COCO (Lin et al., 2014) and Pascal VOC (Everingham et al., 2010), are also the most commonly used datasets in active learning and contain label errors (Schubert et al., 2023; Rottmann and Reese, 2023). This means that active learning methods developed on these datasets are also evaluated based on noisy labels. To consider label errors during active learning experiments, we introduce a review module. Active learning can be viewed as an alternating process of training a model and labeling data based on informativeness according to the model. Starting with an initially labeled set of images \mathcal{L} , once a model is trained based on \mathcal{L} , the test performance is measured. Object detectors are usually evaluated in terms of mean average precision (mAP , see (Everingham et al., 2010)). New images selected to be labeled (Q) are queried from a pool \mathcal{U} of previously unlabeled images. After obtaining labels for the queried images Q by an oracle, we introduce a review step where an oracle reviews the combined set $\mathcal{R} = \mathcal{L} \cup Q$. The model is then trained on the reviewed data and the cycle is repeated T times. The active learning cycle is visualized in fig. 1 where acquisition and review of data are two independent modules.

Queries. Active learning research usually revolves around the development of model architectures, loss functions or selection strategies used in the query step. Different query approaches are then compared for different annotation budgets in terms of the achieved test performance which is often measured in terms of mAP in object detection. In the following, we investigate two different query strategies: random selection and selection of samples with large entropy of the softmax output of the object detector. For the former, images are randomly chosen from \mathcal{U} . For the latter, images are selected based on the predictive classification uncertainty according to the current model in the following sense. On a given image x , a neural network predicts a fixed number N_0 of bounding box predictions

$$\hat{b}^i = \{(x^i, y^i, \hat{w}^i, \hat{h}^i, \hat{s}^i, \hat{p}_1^i, \dots, \hat{p}_C^i)\}, \quad (1)$$

where $i = 1, \dots, N_0$. Here, $\hat{x}^i, \hat{y}^i, \hat{w}^i, \hat{h}^i$ represent the localization, \hat{s}^i the objectness score and $\hat{p}_1^i, \dots, \hat{p}_C^i$ the class probabilities for all classes $\{1, \dots, C\}$. Only the set of boxes that remain after score-thresholding (with

threshold s_ε) and non-maximum suppression (NMS) are used to compute prediction-wise entropies. The entropy $H(\hat{b}^i)$ of prediction \hat{b}^i is given by $H(\hat{b}^i) = -\sum_{c=1}^C \hat{p}_c^i \cdot \log(\hat{p}_c^i)$. Moreover, we incorporate a class-weighting (Brust et al., 2018) for computing instance-wise uncertainty scores. Finally, the instance-wise entropies of a single image are summed up to obtain an image-wise query score and the images are sorted in descending order by this query score. Note, that both queries are independent of labels and, therefore, of any label errors (random selection is independent of predictions and labels as well as the image-wise query score for the entropy only method is determined based on the predictions only). Note, that query algorithms are based on unlabeled data and oracle noise does not directly impact the selection of images. However, *oracle noise does influence model training*.

Review Module. In order to account for noisy oracles, we allow for incorrect annotations given in response to an active learning query. To counter-act the negative influence of noisy annotations, we introduce a review module into the active learning cycle in which proposals for label reviews are given and part of the annotation budget can be used to clean up a portion of annotation errors. In the following, we introduce the detection of two different kinds of label errors: missing labels (*misses*) and labels with incorrect class assignments (*flips*).

For one active learning cycle, we allow for the consumption of a fixed annotation budget \mathcal{C} . This budget \mathcal{C} is split up into a fraction $\mathcal{C}_Q = (1 - \lambda)\mathcal{C}$ used for querying new annotations and $\mathcal{C}_R = \lambda\mathcal{C}$ used for reviewing data.

After the query, Q is automatically labeled and, together with \mathcal{L} , forms the set of active images for the next cycle. Before the next training cycle starts, we regard $\mathcal{R} = \mathcal{L} \cup Q$ as the set of annotations which are potentially reviewed. Inspired by (Schubert et al., 2023), we introduce a post-processing label error detection method where the detection of *misses* and *flips* are two independent tasks. When both types of label errors are simultaneously present, we use a parameter $\alpha \in [0, 1]$ to distribute the review budget \mathcal{C}_R between reviewing *misses* ($\alpha\mathcal{C}_R$) and *flips* ($(1 - \alpha)\mathcal{C}_R$). In the following, we introduce two different review functions: a random review and a review (Schubert et al., 2023) based on the computed loss function.

An illustration of our label error detection method can be found in fig. 2. We consider the set of all predictions on images from \mathcal{R} and the corresponding noisy labels $\tilde{\mathcal{Y}}$. To detect *misses*, we select those predictions that are identified as false positive predictions according to the noisy ground truth $\tilde{\mathcal{Y}}$. To get an order for the review, we sort the false positives in descend-

ing order by the objectness score \hat{s} for the highest loss based review and in random order for the random review. Large values of \hat{s} on false positives (where objectness is supposed to be small) amounts to a large objectness loss.

For the *flips*, every label from $\tilde{\mathcal{Y}}$ is assigned to the most overlapping prediction if the *IoU* of the two boxes is greater than or equal to IoU_ε . Then, the cross-entropy loss of the possibly incorrect label and the predicted probability distribution is used as a review score for every given label. In case there is no sufficiently overlapping prediction, the respective label is not considered for review. Note, that for given label class \tilde{c} , the cross-entropy loss of the assigned prediction \hat{b} is $CE(\hat{b}|\tilde{c}) = -\sum_{c=1}^C \delta_{c\tilde{c}} \log(\hat{p}_c) = -\log(\hat{p}_{\tilde{c}})$, where δ_{ij} is the Kronecker symbol, i.e., $\delta_{ij} = 1$ if $i = j$ and $\delta_{ij} = 0$ otherwise. That is, the label with assigned prediction that has the lowest corresponding class probability $\hat{p}_{\tilde{c}}$ generates the highest loss. In case of the random review method, we randomly select assigned labels for review by uniformly sampling over all labels.

4 NUMERICAL EXPERIMENTS

In this section, we describe our active learning setup, label noise generation and the automated data review as well as involved hyperparameters. Afterwards, we show results for both query functions with and without review for two different datasets and object detectors in terms of *mAP*. We also measure the performance of the review proposal mechanism in terms of precision.

4.1 Experimental Setup

For our active learning setup, automated labeling and reviewing is desirable. Therefore, we simulate label errors for all training images of the underlying datasets. We do not include evaluations of the active learning experiments on the widely used MS-COCO or Pascal VOC datasets for the following reasons. For an automated review procedure, the frequently occurring label errors in both datasets (Schubert et al., 2023) would lead to strongly biased results. Evaluations on either dataset would require manual annotation review after each active learning cycle for several repetitions of the same experiment. This manual review after each cycle is necessary in practice, however, out of scope for an experimental evaluation of the proposed method.

Datasets and Models. We make use of the EMNIST-Det dataset (Riedlinger et al., 2022) with 20,000

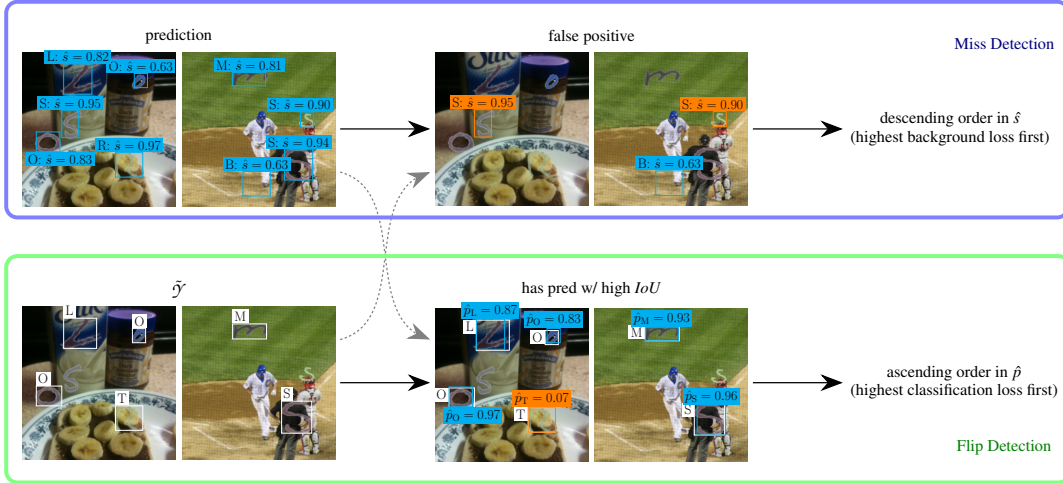


Figure 2: Schematic illustration of the label error detection mechanism using highest loss for missed labels (top, blue) and label flips (bottom, green). The detection of misses considers false positive predictions of the object detector while the detection of flips considers ground truth boxes, each of them matching at least one of the predictions in localization.

training images and 2,000 test images as well as BDD100k (Yu et al., 2020) (short: BDD), where we filter the training and validation split, such that we only use daytime images with clear weather conditions, resulting in 12,454 training images. Furthermore, the validation set is split into equally-sized test and validation sets, each consisting of 882 images. Since EMNIST-Det is a simpler task to learn compared to BDD, we apply a RetinaNet (Lin et al., 2017) and a Faster R-CNN (Ren et al., 2015) with a ResNet-18 (He et al., 2016) backbone for EMNIST-Det, as well as a Faster R-CNN with a ResNet-101 backbone for BDD. Note, that this setup was introduced and used in related work (Riedlinger et al., 2022; Schubert et al., 2023).

Based on clean training data, the test performance for EMNIST-Det in terms of mAP is 91.2% for Faster R-CNN and 90.9% for RetinaNet. For Faster R-CNN, the test performance decreases to 90.2% with simulated *misses* in the training data and to 89.2% with simulated *flips*. Simulating *misses* and *flips* simultaneously yields a test performance of 89.4% for Faster R-CNN and 89.3% for RetinaNet. For BDD and Faster R-CNN, a test performance of 50.0% is obtained for unmodified training data and 48.9% for training data including *misses* and *drops*.

Simulation of Label Errors. For the simulation of *misses* and *flips*, we follow (Schubert et al., 2023). An illustration of the label error injection scheme can be found in fig. 3. Any dataset is equipped with a set of G labels, i.e.,

$$\mathcal{Y} = \{b^i = (x^i, y^i, w^i, h^i, c^i) : i = 1, \dots, G\}. \quad (2)$$

Let $I = \{1, \dots, G\}$ be the set of indices of all boxes $b^i \in \mathcal{Y}$, $i = 1, \dots, G$. Furthermore, we choose a pa-

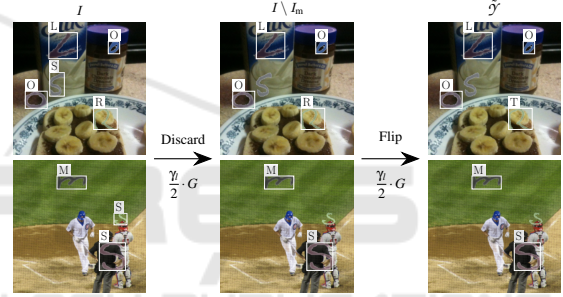


Figure 3: Schematic illustration of noise injection into a clean dataset. Misses are generated by randomly discarding $\gamma_l/2 \cdot G$ of the present annotations. Afterwards, $\gamma_l/2 \cdot G$ of the remaining annotations receive random class flips to one of the $C - 1$ incorrect classes. This results in the noisy ground truth $\tilde{\mathcal{Y}}$ used in our experiments.

parameter $\gamma_l \in [0, 1]$ representing the relative frequency of label errors during data acquisition.

For generating label misses, we randomly choose a subset $I_m \subset I$ of size $\frac{\gamma_l}{2} \cdot G$, representing missed ground truth boxes which are discarded from \mathcal{Y} . The number of remaining annotations that receive a class flip is again $\frac{\gamma_l}{2} \cdot G$, where the class flip is determined uniformly over the $C - 1$ incorrect class assignments. The potentially flipped class for each label b^i is denoted by \tilde{c}^i . Finally, we denote the training set including label errors by $\tilde{\mathcal{Y}} = \{(x^i, y^i, w^i, h^i, \tilde{c}^i) : i \in I \setminus I_m\}$. Note, that a single label is perturbed by only one type of label error at most. In addition, label errors are not simulated on the test dataset to ensure an unbiased evaluation of test performance.

Automated Review of Label Errors Since the oracle is noisy with error frequency γ_r , the review is also error-prone, i.e., *misses* are detected with proba-

Table 1: Overview of important training, review and active learning hyperparameters for all datasets and networks.

		Active Learning				Review				Training	
		$ \mathcal{U}_{init} $	C	s_ϵ	T	γ_l	γ_r	IoU_ϵ	α	batch size	training iters
Faster R-CNN	EMNIST-Det	150	200	0.7	20	0.2	0.05	0.3	0.5	4	25,000
RetinaNet	EMNIST-Det	150	200	0.5	20	0.2	0.05	0.3	0.5	4	38,000
Faster R-CNN	BDD100k	625	10,000	0.7	7	0.2	0.05	0.3	0.5	4	170,000




Figure 4: EMNIST-Det ablation study of the ratio between labeled and reviewed bounding boxes for Faster R-CNN where both label error types are present. We compare the random query without review with random query and highest loss review (RHL) with the chosen ratio λ in brackets.

bility $1 - \gamma_r$ and still overlooked with probability γ_r . The *flips*, whether the label error proposal was a false alarm or not, are corrected with probability $1 - \gamma_r$ and randomly misclassified with probability γ_r .

Implementation Details. We implemented our active learning methods in the open source MMDetection toolbox (Chen et al., 2019). For the label error simulation, we choose the relative frequency of label errors $\gamma_l = 0.2$, the relative frequency of label errors during review $\gamma_r = 0.05$, the value for score-thresholding $s_\epsilon = 0.7$ for Faster R-CNN and $s_\epsilon = 0.5$ for RetinaNet as well as the IoU -value $IoU_\epsilon = 0.3$ that assigns predictions with labels. We choose $\gamma_r < \gamma_l$, assuming that the oracle is more engaged in viewing and evaluating single boxes during the review compared to labeling from scratch, with all boxes having to be located and classified on a new image. As hyperparameters for the active learning cycle, the initially labeled set consists of 150 randomly picked images EMNIST-Det and of 625 randomly picked images for BDD. The budget for a single active learning step C is 200 for EMNIST-Det and $C = 10,000$ for BDD. Labeling a single box has cost 1, as does reviewing a label error proposal, whether *miss* or *flip* and also whether a label error was identified or not. If *misses* and *flips* are simultaneously present in the experiment, we set the ratio between reviewing *misses* and *flips* $\alpha = 0.5$. Finally, the number of active learning steps for EMNIST-Det is $T = 20$ and for BDD $T = 7$. For an overview of

training, review and active learning hyperparameters, see table 1.

4.2 Results

In the following, we show active learning results for EMNIST-Det and BDD. Therefore, we compare six different methods, the random and entropy query, both without review, as well with random review or review by highest loss. Furthermore, we present performance results for both review methods in terms of precision over the whole active learning course.

Ablation for the Ratio of Queried and Reviewed Bounding Boxes. For the methods with review, the fraction λ of the amount of new data queries and the amount of bounding box reviews plays a significant role. Therefore, we repeat the same experiment for EMNIST-Det and Faster R-CNN with different values for λ , see fig. 4. These results are based on training data with simulated *flips* and *misses*. The gray and yellow lines indicate the 100% and 90% reference performance mark of the model trained with the entire (noisy) dataset. The random query with highest loss review (RHL) is visually almost identical for $\lambda = 0.1$, $\lambda = 0.16$ and $\lambda = 0.2$. All these three methods outperform the random query without review. The random query without review performs similar to the random query with highest loss review with $\lambda = 0.3$ and outperforms the random query with $\lambda = 0.4$, i.e., at about $\lambda = 0.3$ is the break-even-point, at which it is no longer worthwhile to review more bounding boxes instead of labeling new ones. We hypothesize that this tipping point is strongly dependent on our chosen setup with a relative frequency of label errors of $\gamma_l = 0.2$. Since the red curve seems to be most favorable, we set the fraction between queries and reviews to $\lambda = 0.2$ in all of the following experiments.

Active Learning with Different Label Error Types.

We first investigate active learning curves for EMNIST-Det and Faster R-CNN. We consider active learning curves where (a) we simulated only *misses*, (b) only *flips*, and (c) both label error types occur equally often in the training dataset, each with noise rate $\gamma_l = 0.1$ (recall fig. 3). We compare both query strategies, random and entropy, without review, with highest loss (HL) review and with random (R) review,




Figure 5: EMNIST-Det active learning curves, where only flips are present (left), only misses (center) as well as where flips and misses are simultaneously present (right).

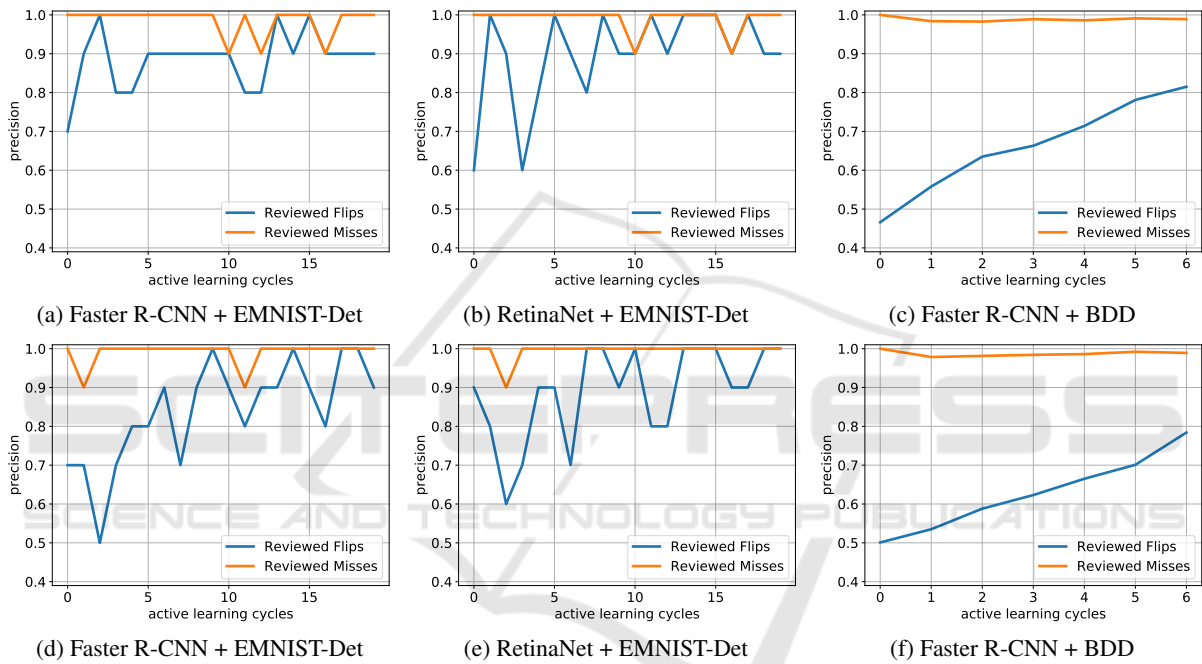


Figure 6: Review quality results for random highest loss (top) and entropy highest loss (bottom). Misses and flips are simultaneously present in all experiments.

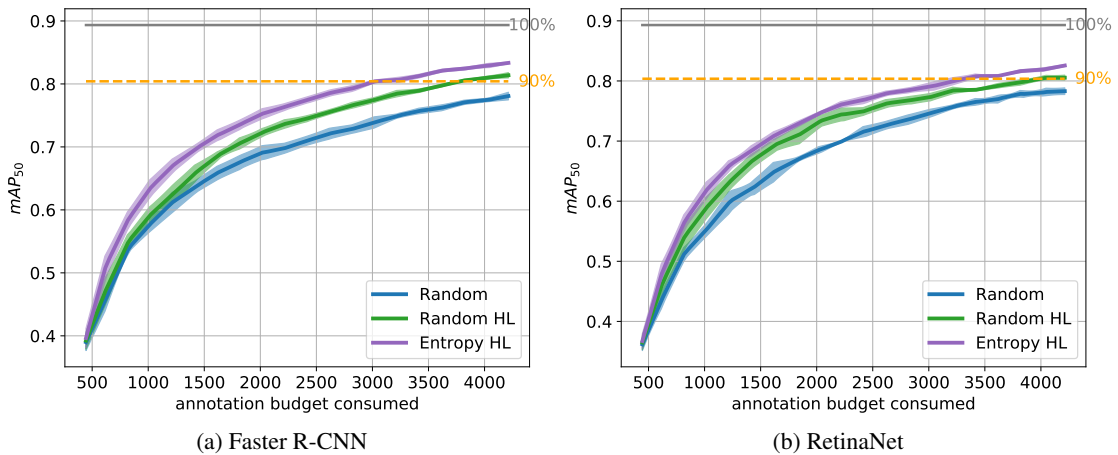


Figure 7: Comparison of EMNIST-Det active learning curves for Faster R-CNN (left) and RetinaNet (right) where both label error types are present.

respectively. The obtained active learning curves are averaged over four random initializations and evaluated in terms of the total annotation budget consumed. Note, that for the active learning methods without review the total amount of annotation budget is equal to the amount of (possibly incorrect) labeled bounding boxes. For those methods incorporating a review step the total amount of annotation budget represents the sum of labeled and reviewed bounding boxes. Figure 5 shows active learning curves in terms of test performance with point-wise standard deviations. For all three active learning curves, we observe that the entropy method outperforms random at every point. Furthermore, the queries without review perform superior to the respective query with random review. Entropy HL and random HL, i.e., both methods with highest loss review clearly outperform the strategies without review and with random review. We conclude that the success of reviewing queries strongly depends on the performance of the review methods and that random review is too expensive in terms of annotation budget. From this we conclude that it is more worthwhile to acquire new (noisy) labels than to randomly review the active labels, at least for the given amount of noise we studied.

All in all, the distance between the active learning curves of the six methods is significantly larger for label *flips* as compared to *misses*. Moreover, the maximum performance with simulated label errors is also inferior for the *flips* compared to the *misses*. We hypothesize that the reason for this is the sub-sampling from the negatively associated anchors (Ren et al., 2015) during training. This mechanism leads to only partial learning from the *misses*, whereas an incorrect foreground class induced from *flips* has a negative impact on every gradient step.

The significant difference of the active learning curves of the respective queries with random review and the highest loss review can be attributed to the high precision of the highest loss review. The random review has an expected precision of γ_l . Figure 6 shows the precision for the highest loss review applied after random query in (a) and after entropy query in (d) across the span of all active learning cycles. In both plots, (a) and (d), *flips* and *misses* are simultaneously present, i.e., both plots correspond to the respective method from fig. 5 (c). The blue lines visualize the precision for the review identifying a *flip* and the orange lines analogously for the *misses*. Here, the precision for detecting *flips* is always above 50% and tends to improve as the active learning experiment progresses, whereas the precision for the detection of *misses* is even consistently above 90%. In general, *flips* are more difficult to detect compared to

Table 2: Mean average precision values in % with standard deviations in brackets for 2000 and 4000 queried and reviewed annotations for Faster R-CNN (top) and RetinaNet (bottom) for EMNIST-Det. Note, that in every experiment both label error types are present; the upper half represents fig. 7 (a) and the bottom half fig. 7 (b).

Network	Method	$mAP@2000$	$mAP@4000$
Faster R-CNN	Random	68.97(± 1.09)	77.41(± 0.28)
	Random HL	72.13(± 0.64)	80.92(± 0.23)
	Entropy HL	75.13(± 0.79)	82.87(± 0.29)
RetinaNet	Random	68.38(± 0.52)	78.04(± 0.51)
	Random HL	72.94(± 1.06)	80.40(± 0.34)
	Entropy HL	74.39(± 0.28)	81.86(± 0.27)

misses due to the different construction of the detection methods of either label error type, recall fig. 2.

Comparing Faster R-CNN with RetinaNet on EMNIST-Det. For the following results, we compare only the random query without review with the random and entropy query, both with highest loss review. Figure 7 shows active learning curves for these methods for Faster R-CNN in (a) and for RetinaNet in (b). Note, that in both cases both label error types are present. Moreover, (a) is a trimmed version of fig. 5 (c) to make it more convenient to compare the results from both detectors visually. We observe that the curves for Faster R-CNN and RetinaNet look very similar over the entire active learning course. All curves start at just below 40% mAP and the respective methods end at similar test performances. The ranking of the methods is always the same: entropy HL outperforms random HL and random without review. Also, random HL outperforms random without review. For RetinaNet, random HL seems to be closer to entropy HL as compared to Faster R-CNN.

These observations are also supported by table 2, wherein we stated the mAP values with standard deviations in parentheses. We compare performance for the total annotation budget consumed equal to 2000 and 4000 from the active learning curves shown in fig. 7. In particular, for entropy with highest loss review, the $mAP@2000$ for Faster R-CNN is 0.74 percent points (pp) higher and even 1.01 $mAP@4000$ pp higher. For Faster R-CNN, the difference between entropy HL and random HL is 3 pp for $mAP@2000$ and 1.95 for $mAP@4000$. For RetinaNet, the difference is only 1.45 pp for $mAP@2000$ and 1.46 for $mAP@4000$.

Comparing the quality of the highest loss review, the results for RetinaNet are highly correlated to the results of Faster R-CNN, see fig. 6. For RetinaNet, the precision for the highest loss review in combination with random query is visualized in (b) and with the entropy query in (e). Here, the precision for detecting the *misses* is at or above 90%. The precision for the detection of *flips* is always greater or equal to 60% and from active learning cycle 7 onward even

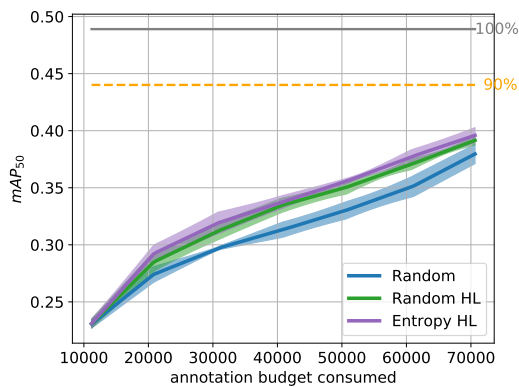


Figure 8: BDD active learning curves for Faster R-CNN where both label error types are present.

always above 80%. We observe that the precision of the highest loss review increases while the experiments progress, i.e., object detectors trained with more data generate better label error proposals. We hypothesize that with more data, overfitting can be more effectively prevented and that the object detectors will generalize better, thus label errors in the active labels will not be as significant when sufficient data is available.

Results for BDD with Faster R-CNN. Figure 8 shows active learning curves for the random query without review, as well as for the random and entropy query both with highest loss review. Comparable to the results for EMNIST-Det, the ranking of the methods is identical over the entire active learning course. The random query without review is inferior to both queries with review and entropy HL is superior to random HL. Note, that the distance between the two queries with review is marginal. In contrast, there is a significant difference between either one and the random query without review.

The review quality of the highest loss review for the random query is shown in fig. 6 (c) and for the entropy query in (f). Again, the *misses* are detected at all times with a precision of nearly 100%. Starting at just under 50%, the precision for identifying *flips* increases steadily over the active learning course ending at close to 80%. We conclude that involving a label review in the active learning cycle is also highly beneficial in the more complex BDD real world dataset. Analogous to the results for EMNIST-Det, the highest loss review becomes more precise as the experiments progress and the number of active labels increases.

5 CONCLUSION

In this work, we considered label errors in active learning cycles for object detection for the first time,

where we assumed a noisy oracle during the annotation process. We realized this assumption by simulating two types of label errors for the training data of datasets which are reasonably free of intrinsic label errors. These types of label errors are missing bounding box labels as well as bounding box labels with an incorrect class assignment. We introduce a review module to the active learning cycle, that takes as input the currently labeled images and the corresponding predictions of the most recently trained object detector. Furthermore, we detect both types of label errors by a random review method and a method based on the highest loss of the model’s predictions and the corresponding noisy labels. We observe that the incorporation of random review leads to an even worse test performance compared to the corresponding query without review. Nevertheless, we show that the combination of query strategies, like random selection or instance-wise entropy, with an accurate review yields a significant performance increase. For both query strategies, the improvement obtained by including the highest loss review persists during the whole active learning course for different dataset-network-combinations. We make our code for reproducing results and further development publicly available at https://github.com/mschubert95/active_learning_with_label_errors.

ACKNOWLEDGEMENTS

We gratefully acknowledge financial support by the German Federal Ministry of Education and Research in the scope of project “UnREAL”, grant no. 01IS22069. The authors gratefully acknowledge the Gauss Centre for Supercomputing e.V. for funding this project by providing computing time through the John von Neumann Institute for Computing on the GCS Supercomputer JUWELS at Julich Supercomputing Centre.

REFERENCES

- Brust, C.-A., Käding, C., and Denzler, J. (2018). Active learning for deep object detection. *arXiv preprint arXiv:1809.09875*.
- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., and Zagoruyko, S. (2020). End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer.
- Chen, K., Wang, J., Pang, J., Cao, Y., Xiong, Y., Li, X., Sun, S., Feng, W., Liu, Z., Xu, J., Zhang, Z., Cheng, D., Zhu, C., Cheng, T., Zhao, Q., Li, B., Lu, X., Zhu, R., Wu, Y., Dai, J., Wang, J., Shi, J., Ouyang, W., Loy,

- C. C., and Lin, D. (2019). MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*.
- Choi, J., Elezi, I., Lee, H.-J., Farabet, C., and Alvarez, J. M. (2021). Active Learning for Deep Object Detection via Probabilistic Modeling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10244–10253, Montreal, QC, Canada. IEEE.
- Cohen, G., Afshar, S., Tapson, J., and Van Schaik, A. (2017). Emnist: Extending mnist to handwritten letters. In *2017 international joint conference on neural networks (IJCNN)*, pages 2921–2926. IEEE.
- Everingham, M., Van Gool, L., Williams, C. K., Winn, J., and Zisserman, A. (2010). The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88:303–338.
- Gupta, G., Sahu, A. K., and Lin, W.-Y. (2019). Noisy Batch Active Learning with Deterministic Annealing. *arXiv preprint arXiv:1909.12473*.
- Haussmann, E., Fenzi, M., Chitta, K., Ivanecky, J., Xu, H., Roy, D., Mittel, A., Koumchatzky, N., Farabet, C., and Alvarez, J. M. (2020). Scalable Active Learning for Object Detection. In *2020 IEEE Intelligent Vehicles Symposium (IV)*, pages 1430–1435.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Hu, Z., Gao, K., Zhang, X., Wang, J., Wang, H., and Han, J. (2022). Probability differential-based class label noise purification for object detection in aerial images. *IEEE Geoscience and Remote Sensing Letters*, 19:1–5.
- Kim, K. I. (2022). Active Label Correction Using Robust Parameter Update and Entropy Propagation. In *European Conference on Computer Vision*, pages 1–16. Springer.
- Koksal, A., Ince, K. G., and Alatan, A. (2020). Effect of annotation errors on drone detection with YOLOv3. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 1030–1031.
- Lin, T.-Y., Goyal, P., Girshick, R., He, K., and Dollár, P. (2017). Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988.
- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. (2014). Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer.
- Redmon, J. and Farhadi, A. (2018). Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*.
- Ren, S., He, K., Girshick, R., and Sun, J. (2015). Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28.
- Riedlinger, T., Schubert, M., Kahl, K., Gottschalk, H., and Rottmann, M. (2022). Towards Rapid Prototyping and Comparability in Active Learning for Deep Object Detection. *arXiv preprint arXiv:2212.10836*.
- Rottmann, M. and Reese, M. (2023). Automated detection of label errors in semantic segmentation datasets via deep learning and uncertainty quantification. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3214–3223.
- Roy, S., Unmesh, A., and Nambodiri, V. P. (2018). Deep active learning for object detection. In *BMVC*, volume 362, page 91.
- Schilling, M. P., Scherr, T., Munke, F. R., Neumann, O., Schutera, M., Mikut, R., and Reischl, M. (2022). Automated Annotator Variability Inspection for Biomedical Image Segmentation. *IEEE access*, 10:2753.
- Schmidt, S., Rao, Q., Tatsch, J., and Knoll, A. (2020). Advanced active learning strategies for object detection. In *2020 IEEE Intelligent Vehicles Symposium (IV)*, pages 871–876. IEEE.
- Schubert, M., Riedlinger, T., Kahl, K., Kröll, D., Schoenen, S., Šegvić, S., and Rottmann, M. (2023). Identifying label errors in object detection datasets by loss inspection. *arXiv preprint arXiv:2303.06999*.
- Settles, B. (2009). Active learning literature survey.
- Yan, S., Chaudhuri, K., and Javidi, T. (2016). Active learning from imperfect labelers. *Advances in Neural Information Processing Systems*, 29.
- Yan, Y., Rosales, R., Fung, G., Subramanian, R., and Dy, J. (2014). Learning from multiple annotators with varying expertise. *Machine Learning*, 95(3):291–327.
- Younesian, T., Epema, D., and Chen, L. Y. (2020). Active learning for noisy data streams using weak and strong labelers. *arXiv preprint arXiv:2010.14149*.
- Younesian, T., Zhao, Z., Ghiassi, A., Birke, R., and Chen, L. Y. (2021). Qactor: Active learning on noisy labels. In *Asian Conference on Machine Learning*, pages 548–563. PMLR.
- Yu, F., Chen, H., Wang, X., Xian, W., Chen, Y., Liu, F., Madhavan, V., and Darrell, T. (2020). Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2636–2645.