Object Detection on Streaming LiDAR Data with Active Learning

Casey Baker, Brian Daniel, Elliot Immel, Christopher Bogart

Expedition Technology, Inc

Keywords: Active Learning, LiDAR, Object Detection

ABSTRACT

Geospatial intelligence analysts rely on automated object detection methods to make efficient and informed decisions in dynamic environments. Deep learning methods achieve state-of-the-art accuracy and speed for object detection but require large amounts of manually labeled data for training. Active learning approaches aim to reduce this manual effort by nominating the most informative subset of the data for human labeling. Active learning on streaming data occurs when a deployed inference model makes on-the-fly decisions of whether to nominate a sample to be labeled by analysts. These labeling determinations are typically made using uncertainty-based metrics, which prioritize labeling samples where the model's predictive confidence is below a set threshold. However, uncertainty-based methods tend to select highly similar samples with redundant information, which may not improve model performance. Alternatively, diversity-based methods prioritize samples with non-redundant information, but these methods have generally been limited to settings with 2-dimensional non-streaming data. In this work, we develop diversity-based active learning methods for object detection on streaming LiDAR point clouds. We compare the accuracy and labeling efficiency of our methods to standard uncertainty-based methods and find that our novel methods achieve superior accuracy with fewer labeled samples. Our findings can be applied to diverse fields of computer vision, including remote sensing, to improve the efficiency of their data processing pipelines.

1 INTRODUCTION

Geospatial intelligence analysts rely on continuous information from airborne sensors to make decisions that impact national security. Specifically, these analysts observe the dynamic distributions and locations of target objects in an environment to evaluate potential threats. Machine learning can help analysts quickly survey new data by automatically detecting these targets. These Automatic Target Recognition (ATR) algorithms need to be fast and accurate to provide useful real-time decision support. These criteria are met by deep learning methods, which achieve state-of-the-art speed and accuracy for object recognition tasks across a variety of imaging domains [1-3]. However, deep learning is very resource-intensive, requiring large amounts of labeled ground-truth data to train models. Although large, labeled datasets exist for research purposes as standardized benchmarks, models trained on these datasets do not generalize to real-world tasks, which have different imaging parameters, ontologies, and environments. Therefore, analysts must train models with their own manually annotated data to achieve maximum accuracy.

The process of labeling data is tedious, time consuming, and prone to human error and bias. These drawbacks are especially true for object detection tasks, where annotating data is more laborious than for image classification tasks [4]; ground truth labels for object detection tasks require localized annotations for each object in a scene whereas classification tasks require just one non-localized label per scene. Labeling objects in 3D scenes adds further complexity and labor. Additionally, in real-world scenarios, the labeling effort would be ongoing: continuous data collection in dynamic conditions requires models to be updated routinely to maintain accuracy over time. Therefore, analysts need methods to distinguish which incoming data should be manually labeled for training. Our work aims to mitigate these labeling challenges by reducing the number of ground truth labels required to train and update 3D object detection models.

Active learning is a popular technique to reduce the amount of ground truth data necessary for training by iteratively selecting a small subset of unlabeled data that would maximally improve the model's performance [5]. Each iteration contains three steps: 1) the machine learning model produces a score for each unlabeled sample that represents the sample's potential to improve the model, 2) subset of samples is selected based on these scores and sent to a human analyst to label, 3) the labeled samples are added to the training dataset and the model is updated. These scores can be based on either model uncertainty, sample diversity, or a hybrid of the two. Uncertainty-based methods choose samples in which the model exhibits low confidence in its predictions [6]. As a result, the selected samples are close to the model's decision boundary but are usually similar to each other. Similar samples would likely contain redundant features, making it inefficient for reducing labeling efforts [7-9]. Conversely, diversity-based methods choose samples that reflect the dataset's heterogeneity by selecting varied samples [10,11]. However, the model may already confidently and accurately process them, so including them in training may not contribute to maximal performance gains. Hybrid methods moderate these drawbacks by choosing a diverse subset of the most uncertain samples [12].

The majority of existing research on active learning was conducted on 2D data [13,14]; however, some initial studies on 3D data have shown that active learning can reduce the amount of training data by 50-60% on object detection tasks [12,15,16]. Most of this research, however, assumed all unlabeled data was simultaneously available, failing to account for more realistic scenarios with ongoing data collection and a continuous incoming data stream. Methods that can accommodate data streams are especially important for DoD use-cases, where data acquisition and analysis are often done in tandem. Active learning on streaming data would require a deployed inference model to decide on-the-fly whether a newly collected data sample would be useful to bolster model performance (Figure 1). In the streaming setting, these labeling determinations are typically made using uncertainty-based metrics because diversity metrics generally begin with a clustering algorithm to divide the unlabeled data into disparate groups, which requires many unlabeled data samples to be simultaneously available [17,18]. Streaming active learning methods would benefit from novel diversity metrics that can be computed on individual samples.

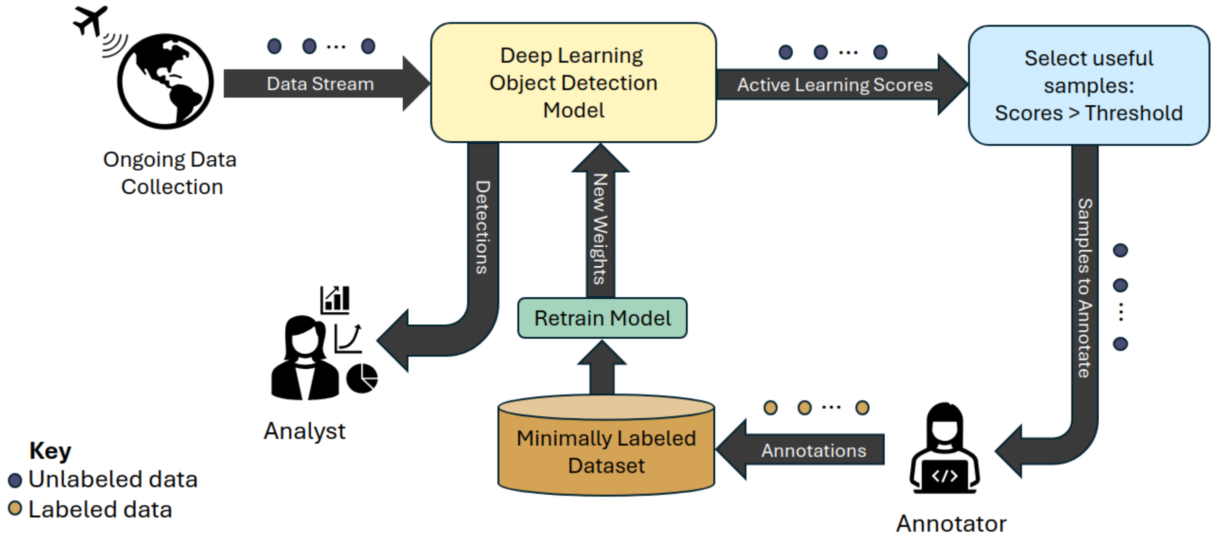

Figure 1: Streaming Active Learning Pipeline.

In this work, we develop a novel active learning metric, Normalized Object Distribution Entropy (NODE), that can be implemented on data streams. In contrast to most existing diversity metrics, which use inter-sample diversity, NODE uses intra-sample diversity to accommodate streaming data. Compared to leading uncertainty-based metrics, NODE reduces annotation costs while maintaining performance. NODE also prioritizes samples with more rare objects in imbalanced data sets. Our work is some of the first to implement diversity-based active learning in a streaming setting for 3D object detection. Our findings can be applied to many fields that need automated 2D and 3D object detection methods such as geospatial intelligence, robotics, and autonomous driving.

2 RELATED WORK

2.1 Pool-Based Active Learning

Pool-based active learning is a technique that applies when a user has a large, unlabeled dataset readily available (i.e., a previously collected dataset). Therefore, samples can be compared against each other to select the most informative subset during each iteration of active learning. This is contrary to stream-based active learning, which makes labeling decisions individually and on-the-fly. When using uncertainty metrics, pool-based active learning methods can calculate the uncertainty score of all unlabeled samples, rank these scores, and choose to label the most uncertain subset. When using diversity metrics, pool-based methods can apply a clustering algorithm to sort all unlabeled samples into disparate groups to ensure selected samples represent the dataset's heterogeneity.

Uncertainty metrics for active learning aim to select the samples where the model is least certain about its prediction. This certainty is typically estimated using the values of the final softmax layer of the model. In the case of object detection, this softmax layer represents the predicted probability that a detected object belongs to each class. The most straightforward approach to measuring uncertainty would be to take the highest value from the softmax layer, which designates the most probable class, and treat this value as the model's predictive certainty [19,20]. During each active learning iteration, the samples with the lowest certainty would be selected for labeling.

To leverage more available information, other uncertainty methods utilize the predicted probability values for the other classes, not just the most probable one. A common uncertainty metric, margin uncertainty, calculates uncertainty as the difference between the highest and second highest predicted probabilities in the softmax layer [21, 22]. Large differences between the highest and second highest probability reflect higher model certainty. Smaller differences indicate that the model has trouble distinguishing the two classes. Another common uncertainty metric is to compute the Shannon Entropy of the softmax layer [19, 21-24]. Shannon Entropy is highest when computed on a uniform probability distribution and lower on non-uniform distributions. Therefore, if the maximum softmax value is high, the model is confident in its prediction, and the Shannon Entropy will be low. Prior work on pool-based active learning for LiDAR point clouds showed that selecting samples Shannon Entropy resulted in the most accurate model [25].

Diversity-based methods project each sample into latent space so that distances between samples can be efficiently measured [12]. These methods aim to select samples with diverse features. One common diversity method utilizes K-medoids clustering, which iteratively groups the samples into distinctive clusters [26, 27]. The samples located closest to the center of each cluster are then selected for labeling. An alternative diversity approach is called Core-Set [28]. This method chooses the samples that minimizes the distance between unselected samples and their closest selected sample. While these methods only consider the unlabeled samples during the selection process, an approach called Furthest Nearest Neighbors also leverages the labeled samples [29]. This approach chooses the unlabeled samples that are the furthest distance from the labeled samples in latent space.

2.2 Stream-Based Active Learning

As previously stated, stream-based active learning occurs when data collection and analysis must be done in parallel and labeling decisions need to be made on-the-fly. In other words, unlabeled samples cannot be compared to each other when making labeling decisions; samples that are not selected for labeling are discarded. This is necessary for instances with data storage and time constraints [17]. Additionally, continuously filtering samples during ongoing data collection helps maintain the accuracy of deployed models in the presence of concept drift [18, 30, 31]. Concept drift occurs when the data distribution changes over time, which diminishes the performance of models trained on older data.

Stream-based active learning is a newer area of research compared to pool-based active learning. Consequently, there is less previous work for this branch of active learning, especially for object detection tasks. The work that has been done has focused largely on adapting existing uncertainty methods for a streaming scenario [17]. In general, the uncertainty metrics used in the streaming case are the same as those used in pool-based active learning (e.g. Shannon Entropy, margin, etc.). However, instead of ranking the metrics across all unlabeled samples and choosing the most uncertain, stream-based methods set an uncertainty threshold [18, 32]. Any sample whose uncertainty score exceeds the threshold is selected for labeling. This threshold can be a fixed value or adjustable based on model performance.

There is little existing research on diversity metrics for stream-based active learning. However, recent work developed a diversity metric for stream-based active learning on 3D object segmentation [22]. This metric, called Voxel Confusion Degree, divides a single sample into smaller samples and then selects the sub-sample with the highest predicted class diversity. Additionally, a diversity metric for stream-based active learning was recently developed for 2D image classification tasks [33]. This method projects all labeled data into latent space and computes the mean and standard deviation of the distance between the samples in the labeled dataset. This mean and standard deviation are used to compute the Z-score for each unlabeled sample. If the Z-score surpasses a set threshold, the sample is selected for labeling. To our knowledge, our work is the first to study active learning for object detection in 3D streaming data.

3 EXPERIMENT

3.1 Dataset

We used the same dataset in this work as was used in our prior work on pool-based active learning methods [12, 25]. This dataset consisted of 41 LiDAR point clouds with 7863 target objects. These objects are distributed across seven classes, with class 1 comprising ~50% of the data. We randomly divided the point clouds into training, validation, and test sets, which maintain the approximate class distribution of the entire dataset (Figure 2, Table 1). We preprocessed the data by dividing the point clouds into 30m x 30m overlapping tiles. Finally, we randomly selected 1000 tiles from the training set to serve as our initial labeled training data and randomly ordered the remaining 6589 training tiles to simulate an unlabeled data stream.

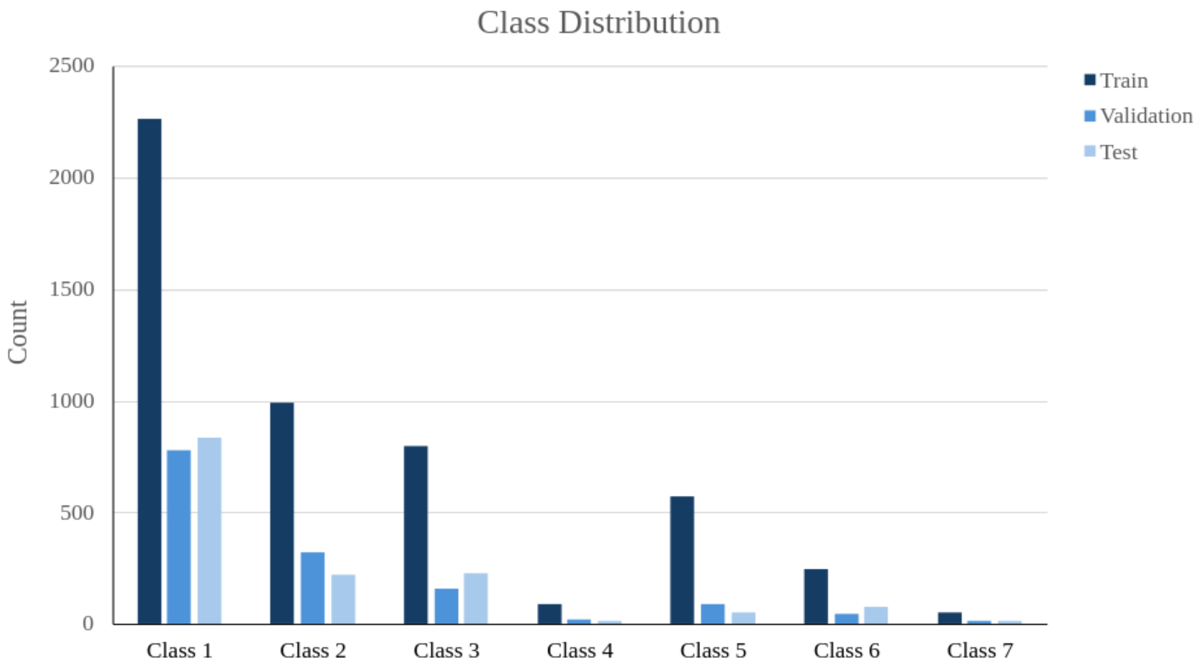

Figure 2: Class distribution for training, validation, and test sets

Table 1: Dataset Class Distribution

Split	Class 1	Class 2	Class 3	Class 4	Class 5	Class 6	Class 7	Total
Train	2261	993	796	87	570	242	53	5002
Validation	776	323	159	22	87	44	15	1426
Test	837	218	226	16	49	73	16	1435
Total	3874	1534	1181	125	706	359	84	7863

3.2 Model Architecture

We used a modified VoxelNet architecture as our object detector because of its robust prior performance on point clouds [34]. Our modifications to VoxelNet are described in detail in our previous work [12] but can be summarized by three main changes: 1) we removed a concatenation layer from the feature extractor portion of the model to reduce computational load, 2) we used Focal Loss and weighted sampling to adapt VoxelNet to imbalanced data sets, and 3) we used stochastic weight averaging and weight regularization to reduce overfitting. Prior to training, we augmented and divided each tile into 0.75m x 0.75m x 1m voxels. Augmentations included rotating or cropping point clouds and randomly dropping voxels. We used the Adam optimizer during training and trained the model from scratch for 150 epochs during each active learning iteration. During post-processing, we used non-maximum suppression with an Intersection-over-Union (IoU) threshold of 0.3 to remove duplicate detections.

3.3 Active Learning Pipeline

We first trained an initial model on 1000 labeled training tiles. We then simulated a streaming active learning paradigm. We presented unlabeled tiles to the initial model one at a time and calculated an active learning score for each tile. If the score exceeded a predetermined threshold, we added that tile to a cache to be labeled. If the score did not exceed the threshold, the tile was discarded. When 275 tiles had been added to the cache, we added these tiles to the training data and retrained the model from scratch. We repeated this process four more times for a total of five active learning iterations.

We used the same set of labeled tiles as our starting training data for every trial. During training, we randomly disregarded 50% of the background-only tiles (e.g. tiles without any objects) to reduce the likelihood of false negative detections. The active learning score thresholds were hyperparameters that we tuned to optimize performance on the validation set. For analysis, we ran eight trials of each active learning metric.

3.4 Active Learning Metrics

We developed a novel diversity-based active learning metric, Normalized Object Distribution Entropy (NODE), and compared its performance to a common uncertainty-based method, entropy, as well as a control of random selection.

Random: For our random control, we allotted a labeling budget of 25% to match the approximate proportion of selected tiles from the entropy and NODE methods. For each unlabeled tile in the data stream, we generated a random number between 0 and 1. If a tile's random number was less than 0.25, we added the tile to the cache. This method equates to selecting unlabeled tiles with a 25% probability.

Uncertainty - Entropy: Our uncertainty-based method, Entropy, is defined using the Shannon Entropy equation:

$$H(x) = -\sum p(x)\log\big(p(x)\big) \qquad (1)$$

In this case, $p(x)$ is the softmax vector for a single detected object in a tile. This vector has dimension 1 x (N+1), where N is the number of classes in a dataset, and the extra value is to accommodate background detections. Entropy scores are highest when computed on a uniform distribution, but their peak value varies with N. To generalize our method to different N values, we normalize object entropy scores by the entropy of the uniform distribution of length N. To calculate a tile-wide entropy score, we averaged the entropy scores for each detected object in the tile. We selected tiles that had an entropy score above 0.25 for labeling. This threshold was selected because it maximized performance on the validation set.

NODE: Our novel diversity-based method, NODE, also uses the Shannon Entropy equation (1). However, instead of using the softmax vector for $p(x)$, we use the distribution of detected objects in each tile. In this case, $p(x)$ has dimension 1 x N, and each value of the vector represents the proportion of objects belonging to a particular class. This method priorities tiles that have diverse objects over tiles with multiple objects of the same class. We first normalize the entropy value by the N-dimensional uniform distribution entropy value. Then, we normalize this value by the number of detected objects in a tile to encourage the model to select tiles with diverse, but few, objects to reduce the manual labeling effort. We selected tiles that had a NODE score above 0.025 for labeling. This threshold was selected because it maximized performance on the validation set.

3.5 Evaluation Metrics

The most widely used and accepted object detection metric is Average Precision (AP). AP is a recall-weighted average of precision. The definition of precision is:

$$\text{Precision} = \frac{True\ Positive}{True\ Positive + False\ Positive} \qquad (2)$$

Recall, also known as probability of detection, is defined as:

$$\text{Recall} = \frac{True\ Positive}{True\ Positive + False\ Negative} \qquad (3)$$

Based on (2) and (3), AP is defined as:

$$\text{AP} = \int_0^1 precision(recall)d(recall) \qquad (4)$$

AP is a great metric for single-class detection scenarios, but with multiple classes, there are two approaches to summarize performance, class-agnostic and class-specific. Class-agnostic AP, often referred to as mean Average Precision (mAP), is defined as:

$$\text{mAP} = \frac{1}{N}\sum_{i=1}^{N} AP_i \qquad (5)$$

Where N is the number of classes, and each class is equally weighted. Conversely, class-specific AP (csAP) is a weighted average of AP where the weights are the number of objects in each class. If we define h_i as the number of annotations in class i, then class-specific AP is defined as:

$$\text{csAP} = \frac{1}{N}\sum_{i=1}^{N} h_i\, AP_i \qquad (6)$$

We evaluate our object detector by AP with Intersection over Union (IoU) threshold of 0.5, known in the literature as AP50. The mAP and csAP metrics are used to evaluate the Active Learning methods in this work.

4 RESULTS

The results of our experiment described above are shown plotted with class-specific and class agnostic metrics versus number annotated tiles (left) and number of objects (right) in Figure 3. Error bars are standard error measurements. The first observation to note is the random selection of tiles for each active learning training run produces poor results compared to NODE and Uncertainty. Comparison of our method to Uncertainty is a little more nuanced.

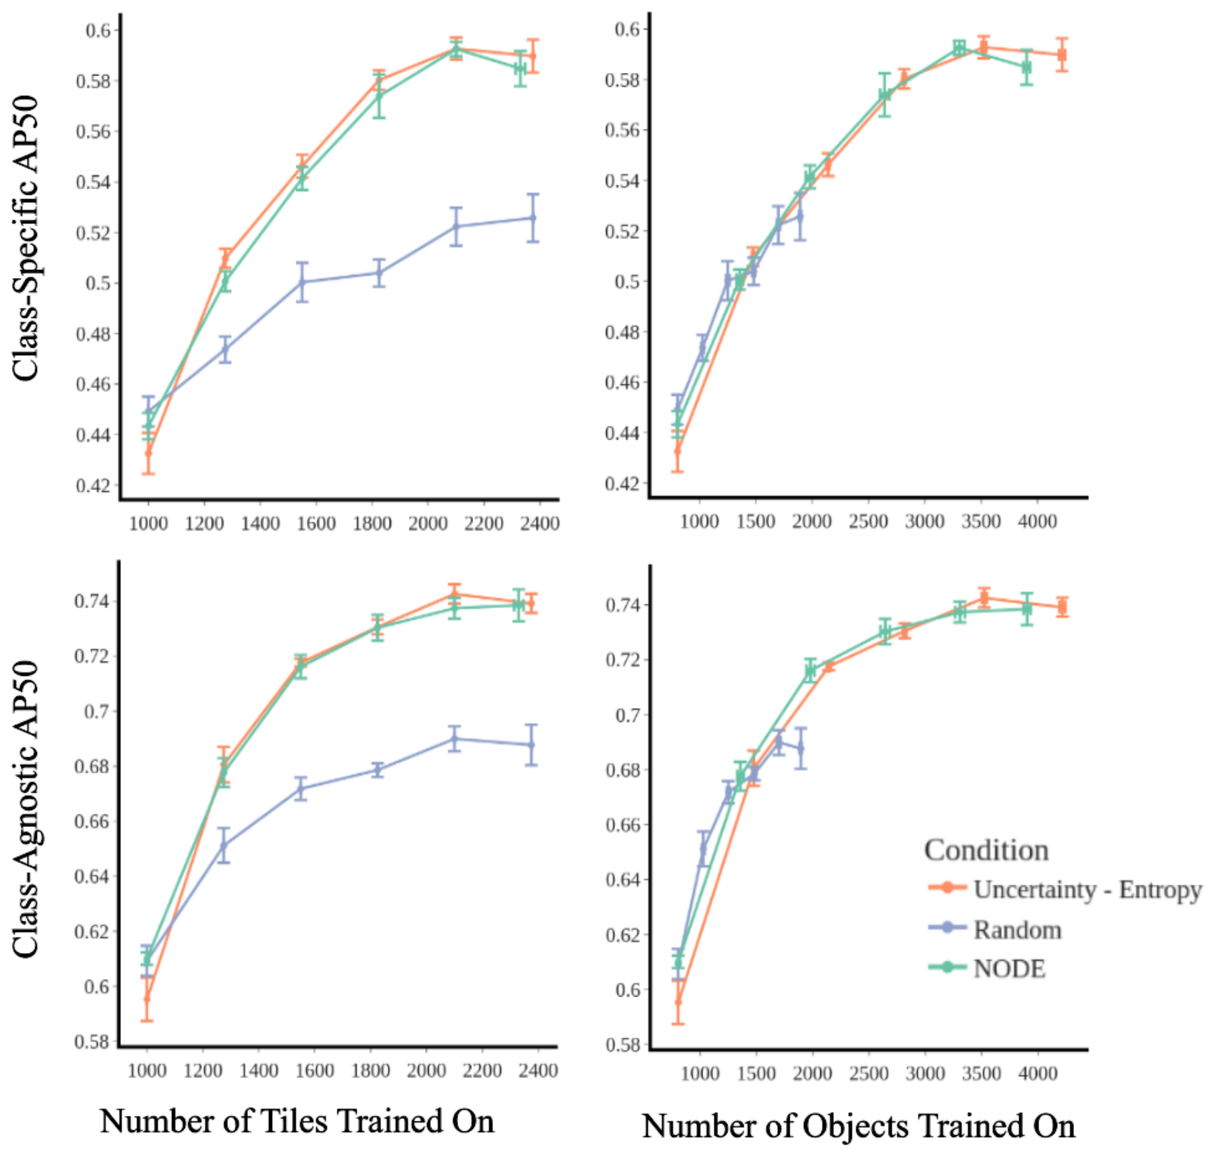

Figure 3: Average Precision metrics plotted versus annotated tiles (left) and number of annotations (right).

NODE is just as efficient in creating an optimal model as Uncertainty when considering the number of tiles (i.e., samples) that must be annotated. However, the NODE method finds tiles with fewer and more class-diverse objects, making the number of objects found more impactful for the model. Analysts would have to annotate the same number of tiles, but they would be able to annotate faster since there are fewer objects. The overall performance of the models reaches the same level within experimental uncertainties.

To illustrate this effect, Figure 4 shows the number of objects of each class in the Active Learning Iterations. The Uncertainty method selects more of the most common classes (classes 1 & 5) and selects fewer of the rarer classes. Specifically, NODE selected more of classes 4, 6, and 7, which improves detection performance as shown in Figure 5.

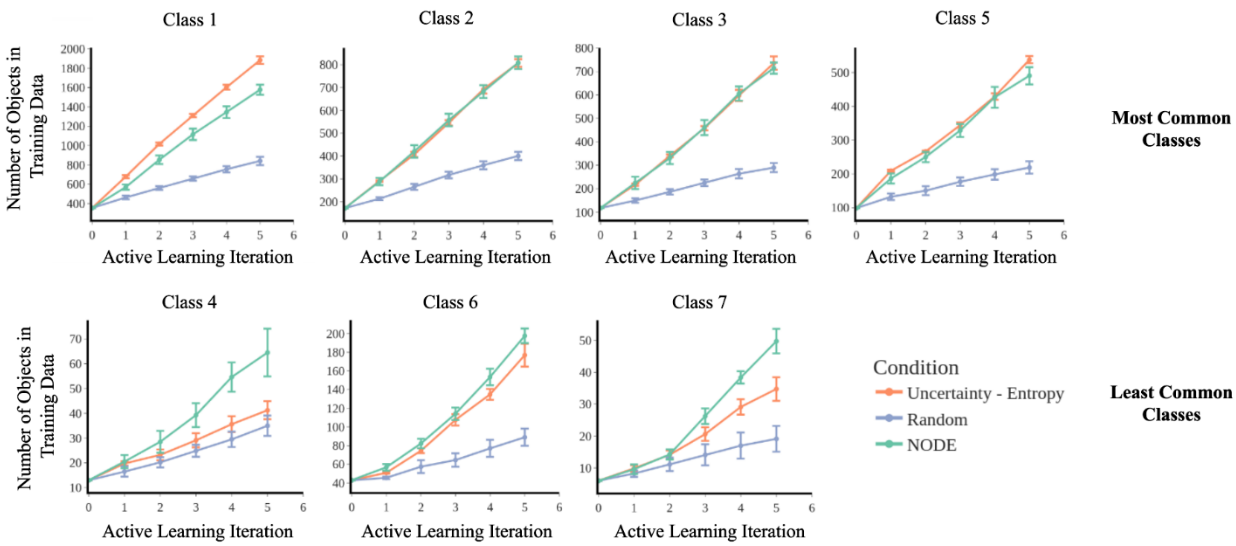

Figure 4: Number of objects in the Active Learning Iterations by each class. Classes with the most objects are in the top row, classes that have fewer objects are in the bottom row.

The NODE method results in higher accuracy on the rarest classes, 4 and 7, while still matching the performance of the Uncertainty method for the more common classes 1, 2, 3, 5, and 6.

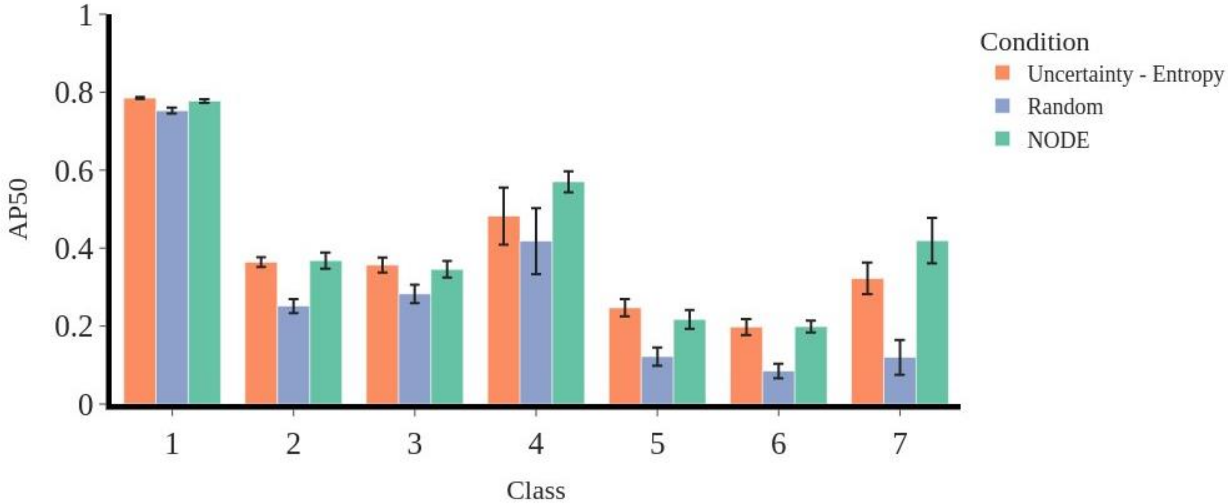

Figure 5: AP50 results per class after 5 iterations of Active Learning.

5 CONCLUSION

As machine learning becomes more prevalent in DoD analysis pipelines, active learning offers a promising solution to maximize labeling efficiency, which in turn helps deploy working models as quickly as possible. Most active learning research to date has targeted situations where large unlabeled datasets are already available. These methods are not optimal for more realistic DoD scenarios where data is collected and analyzed in parallel. Streaming active learning methods exist but have focused largely on the development of uncertainty-based methods. In this work, we develop one of the first diversity-based active learning metrics for streaming data. Our metric, NODE, can reduce the manual labeling effort required during active learning, particularly during early iterations, and improve performance on rarer classes in imbalanced data sets.

Future work should focus on combining NODE with existing uncertainty methods to develop hybrid methods for streaming scenarios. For example, a new hybrid metric could be created using a weighted sum of NODE and an uncertainty metric such as entropy. Alternatively, given the success of NODE in early iterations, future work could implement NODE in the first few iterations of active learning and then switch to an uncertainty-based method once the model is more mature. Overall, our work can help analysts efficiently label data, especially in cases where the target of interest is rare.

REFERENCES

1. Zhao, Zhong-Qiu, Peng Zheng, Shou-tao Xu, and Xindong Wu. "Object detection with deep learning: A review." *IEEE transactions on neural networks and learning systems* 30, no. 11 (2019): 3212-3232.
2. Zaidi, Syed Sahil Abbas, Mohammad Samar Ansari, Asra Aslam, Nadia Kanwal, Mamoona Asghar, and Brian Lee. "A survey of modern deep learning based object detection models." *Digital Signal Processing* 126 (2022): 103514.
3. Ren, Shaoqing, Kaiming He, Ross Girshick, and Jian Sun. "Faster R-CNN: Towards real-time object detection with region proposal networks." *IEEE transactions on pattern analysis and machine intelligence* 39, no. 6 (2016): 1137-1149.
4. Takezoe, Rinyoichi, Xu Liu, Shunan Mao, Marco Tianyu Chen, Zhanpeng Feng, Shiliang Zhang, and Xiaoyu Wang. "Deep active learning for computer vision: Past and future." *APSIPA Transactions on Signal and Information Processing* 12, no. 1 (2023).
5. Settles, Burr. "Active learning literature survey." (2009).
6. Ash, Jordan T., Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. "Deep batch active learning by diverse, uncertain gradient lower bounds." *arXiv preprint arXiv:1906.03671* (2019).
7. Agarwal, Sharat, Himanshu Arora, Saket Anand, and Chetan Arora. "Contextual diversity for active learning." In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVI 16*, pp. 137-153. Springer International Publishing, 2020.
8. Hekimoglu, Aral, Adrian Brucker, Alper Kagan Kayali, Michael Schmidt, and Alvaro Marcos-Ramiro. "Active Learning for Object Detection with Non-Redundant Informative Sampling." *arXiv preprint arXiv:2307.08414* (2023).
9. Wu, Jiaxi, Jiaxin Chen, and Di Huang. "Entropy-based active learning for object detection with progressive diversity constraint." In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9397-9406. 2022.
10. Ren, Pengzhen, Yun Xiao, Xiaojun Chang, Po-Yao Huang, Zhihui Li, Brij B. Gupta, Xiaojiang Chen, and Xin Wang. "A survey of deep active learning." *ACM computing surveys (CSUR)* 54, no. 9 (2021): 1-40.
11. Sener, Ozan, and Silvio Savarese. "Active learning for convolutional neural networks: A core-set approach." *arXiv preprint arXiv:1708.00489* (2017).
12. Moses, Aimee, Christopher W. Bogart, Kelsey O'Haire, Lidia Solorzano, and Elinor Yeo. "Diversity-based active learning: creating a representative object detection dataset in 3D point clouds." In *Geospatial Informatics XIII*, vol. 12525, pp. 136-155. SPIE, 2023.
13. Aghdam, Hamed H., Abel Gonzalez-Garcia, Joost van de Weijer, and Antonio M. López. "Active learning for deep detection neural networks." In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 3672-3680. 2019.

14. Wu, Mingfei, Chen Li, and Zehuan Yao. "Deep active learning for computer vision tasks: methodologies, applications, and challenges." *Applied Sciences* 12, no. 16 (2022): 8103.
15. Ghita, Ahmed, Bjørk Antoniussen, Walter Zimmer, Ross Greer, Christian Creß, Andreas Møgelmose, Mohan M. Trivedi, and Alois C. Knoll. "ActiveAnno3D--An Active Learning Framework for Multi-Modal 3D Object Detection." *arXiv preprint arXiv:2402.03235* (2024).
16. Feng, Di, Xiao Wei, Lars Rosenbaum, Atsuto Maki, and Klaus Dietmayer. "Deep active learning for efficient training of a lidar 3d object detector." In *2019 IEEE Intelligent Vehicles Symposium (IV)*, pp. 667-674. IEEE, 2019.
17. Cacciarelli, Davide, and Murat Kulahci. "Active learning for data streams: a survey." *Machine Learning* 113, no. 1 (2024): 185-239.
18. Žliobaitė, Indrė, Albert Bifet, Bernhard Pfahringer, and Geoffrey Holmes. "Active learning with drifting streaming data." *IEEE transactions on neural networks and learning systems* 25, no. 1 (2013): 27-39.
19. Wang, Keze, 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, no. 12 (2016): 2591-2600.
20. Kao, Chieh-Chi, Teng-Yok Lee, Pradeep Sen, and Ming-Yu Liu. "Localization-aware active learning for object detection." In Computer Vision–ACCV 2018: 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part VI 14, pp. 506-522. Springer International Publishing, 2019.
21. Nguyen, Vu-Linh, Mohammad Hossein Shaker, and Eyke Hüllermeier. "How to measure uncertainty in uncertainty sampling for active learning." Machine Learning 111, no. 1 (2022): 89-122.
22. Xie, Binhui, Shuang Li, Qingju Guo, Chi Liu, and Xinjing Cheng. "Annotator: A generic active learning baseline for lidar semantic segmentation." Advances in Neural Information Processing Systems 36 (2023).
23. Beluch, William H., Tim Genewein, Andreas Nürnberger, and Jan M. Köhler. "The power of ensembles for active learning in image classification." In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 9368-9377. 2018.
24. Gal, Yarin, Riashat Islam, and Zoubin Ghahramani. "Deep bayesian active learning with image data." In International conference on machine learning, pp. 1183-1192. PMLR, 2017.
25. Moses, Aimee, Srikanth Jakkampudi, Cheryl Danner, and Derek Biega. "Localization-based active learning (LOCAL) for object detection in 3D point clouds." In Geospatial Informatics XII, vol. 12099, pp. 44-58. SPIE, 2022.
26. Shuyang, Zhao, Toni Heittola, and Tuomas Virtanen. "An active learning method using clustering and committee-based sample selection for sound event classification." In 2018 16th International Workshop on Acoustic Signal Enhancement (IWAENC), pp. 116-120. IEEE, 2018.

27. Xu, Zhao, Kai Yu, Volker Tresp, Xiaowei Xu, and Jizhi Wang. "Representative sampling for text classification using support vector machines." In Advances in Information Retrieval: 25th European Conference on IR Research, ECIR 2003, Pisa, Italy, April 14–16, 2003. Proceedings 25, pp. 393-407. Springer Berlin Heidelberg, 2003.
28. Sener, Ozan, and Silvio Savarese. "Active learning for convolutional neural networks: A core-set approach." arXiv preprint arXiv:1708.00489 (2017).
29. Geifman, Yonatan, and Ran El-Yaniv. "Deep active learning over the long tail." arXiv preprint arXiv:1711.00941 (2017).
30. Kurlej, Bartosz, and Michal Wozniak. "Active learning approach to concept drift problem." Logic Journal of IGPL 20, no. 3 (2012): 550-559.
31. Liu, Sanmin, Shan Xue, Jia Wu, Chuan Zhou, Jian Yang, Zhao Li, and Jie Cao. "Online active learning for drifting data streams." IEEE Transactions on Neural Networks and Learning Systems 34, no. 1 (2021): 186-200.
32. Zhang, Hang, Weike Liu, Jicheng Shan, and Qingbao Liu. "Online active learning paired ensemble for concept drift and class imbalance." IEEE Access 6 (2018): 73815-73828.
33. Rajendran, Prajit T., Huascar Espinoza, Agnes Delaborde, and Chokri Mraidha. "Stream-Based Active Learning with Adaptive Uncertainty and Diversity Thresholds." In 2023 IEEE 33rd International Workshop on Machine Learning for Signal Processing (MLSP), pp. 1-6. IEEE, 2023.
34. Zhou, Yin, and Oncel Tuzel. "Voxelnet: End-to-end learning for point cloud based 3d object detection." In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 4490-4499. 2018.