

Impact of Stop Sets on Stopping Active Learning for Text Classification

Luke Kurlandski
Department of Computer Science
The College of New Jersey
Ewing, New Jersey 08618
Email: kurlanl1@tcnj.edu

Michael Bloodgood
Department of Computer Science
The College of New Jersey
Ewing, New Jersey 08618
Email: mbloodgood@tcnj.edu

Abstract—Active learning is an increasingly important branch of machine learning and a powerful technique for natural language processing. The main advantage of active learning is its potential to reduce the amount of labeled data needed to learn high-performing models. A vital aspect of an effective active learning algorithm is the determination of when to stop obtaining additional labeled data. Several leading state-of-the-art stopping methods use a stop set to help make this decision. However, there has been relatively less attention given to the choice of stop set than to the stopping algorithms that are applied on the stop set. Different choices of stop sets can lead to significant differences in stopping method performance. We investigate the impact of different stop set choices on different stopping methods. This paper shows the choice of the stop set can have a significant impact on the performance of stopping methods and the impact is different for stability-based methods from that on confidence-based methods. Furthermore, the unbiased representative stop sets suggested by original authors of methods work better than the systematically biased stop sets used in recently published work, and stopping methods based on stabilizing predictions have stronger performance than confidence-based stopping methods when unbiased representative stop sets are used. We provide the largest quantity of experimental results on the impact of stop sets to date. The findings are important for helping to illuminate the impact of this important aspect of stopping methods that has been under-considered in recently published work and that can have a large practical impact on the performance of stopping methods for important semantic computing applications such as technology assisted review and text classification more broadly.

I. INTRODUCTION

Active learning (AL) is a widely used strategy for reducing data labeling costs. AL has been used for many semantic computing applications within natural language processing (NLP), such as text classification (TC) [1]–[7], named entity recognition (NER) [8], [9], and machine translation (MT) [10]–[12]. Problem settings and recent developments in AL are discussed further in [13].

The essential concept behind AL is for the machine learning (ML) model to take an operational role in deciding which data instances should be labeled for learning. The active learning process begins with an unlabeled pool of data, an initially empty pool of labeled training data, and a batch size. In the first step of AL, the model queries an oracle (a human data labeler) for the labels of an initial batch of training data. Following this, the algorithm enters an iterative process: the model is trained on the existing set of labeled data; the model

queries the oracle to label a new batch of training examples; and the labeled and unlabeled sets are adjusted accordingly. This process is shown in Algorithm 1, where b refers to the batch size, L the set of labeled examples, U the unlabeled pool of examples, and B the batch set of examples queried at each iteration of AL.

Algorithm 1 Active Learning Algorithm

Require: U, b {input to algorithm}

$L \leftarrow$ select and annotate b random examples from U

$U \leftarrow U - L$

repeat

 Train model using L

$B \leftarrow$ select and annotate b examples from U using the selection algorithm

$U \leftarrow U - B$

$L \leftarrow L \cup B$

until stopping criterion is met

An essential component to any active learning setup is an effective stopping method that halts the data labeling process. Without a stopping method, the model would request annotations for every training example, defeating the purpose of AL. Past work has used the term “aggressive” to characterize stopping methods that tend to prioritize saving annotations and the term “conservative” to characterize stopping methods that tend to maximize model performance [8]. In general, it is desirable for stopping methods to be:

- capable of reducing the number of labels required without sacrificing final performance,
- capable of adjustment to be more or less aggressive, and
- applicable and consistent across many AL settings, base learners, tasks, and datasets.

A simple stopping method would be to stop after an ad-hoc predetermined labeling budget has been exhausted. However, this runs a large risk of either obtaining a significantly lower-performing model or wasting large amounts of data labeling effort.

Another simple approach would be to use a small set of labeled data to evaluate performance during learning and stop when performance levels off. However, previous work has

found that not only does this method have the drawback of requiring additional labeled data, but also that this method doesn't work as well as state-of-the-art stopping methods that base the stopping decision on the result of clever measurements made on a large set of unlabeled data [14]. Such a large set of unlabeled data has been termed a "stop set" [8].

Two of the most widely used classes of stopping methods are stabilizing predictions-based methods [8], [15], [16] and confidence-based methods [17].

A well-known confidence-based stopping method stops AL when model confidence on a large, representative stop set begins to decline [17]. We discuss confidence-based stopping methods in greater detail in section III.

Stabilizing Predictions (SP) stops AL when the agreement between successively trained models exceeds some threshold. SP computes the agreement between two models by comparing the predictions of each model on a large, representative, randomly selected set of examples that does not have to be labeled. We discuss the SP method in greater detail in section III.

Two new stopping methods were recently proposed in [18] and presented as having better performance than the leading existing stopping methods, however, the use of SP-stopping in those experiments did not use a proper stop set. Instead of a large unbiased representative stop set, the remaining unlabeled pool was used, which is a shrinking and systematically biased stop set.

In this paper we investigate the impact of using the remaining unlabeled pool as a stop set and various randomly selected stop sets of different sizes. We show that the choice of stop set has a substantial impact on performance for both stability-based stopping and confidence-based stopping. Furthermore, the impact is different depending on whether the stopping method is a stability-based method or a confidence-based method.

In extensive experiments with multiple freely and publicly available widely used benchmark datasets, which will allow for others to reproduce our experiments or run further experiments, we find that when SP is used with an unbiased representative stop set as originally proposed in [8], the performance of SP is stronger than that of the confidence-based methods, including those proposed in [18]. The methods proposed in [18] are similar to the stopping criterion proposed in [17]. We experimentally show the extent to which the method from [17] has more conservative performance than the methods proposed in [18]. We find that the methods from [18] and [17] perform better compared with themselves when using an unbiased representative stop set as originally proposed in [17].

Our findings are important for helping to illuminate the impact of stop set choice on stopping method performance. We show experimentally that stop set choice can have a large practical impact on the performance of stopping methods for important semantic computing applications such as technology assisted review, and text classification more broadly. Contributions in this paper include:

- 1) Illuminate the impact of the choice of using the remaining unlabeled pool as a stop set for SP in the experiments of [18] and show how such a choice impacts SP stopping behavior.
- 2) Compare several stopping methods, using multiple stop sets with each stopping method, on several publicly available datasets to experimentally show how performance is affected by stop set choice.
- 3) Experimentally show the extent to which use of the biased remaining unlabeled pool as a stop set will tend to make stability-based methods stop more aggressively and make confidence-based methods stop more conservatively than when using an unbiased representative stop set.

Section II discusses related work, section III provides detailed descriptions of the stopping methods, section IV discusses our experimental setup, section V presents our experimental results and analysis, and section VI concludes.

II. RELATED WORK

Overall, given the large amount of research on stopping methods, there has been relatively less attention given to the impact of different stop sets than to the design and testing of the stopping method computations performed on the stop sets. To our knowledge there is no past work that focuses on investigating the impact of different stop sets on stopping method performance. There has been some consideration of ideas related to stop set impact as an auxiliary topic in several past works, which we explain here. In contrast to all of this past work, the current paper is focused on clarifying the importance and the impact of stop set selection on stopping method performance.

A widely applicable stopping method based on model confidence is presented in [17]. The method is intended to be used in AL settings where uncertainty sampling is the query strategy and there is an expectation of a rise-peak-drop pattern in model confidence as AL proceeds. The method in [17] stops AL based on model confidence measured on a large separate representative set of examples. After each iteration of AL, this method runs the newly learned model over the large set of examples and estimates its confidence. When model confidence consistently drops on the large set of examples, this method indicates AL should stop. Although the term *stop set* is not used, in [17] there is some discussion of the set of examples that should be used for measuring model confidence and it is indicated that the set should be a separate large set of examples to minimize the risk of not being representative. There are little concrete details about the set of examples used in the experiments reported in [17]; it appears the test set might have been used, which would not be ideal as it could lead to overestimating the effectiveness of the stopping method. We experiment with this method using different stop sets and measure the impact of the different stop sets on performance. We refer to this method as Declining Confidence (DC).

A widely applicable stopping method based on stabilizing predictions is presented in [8]. This is a particularly important

method to investigate because it is widely applicable with few if any restrictions on when it can be used, and it has been found to work well in many settings [8], [14], [19]–[24]. The method stops when successively trained models have high agreement in terms of their predictions on the examples in a large randomly selected set of examples. The authors call this set of examples a *stop set* [8]. In the main experiments, a large randomly selected fixed set of examples is used as the stop set. There is a small section of the paper that explores how stopping performance is impacted if the stop set is made smaller and it’s found that if the set becomes too small the performance could begin to degrade, but it’s concluded that the size required to be representative could vary depending on dataset and task so future work is needed to determine how large the stop set should be. We experiment with this method using different stop sets and measure the impact of the different stop sets on performance. We refer to this method as Stabilizing Predictions (SP).

A theoretical analysis of SP stopping is presented in [15]. An equation for the variance in the agreement estimator between successive models is derived, which shows that the variance is inversely proportional to the size of the stop set. The variance estimator is suggested as a principled way to check stop set size in order to make sure reliable estimates of agreement can be obtained.

Using labeled data versus using unlabeled data to inform the stopping decision is considered in [14]. The tradeoff is that although the labels might give some extra information, that has to be balanced with the fact that using a stop set of labeled data for many applications will by necessity have to be kept relatively small whereas a stop set of unlabeled data can be large. It has been found that stopping methods that use unlabeled data in their decision are more effective than methods that use labeled data.

Two stopping methods are presented in [18]. They are similar to the Declining Confidence (DC) method from [17]. The difference is that whereas the DC method requires a consistent drop in confidence, the methods from [18] require a consistent period where confidence is not increasing, i.e., it could be staying the same or dropping. Another difference is that the methods were implemented with two different stop sets: the batch set, B , and the unlabeled pool, U , instead of a large representative set as suggested by [17]. We refer to these methods as Non-increasing Confidence (NC).

In [18] the performance of the NC methods are compared with current state-of-the-art methods, including SP, on a private technology assisted sensitivity review dataset. The conclusion reached in [18] is that the leading method, SP, stops too aggressively on their private dataset and does not produce a model as well-learned as when using their NC methods. However, it is important to observe that SP was not implemented correctly in the experiments in [18]. Instead of using a large unbiased representative stop set, the remaining unlabeled pool was used as the stop set. We hypothesize that this choice can cause SP to stop more aggressively and confidence-based methods such as DC and NC to stop more

conservatively. In the current paper, we show that this is indeed the case and explain the reasons why. We compare the performance of NC with DC, which is not considered in [18], and quantify how much more conservative DC is on several datasets. We also compare the performance of NC with SP implemented in the normal way, which is not considered in [18], and find that then SP has compellingly stronger performance.

III. METHODS FOR STOPPING

In this section, we provide more details about the Stabilizing Predictions, Declining Confidence, and Non-increasing Confidence stopping methods, including details about their usage of stop sets.

A. Stabilizing Predictions

Stabilizing Predictions determines when to stop based upon the agreement between successively trained models [8]. The agreement is computed between the models’ two sets of predictions on a large, randomly selected, fixed stop set of training examples. There are several ways the agreement can be computed. In [8] Cohen’s Kappa statistic [25] is suggested to account for agreement due to chance.

We compute agreement in our implementation of SP as given by equation 1

$$agreement = \frac{A_o - A_e}{1 - A_e} \quad (1)$$

where A_o is the observed agreement between the two models and A_e is the agreement expected between the two models by chance. We compute A_o by the proportion of examples the two models agree upon. We compute A_e as in equation 2

$$A_e = \sum_{k \in \{+1, -1\}} P(k|c_1)P(k|c_2) \quad (2)$$

where c_i are the two models whose agreement is being computed and $P(k|c_i)$ is the probability that model c_i labels an example as being in category k [8], [25].

The SP method determines when to stop based on whether the average of the agreements over a window of size k exceeds a threshold K , where k and K are user-defined parameters. The SP method can be adjusted to be more aggressive by decreasing k and K , and can be adjusted to be more conservative by increasing k and K . Previous research has suggested that $k = 3$ and $K = 0.99$ work well as default values [8].

The stop set of examples on which SP computes the agreement between successively trained models is an important part of the stopping method [8], [14], [15]. Since SP uses this set of examples to make a generalization about the data distribution for the intended application of the learned model, the stop set should be an unbiased representation of the data. If it is systematically biased, the agreement between successively learned models on the elements in the stop set might not be indicative of the agreement between the learned models on application data as a whole. A larger stop set is more likely to be representative and unbiased, but a smaller set is more

computationally efficient for computing agreement between the learned models. An easy way to select the stop set is to allocate a randomly chosen fixed percent of the data to it, such as 50%, as in [14]. We denote such a randomly chosen unbiased stop set as S . Previous research found that the variance in the Kappa statistic [15] is inversely proportional to the size of the stop set. High variance in Kappa could cause SP to behave unpredictably so the variance check from [15] is recommended as a principled way to ensure that the size is large enough to get low-variance estimates of agreement.

Recently, in [18] SP was implemented using the remaining unlabeled pool, U , as the stop set. We will refer to this method as SP_U and the original method as SP or SP_S . The choice of using the remaining unlabeled pool as the stop set is not explained in [18], providing evidence that the community could be under-appreciating the importance of stop set selection. When using the remaining unlabeled pool as the stop set, the set is changing at each iteration of AL. In particular, the size of this stop set is shrinking at each iteration of AL. Therefore, the number of examples in the remaining unlabeled pool at iteration $i - 1$ and iteration i will be different. Cohen’s Kappa agreement needs to be computed on the same set of examples [25]. When we implement SP_U , to ensure equally sized prediction lists, we compute the Kappa agreement between models trained at iteration $i - 1$ and iteration i , but only use the examples in the unlabeled pool at iteration i . In contrast to S , U is not representative of the application space, is systematically biased, and becomes increasingly more biased as AL proceeds. The extent and way that U is biased will depend on the details of the AL setting, including the base learner that is used and the query strategy that is used. In [18] a support vector machine (SVM) base learner is used with uncertainty sampling based on closest-to-hyperplane selection. This is known to result in a set of labeled data and a remaining set of unlabeled data that are biased [26]. Intuitively, this makes sense because the algorithm is systematically selecting the examples that are the most difficult to classify to be labeled, leaving U to consist increasingly of examples that are easier to classify as AL proceeds. We hypothesize that using this biased set as a stop set can be intuitively expected to yield higher agreement scores sooner than if an unbiased set were used. Our experimental results in section V confirm this hypothesis and help shed light on the extent of the impact.

B. Declining Confidence

The Declining Confidence method is based on the premise that once all the informative examples in the unlabeled pool are labeled, further labeling could mislead the model and decrease its confidence [17]. DC determines when to stop AL based on when the model’s confidence decreases for ϵ iterations where ϵ is a user-defined parameter. Previous research has suggested that using $\epsilon = 3$ works well as a default.

Given a stop set of training examples, D , DC measures the currently learned model’s confidence by taking the mean of its confidence scores on every example in D . Using this notation,

the DC method stops when the *Conf* score, as in equation 3, decreases for ϵ consecutive iterations

$$Conf(D) = \frac{1}{|D|} \sum_{d \in D} confidence(d) \quad (3)$$

where *confidence*(d) is a measure of the model’s confidence on datapoint d . In [17] when a support vector machine base learner is used, the confidence is measured by the distance of the datapoint d from the separating hyperplane learned by the current model.

Recently, in [18], two stopping methods were presented that are similar to DC from [17]. One difference is that the methods in [18] use the batch set, B , and the remaining unlabeled pool, U , as their stop sets. In [17], similarly to [8], a large representative set is suggested to be used with DC. To investigate the impact of these different stop set choices, we implement DC using B , U , and a large randomly selected unbiased stop set S . We call these methods DC_B , DC_U , and DC_S , respectively¹. Each method stops when $Conf(B)$, $Conf(U)$, or $Conf(S)$ decreases for $\epsilon = 3$ iterations, respectively. We acknowledge that implementing DC with B and U is contrary to Vlachos’ guidance about how to select the stop set because these sets are biased, changing, and in the case of B , small. We do so for illumination of the impact of stop set selection on the method since these sets were suggested for use as stop sets with the similar NC method in [18].

C. Non-increasing Confidence

The two stopping methods presented in [18], TotalConf and LeastConf, both stop AL if the model’s *Conf* score (equation 3) measured over a stop set of training examples does not increase for ϵ iterations, where ϵ is a user-defined parameter. In [18] the value of ϵ is set to three. Since the methods from [18] stop when the model’s confidence stops increasing, we refer to them as Non-increasing Confidence (NC). The difference between TotalConf and LeastConf is that TotalConf uses U to compute confidence and LeastConf uses B . For brevity, we will call LeastConf NC_B and TotalConf NC_U .

We also implement NC with a large fixed unchanging unbiased representative stop set as recommended in [8], [17] and refer to this implementation of NC stopping as NC_S . We acknowledge that this was not suggested by the original authors in [18]. We do so to illuminate the impact of stop set choice on NC stopping methods and we also hypothesize that using such an unbiased fixed stop set will strengthen the performance of NC stopping, which is confirmed experimentally in section V.

IV. EXPERIMENTAL SETUP

In our experiments, we perform text classification with seven widely used publicly available datasets. To make our results as reproducible as possible, we use the train-test split

¹Observe that neither [18] nor the current paper use B as a stop set with SP because B is a completely new set at each iteration of AL, and therefore, cannot be used with SP.

recognized by the community if such a split exists. Otherwise, we perform 10-fold cross validation.

We use the “bydate” version of the 20NewsGroups² dataset and classify its newsgroup postings into 20 categories. We use the Reuters-21578 Distribution 1.0 ModApte split³ as in [27], [28] and classify the 10 largest categories of business news articles. We classify email spam from the SpamAssassin⁴ corpus and the ham25 version of the TREC⁵ corpus. We use the results from the four largest categories of academic webpages in the WebKB⁶ dataset as in [29], [30], [31], [8]. We use 10-fold cross validation for SpamAssassin, TREC, and WebKB; we macro average our results across folds and categories.

In order to compare directly with the technology assisted sensitivity review experiment conducted in [18] we tried to obtain the dataset that was used in that paper. Unfortunately, this dataset is private and cannot be shared with other researchers (personal communication with Graham McDonald). Instead, we perform TC on the publicly available RCV1-v2 dataset [32] that has been used frequently in past works to simulate TAR [33], [34], [35], [36], [37], [38]. Most of these authors use a subset of the 103 RCV1-v2 categories; we use all 103 of them. RCV1-v2 is computationally demanding, so we do not use it in all of our experiments.

We use a support vector machine with a linear kernel and $C=1$, as done by [8], [17], [18]. We use the closest-to-hyperplane selection algorithm as the query strategy during AL, as done by [8], [17], [18]. We also use random selection, that is, passive learning, for a comparison experiment. We allocate 50% of training examples to the randomly selected stop set, S , as in [14]. We use a batch size of 0.5%, consistent with the findings of [39]. We use a bag-of-words approach that only considers words that appear more than three times. For all datasets other than RCV1-v2⁷, during preprocessing we remove 174 stop words using the *Default English Stopwords List*⁸.

When comparing stopping methods, we report the number of annotations (i.e., labels) acquired (denoted by ANN in our tables), the proportion of annotations of the entire dataset acquired (ANN-P in tables), the learned model’s F-1 measure (F1), F-2 measure (F2), accuracy (ACC), and balanced accuracy (BAC), all at the time the stopping method indicates to stop AL. These performance measures were used in [18], which will help to facilitate comparison. We also display these statistics when the model is given all possible training data (denoted by the column “Final” in our tables). F- β Measure is a weighted harmonic mean between precision and recall, where β is a constant such that recall is β times more important than precision. Accuracy is the proportion of the

model’s predictions that are correct. Balanced accuracy is the average of the recall scores for each class. We report these statistics for individual datasets, as well as a macro average across all datasets. In our tables we boldface entries that are the best-performing for each metric. Lower numbers are better for ANN and ANN-P because that reflects reducing the amount of labeled data required. Higher numbers are better for the other metrics since that reflects better system performance for end-users.

V. RESULTS AND ANALYSIS

This section displays the results of our experiments and analyzes the results. Table I shows the notation used throughout this section.

Symbol	Type	Definition
SP	stopping method	Stabilizing Predictions
DC	stopping method	Declining Confidence
NC	stopping method	Non-increasing Confidence
S	stop set	Standard fixed unbiased set of examples
B	stop set	Batch set of examples
U	stop set	Unlabeled pool of examples

TABLE I

NOTATION USED IN THIS SECTION. STOPPING METHODS AND STOP SET MAY BE COMBINED USING SUBSCRIPT NOTATION. FOR EXAMPLE, SP_S INDICATES THAT THE STABILIZING PREDICTIONS STOPPING METHOD IS USED WITH A STANDARD FIXED UNBIASED STOP SET.

A. Impact of Stop Sets on Stabilizing Predictions

Table II shows the performance of SP_S and SP_U for SVM-closest and SVM-random on all datasets (except RCV1-v2 for computational feasibility).

As expected, active learning significantly outperforms passive learning, achieving similar performance with far fewer annotations. When the uncertainty sampling algorithm is used in active learning, we see that SP_U is more aggressive than SP_S . We know from [15] that as the unlabeled pool shrinks, the variance in the Kappa statistic increases, and Stabilizing Predictions becomes more unstable. It was hypothesized in [18] that the unlabeled pool becoming small in the later stages of AL could cause the more aggressive behavior of SP_U .

However, when examining the size of the unlabeled pool at the point when SP_U stops AL, we see that U is large at the stopping point for all datasets. At the stopping point, U contains between 85% and 96% of the total dataset (91% on average). Since U never gets small, the variance of the Kappa estimator never gets large. Therefore, the decreasing size of U is not the main cause of the more aggressive stopping behavior of SP_U .

When an uncertainty sampling query strategy is used, B and U become biased sets. Since uncertainty sampling selects the most uncertain training examples for labeling, elements in B are systematically biased such that they are difficult for the model to predict correctly. At each round of AL, the elements of B are removed from U . This in turn causes the elements in U to be biased such that they are easier for the model to predict correctly. Since the examples in U are easier to predict, the

²<http://qwone.com/~jason/20Newsgroups/>

³<http://www.daviddlewis.com/resources/testcollections/reuters21578/>

⁴<https://spamassassin.apache.org/old/publiccorpus/>

⁵<https://plg.uwaterloo.ca/cgi-bin/cgiwrap/gvcormac/foo>

⁶<http://www.cs.cmu.edu/afs/cs.cmu.edu/project/theo-20/www/data/>

⁷RCV1-v2 already comes with stop words removed.

⁸<https://www.ranks.nl/stopwords>

Dataset	Stat	Active Learning			Passive Learning		
		SP _U	SP _S	Final	SP _U	SP _S	Final
20News	ANN	974	1088	11314	3300	2696	11314
	ANN-P	0.086	0.096	1.000	0.292	0.238	1.000
	F1	0.714	0.715	0.725	0.623	0.606	0.725
	F2	0.662	0.664	0.681	0.558	0.539	0.681
	ACC	0.972	0.973	0.973	0.964	0.963	0.973
	BAC	0.812	0.814	0.824	0.757	0.748	0.824
Reuters	ANN	585	696	9655	1190	1084	9655
	ANN-P	0.061	0.072	1.000	0.123	0.112	1.000
	F1	0.786	0.782	0.761	0.709	0.708	0.761
	F2	0.746	0.738	0.717	0.654	0.654	0.717
	ACC	0.968	0.969	0.967	0.961	0.96	0.967
	BAC	0.856	0.851	0.840	0.806	0.806	0.840
Spam Assassin	ANN	618	655	10351	899	863	10351
	ANN-P	0.060	0.063	1.000	0.087	0.083	1.000
	F1	0.983	0.985	0.997	0.956	0.957	0.997
	F2	0.989	0.991	0.997	0.960	0.960	0.997
	ACC	0.989	0.991	0.998	0.973	0.973	0.998
	BAC	0.990	0.992	0.998	0.970	0.970	0.998
TREC	ANN	5260	5565	152446	8615	8081	152446
	ANN-P	0.034	0.036	0.999	0.056	0.053	0.999
	F1	0.988	0.989	1.000	0.971	0.97	1.000
	F2	0.986	0.986	1.000	0.967	0.966	1.000
	ACC	0.996	0.997	1.000	0.991	0.991	1.000
	BAC	0.991	0.992	1.000	0.980	0.979	1.000
WebKB	ANN	2059	2216	14176	6707	5268	14176
	ANN-P	0.145	0.156	0.999	0.473	0.371	0.999
	F1	0.843	0.856	0.980	0.817	0.781	0.98
	F2	0.812	0.827	0.979	0.814	0.777	0.979
	ACC	0.963	0.966	0.995	0.958	0.949	0.995
	BAC	0.889	0.898	0.988	0.894	0.873	0.988
Avg	ANN	1899	2044	39588	4142	3598	39588
	ANN-P	0.077	0.085	1.000	0.206	0.172	1.000
	F1	0.863	0.865	0.893	0.815	0.804	0.893
	F2	0.839	0.841	0.875	0.791	0.779	0.875
	ACC	0.978	0.979	0.987	0.969	0.967	0.987
	BAC	0.908	0.909	0.930	0.881	0.875	0.930

TABLE II

THE NUMBER OF ANNOTATIONS AND MODEL PERFORMANCE WHEN SP_U AND SP_S STOP FOR ACTIVE LEARNING AND PASSIVE LEARNING FOR ALL DATASETS (EXCEPT RCV1 FOR COMPUTATIONAL FEASIBILITY). BOLDED ENTRIES ARE THE BEST-PERFORMING VALUE FOR THE RELEVANT STATISTIC.

consecutively trained models during AL will be more likely to agree on the predictions of the examples in U . This will make the Kappa agreement exceed its threshold sooner, and cause SP_U to stop more aggressively.

When a random sampling algorithm is used, that is, passive learning, U should be representative of the dataset except that it is by definition not permitted to have any representation in the training data. This will make U slightly harder to achieve agreement on when compared with an unbiased predetermined random set that is not artificially prevented from having representation in the training data. Therefore, we would expect SP_U to behave more conservatively during passive learning than SP_S. Table II confirms this, where we see that SP_U is more conservative than SP_S during passive learning.

SP_U requires a non-standard approach to compute Cohen’s Kappa agreement between successively trained models, was found to behave over-aggressively by [18] during AL, and uses a biased non-representative stop set. For these reasons, we suggest referring to the original authors of SP and implementing the method with a fixed unbiased representative stop set S . We compare SP_S with the confidence-based stopping methods in the next section.

B. Impact of Stop Sets on Confidence Based Stopping Methods and Comparison with the SP Method

In table III, we compare the stopping results from Declining Confidence and Non-increasing Confidence using the batch set, the unlabeled pool, and a large, randomly selected, fixed

set as stop sets during AL. We also add the performance of SP_S as a column for ease of comparison of stabilizing predictions stopping performance with confidence-based stopping performance.

Dataset	Stat	SP _S	DC _B	DC _U	DC _S	NC _B	NC _U	NC _S	Final
20News	ANN	1088	3395	5793	1707	265	1567	567	11314
	ANN-P	0.096	0.300	0.512	0.151	0.023	0.139	0.050	1.000
	F1	0.715	0.577	0.608	0.598	0.488	0.539	0.530	0.725
	F2	0.664	0.520	0.553	0.544	0.419	0.476	0.470	0.681
	ACC	0.973	0.962	0.963	0.962	0.956	0.958	0.957	0.973
	BAC	0.814	0.740	0.757	0.752	0.688	0.717	0.714	0.824
RCV1	ANN	6497	60375	67547	14422	20946	46463	8787	73697
	ANN-P	0.088	0.819	0.917	0.196	0.284	0.630	0.119	1.000
	F1	0.525	0.520	0.515	0.522	0.508	0.516	0.518	0.515
	F2	0.490	0.502	0.499	0.491	0.479	0.493	0.484	0.500
	ACC	0.982	0.979	0.979	0.981	0.981	0.980	0.981	0.979
	BAC	0.725	0.736	0.735	0.728	0.723	0.730	0.725	0.736
Reuters	ANN	696	6955	8713	1785	758	4202	720	9655
	ANN-P	0.072	0.720	0.902	0.185	0.079	0.435	0.075	1.000
	F1	0.782	0.711	0.746	0.761	0.683	0.706	0.706	0.761
	F2	0.738	0.663	0.703	0.714	0.622	0.656	0.656	0.717
	ACC	0.969	0.963	0.965	0.967	0.960	0.960	0.960	0.967
	BAC	0.851	0.813	0.833	0.839	0.789	0.807	0.808	0.840
Spam Assassin	ANN	655	10337	10337	2111	4579	10337	1424	10351
	ANN-P	0.063	0.999	0.999	0.204	0.442	0.999	0.138	1.000
	F1	0.985	0.999	0.999	0.994	0.996	0.999	0.993	0.997
	F2	0.991	0.999	0.999	0.997	0.998	0.999	0.996	0.997
	ACC	0.991	0.999	0.999	0.996	0.997	0.999	0.995	0.998
	BAC	0.992	0.999	0.999	0.997	0.998	0.999	0.996	0.998
TREC	ANN	5565	143103	152478	13113	113218	22716	12808	152531
	ANN-P	0.036	0.938	1.000	0.086	0.742	0.149	0.084	0.999
	F1	0.989	1.000	1.000	0.994	0.998	0.995	0.994	1.000
	F2	0.986	0.999	1.000	0.992	0.998	0.993	0.992	1.000
	ACC	0.997	1.000	1.000	0.998	0.999	0.998	0.998	1.000
	BAC	0.992	1.000	1.000	0.995	0.999	0.996	0.995	1.000
WebKB	ANN	2216	9981	11753	3846	958	6651	2241	14176
	ANN-P	0.156	0.703	0.828	0.271	0.068	0.469	0.158	0.999
	F1	0.856	0.870	0.940	0.893	0.588	0.825	0.784	0.980
	F2	0.827	0.855	0.929	0.871	0.546	0.798	0.750	0.979
	ACC	0.966	0.977	0.989	0.978	0.917	0.964	0.954	0.995
	BAC	0.898	0.919	0.960	0.925	0.746	0.884	0.857	0.988
Avg	ANN	2786	39024	42770	6164	23454	15323	4425	45273
	ANN-P	0.085	0.747	0.860	0.182	0.273	0.470	0.104	1.000
	F1	0.809	0.78	0.801	0.794	0.71	0.763	0.754	0.830
	F2	0.783	0.756	0.780	0.768	0.677	0.736	0.725	0.812
	ACC	0.980	0.980	0.982	0.980	0.968	0.976	0.974	0.985
	BAC	0.879	0.868	0.881	0.873	0.824	0.856	0.849	0.898

TABLE III

THE NUMBER OF ANNOTATIONS AND MODEL PERFORMANCE WHEN A STOPPING METHOD STOPS. BOLDED ENTRIES ARE THE BEST-PERFORMING VALUE FOR THE RELEVANT STATISTIC.

Table III shows that an unbiased stop set, S , is the best stop set to use with DC. Table III shows that when implemented with U and B , DC uses far more annotations than is required and does not achieve any meaningful performance boost compared to DC_S. Comparing DC_S with SP, from the average row, we see that DC_S uses more than twice as many annotations as SP and all performance measures except ACC are worse than those of SP.

For NC, we see that the methods presented in [18] have very different behavior on different datasets. For the individual datasets in table III, we see that NC_B does not learn an effective model for 20NewsGroups, Reuters, and WebKB, while it requests wastefully large numbers of annotations for RCV1, SpamAssassin, and TREC. NC_U does not learn a good model for 20NewGroups and requests many more annotations than are required for all other datasets. In both the individual dataset results and the overall average results, we see that NC performs best when implemented with a large representative stop set, S . This strategy was first implemented in this work and was not considered in [18], yet it appears to be the highest-performing variant of NC. In the average row, we see that NC_S attains essentially the same model performance as DC_S, but

uses about half of the annotations. Therefore, it is reasonable to conclude that NC_S is the best out of all the confidence-based methods. However, its performance is still not as strong as SP. The average row shows that SP_S requires less labeled data and produces higher-performing models across all performance measures than NC_S .

Figure 1 shows a representative learning curve for the 20NewsGroups dataset. This learning curve captures common trends that occur at the level of an individual fold or category: the methods that use U fail to stop, the methods that use B stop too early, and the methods that use S stop at reasonable points. The methods that use S are more stable than the other methods. Out of these methods, SP_S is the most efficient at saving annotations and producing a high performing model.




Fig. 1. Learning curve for the 20NewsGroups dataset, classifying articles as belonging to the talk-religion category or not. In this plot, DC_B and NC_B stopped at the same point, so their vertical lines are perfectly on top of each other. The same occurred for DC_U and NC_U .

VI. CONCLUSION

Active learning is used to produce high-performing models while using significantly less labeled training data than when using passive learning. Without an effective stopping method, AL could stop too late, wasting unnecessary data labeling effort, or AL could stop too early, producing a poorly learned model. Several AL stopping methods determine when to stop on the basis of the results of computations performed on a set of examples called the stop set. The choice of the stop set has received relatively less attention than development of the computations that the stopping methods perform on the stop set. In this work, we experiment with several different possible options for a stop set and find that the choice of the stop set can have a significant impact on stopping active learning. It is important to communicate these findings to prevent suboptimal stop sets from being used and suboptimal stopping methods from being used.

Authors of the widely used Stabilizing Predictions method and the Declining Confidence method had originally suggested

that representative unbiased sets be used as stop sets for their methods. Recent work used the remaining unlabeled pool and the batch set as stop sets in their new confidence-based methods and also used the remaining unlabeled pool as the stop set when implementing the Stabilizing Predictions method as a baseline for comparison. The remaining unlabeled pool and the batch set are systematically biased as active learning proceeds.

We find that different stop set choices impact different stopping methods in different ways. When Stabilizing Predictions is implemented with the unlabeled pool instead of an unbiased representative stop set, it becomes more aggressive compared with when it is implemented with an unbiased representative stop set. When confidence-based methods are implemented with the unlabeled pool instead of a large unbiased stop set, they become more conservative, sometimes failing to stop at all. When confidence-based methods are implemented with the batch set, they become less consistent in their behavior and sometimes may become overly conservative or overly aggressive. For both stabilizing predictions and confidence-based methods, we find that the best-performing stop set to use is a random unbiased set. Using one of the biased sets produces a less reliable and lower-performing method. Comparing confidence-based methods with stabilizing predictions methods, our results indicate that when the methods use randomly selected stop sets, the stabilizing predictions method saves many annotations and produces models with higher performance than the confidence-based methods across several publicly available text classification datasets.

ACKNOWLEDGMENT

This work was supported in part by The College of New Jersey Support of Scholarly Activities (SOSA) program. The authors acknowledge use of the ELSA high performance computing cluster at The College of New Jersey for conducting the research reported in this paper. This cluster is funded by the National Science Foundation under grant number OAC-1828163.

REFERENCES

- [1] D. D. Lewis and W. A. Gale, "A sequential algorithm for training text classifiers," in *SIGIR '94: Proceedings of the 17th annual international ACM SIGIR conference on Research and development in information retrieval*. New York, NY, USA: Springer-Verlag New York, Inc., 1994, pp. 3–12.
- [2] S. Tong and D. Koller, "Support vector machine active learning with applications to text classification," *Journal of Machine Learning Research (JMLR)*, vol. 2, pp. 45–66, 2001.
- [3] C. Schröder and A. Niekler, "A survey of active learning for text classification using deep neural networks," *arXiv preprint arXiv:2008.07267*, 2020.
- [4] Y.-F. Yan, S.-J. Huang, S. Chen, M. Liao, and J. Xu, "Active learning with query generation for cost-effective text classification," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 04, 2020, pp. 6583–6590.
- [5] Z. Zhan, Z. Hou, Q. Yang, J. Zhao, Y. Zhang, and C. Hu, "Knowledge attention sandwich neural network for text classification," *Neurocomputing*, vol. 406, pp. 1–11, 2020.
- [6] C. A. Flores, R. L. Figueroa, and J. E. Pezoa, "Active learning for biomedical text classification based on automatically generated regular expressions," *IEEE Access*, vol. 9, pp. 38 767–38 777, 2021.

- [7] M. Bloodgood, "Support vector machine active learning algorithms with query-by-committee versus closest-to-hyperplane selection," in *Proceedings of the 2018 IEEE 12th International Conference on Semantic Computing (ICSC)*. Laguna Hills, CA, USA: IEEE, January 2018, pp. 148–155. [Online]. Available: <https://doi.org/10.1109/ICSC.2018.00029>
- [8] M. Bloodgood and K. Vijay-Shanker, "A method for stopping active learning based on stabilizing predictions and the need for user-adjustable stopping," in *Proceedings of the Thirteenth Conference on Computational Natural Language Learning (CoNLL-2009)*. Boulder, Colorado: Association for Computational Linguistics, June 2009, pp. 39–47. [Online]. Available: <http://www.aclweb.org/anthology/W09-1107>
- [9] T. Cai, Y. Zhou, and H. Zheng, "Cost-quality adaptive active learning for chinese clinical named entity recognition," in *2020 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*. IEEE, 2020, pp. 528–533.
- [10] M. Bloodgood and C. Callison-Burch, "Bucking the trend: Large-scale cost-focused active learning for statistical machine translation," in *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*. Uppsala, Sweden: Association for Computational Linguistics, July 2010, pp. 854–864. [Online]. Available: <http://www.aclweb.org/anthology/P10-1088>
- [11] A. Miura, G. Neubig, M. Paul, and S. Nakamura, "Selecting syntactic, non-redundant segments in active learning for machine translation," in *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. San Diego, California: Association for Computational Linguistics, June 2016, pp. 20–29. [Online]. Available: <http://www.aclweb.org/anthology/N16-1003>
- [12] R. Hazra, P. Dutta, S. Gupta, M. A. Qaathir, and A. Dukkipati, "Active² learning: Actively reducing redundancies in active learning methods for sequence tagging and machine translation," in *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Online: Association for Computational Linguistics, Jun. 2021, pp. 1982–1995. [Online]. Available: <https://www.aclweb.org/anthology/2021.naacl-main.159>
- [13] H. Hino, "Active learning: Problem settings and recent developments," *arXiv preprint arXiv:2012.04225*, 2020.
- [14] G. Beatty, E. Kochis, and M. Bloodgood, "The use of unlabeled data versus labeled data for stopping active learning for text classification," in *Proceedings of the 2019 IEEE 13th International Conference on Semantic Computing (ICSC)*. Newport Beach, CA, USA: IEEE, January 2019, pp. 287–294. [Online]. Available: <https://doi.org/10.1109/ICOSC.2019.8665546>
- [15] M. Bloodgood and J. Grothendieck, "Analysis of stopping active learning based on stabilizing predictions," in *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*. Sofia, Bulgaria: Association for Computational Linguistics, August 2013, pp. 10–19. [Online]. Available: <http://www.aclweb.org/anthology/W13-3502>
- [16] M. Altschuler and M. Bloodgood, "Stopping active learning based on predicted change of f measure for text classification," in *Proceedings of the 2019 IEEE 13th International Conference on Semantic Computing (ICSC)*. Newport Beach, CA, USA: IEEE, January 2019, pp. 47–54. [Online]. Available: <https://doi.org/10.1109/ICOSC.2019.8665646>
- [17] A. Vlachos, "A stopping criterion for active learning," *Computer Speech and Language*, vol. 22, no. 3, pp. 295–312, 2008.
- [18] G. McDonald, C. Macdonald, and I. Ounis, "Active learning stopping strategies for technology-assisted sensitivity review," in *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, ser. SIGIR '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 20532056. [Online]. Available: <https://doi.org/10.1145/3397271.3401267>
- [19] T. K. Saha, M. Al Hasan, C. Burgess, M. A. Habib, and J. Johnson, "Batch-mode active learning for technology-assisted review," in *2015 IEEE International Conference on Big Data (Big Data)*. IEEE, 2015, pp. 1134–1143.
- [20] A. Burka and K. J. Kuchenbecker, "How much haptic surface data is enough?" in *2017 AAAI Spring Symposium Series*, 2017.
- [21] G. Wiedemann, "Proportional classification revisited: Automatic content analysis of political manifestos using active learning," *Social Science Computer Review*, vol. 37, no. 2, pp. 135–159, 2019. [Online]. Available: <https://doi.org/10.1177/0894439318758389>
- [22] K. Shmilovich, R. A. Mansbach, H. Sidky, O. E. Dunne, S. S. Panda, J. D. Tovar, and A. L. Ferguson, "Discovery of self-assembling π -conjugated peptides by active learning-directed coarse-grained molecular simulation," *The Journal of Physical Chemistry B*, vol. 124, no. 19, pp. 3873–3891, 2020.
- [23] P. Ren, Y. Xiao, X. Chang, P.-Y. Huang, Z. Li, B. B. Gupta, X. Chen, and X. Wang, "A survey of deep active learning," *ACM Computing Surveys (CSUR)*, vol. 54, no. 9, pp. 1–40, 2021.
- [24] Z. Pullar-Strecker, K. Dost, E. Frank, and J. Wicker, "Hitting the target: Stopping active learning at the cost-based optimum," *arXiv preprint arXiv:2110.03802*, 2021.
- [25] J. Cohen, "A coefficient of agreement for nominal scales," *Educational and psychological measurement*, vol. 20, no. 1, pp. 37–46, 1960.
- [26] M. Bloodgood and K. Vijay-Shanker, "Taking into account the differences between actively and passively acquired data: The case of active learning with support vector machines for imbalanced datasets," in *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics, Companion Volume: Short Papers*. Boulder, Colorado: Association for Computational Linguistics, June 2009, pp. 137–140. [Online]. Available: <http://www.aclweb.org/anthology/N/N09/N09-2035.pdf>
- [27] T. Joachims, "Text categorization with support vector machines: Learning with many relevant features," in *ECML*, ser. Lecture Notes in Computer Science, C. Nedellec and C. Rouveirol, Eds., vol. 1398. Springer, 1998, pp. 137–142.
- [28] S. Dumais, J. Platt, D. Heckerman, and M. Sahami, "Inductive learning algorithms and representations for text categorization," in *CIKM '98: Proceedings of the seventh international conference on Information and knowledge management*, Bethesda, Maryland, United States, 1998, pp. 148–155.
- [29] J. Zhu, H. Wang, and E. Hovy, "Learning a stopping criterion for active learning for word sense disambiguation and text classification," in *In Proceedings of the International Joint Conference on Natural Language Processing (IJCNLP)*, 2008, pp. 366–372.
- [30] —, "Multi-criteria-based strategy to stop active learning for data annotation," in *Proceedings of the 22nd International Conference on Computational Linguistics (Coling 2008)*, Manchester, UK, August 2008, pp. 1129–1136. [Online]. Available: <http://www.aclweb.org/anthology/C08-1142>
- [31] A. McCallum and K. Nigam, "A comparison of event models for naive bayes text classification," in *Proceedings of AAAI-98, Workshop on Learning for Text Categorization*, 1998. [Online]. Available: <http://www.kamalnigam.com/papers/multinomial-aaaiws98.pdf>
- [32] D. D. Lewis, Y. Yang, T. Russell-Rose, and F. Li, "Rcv1: A new benchmark collection for text categorization research," *Journal of machine learning research*, vol. 5, no. Apr, pp. 361–397, 2004.
- [33] G. V. Cormack and M. R. Grossman, "Autonomy and reliability of continuous active learning for technology-assisted review," *ArXiv*, vol. abs/1504.06868, 2015.
- [34] —, "Multi-faceted recall of continuous active learning for technology-assisted review," in *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2015, pp. 763–766.
- [35] —, "Engineering quality and reliability in technology-assisted review," *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, 2016.
- [36] E. Yang, D. D. Lewis, and O. Frieder, "A regularization approach to combining keywords and training data in technology-assisted review," in *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law*, 2019, pp. 153–162.
- [37] —, "On minimizing cost in legal document review workflows," *arXiv preprint arXiv:2106.09866*, 2021.
- [38] E. Yang, S. MacAvaney, D. D. Lewis, and O. Frieder, "Goldilocks: Just-right tuning of bert for technology-assisted review," *arXiv preprint arXiv:2105.01044*, 2021.
- [39] G. Beatty, E. Kochis, and M. Bloodgood, "Impact of batch size on stopping active learning for text classification," in *Proceedings of the 2018 IEEE 12th International Conference on Semantic Computing (ICSC)*. Laguna Hills, CA, USA: IEEE, January 2018, pp. 306–307. [Online]. Available: <https://doi.org/10.1109/ICSC.2018.00059>