

Stopping Active Learning based on Predicted Change of F Measure for Text Classification

Michael Altschuler
Department of Computer Science
The College of New Jersey
Ewing, New Jersey 08618
Email: altschm1@tcnj.edu

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

Abstract—During active learning, an effective stopping method allows users to limit the number of annotations, which is cost effective. In this paper, a new stopping method called Predicted Change of F Measure will be introduced that attempts to provide the users an estimate of how much performance of the model is changing at each iteration. This stopping method can be applied with any base learner. This method is useful for reducing the data annotation bottleneck encountered when building text classification systems.

I. INTRODUCTION

The use of active learning to train machine learning models has been used as a way to reduce annotation costs for text and speech processing applications [1], [2], [3], [4], [5]. Active learning has been shown to have a particularly large potential for reducing annotation cost for text classification [6], [7]. Text classification is one of the most important fields in semantic computing and it has been used in many applications [8], [9], [10], [11], [12].

A. Active Learning

Active learning is a form of machine learning that gives the model the ability to select the data on which it wants to learn from and to choose when to end the process of training.

In active learning, the model is first provided a small batch of annotated data to be trained on. Then, in each following iteration, the model selects a small batch and removes this batch from a large unlabeled set of examples. This batch is then annotated and added to the labeled set of examples. Models are trained at each iteration on the growing set of labeled data. By carefully selecting the batches of data to label at each iteration, learning efficiency can be increased. However, in order to enable the potential benefits enabled by active learning, a method for stopping the process must be developed.

Active learning is most applicable in fields where there is a lot of unlabeled data readily available, but the cost of annotations is very high. Ideally, the model is able to maximize its performance on the least amount of data to minimize the annotation cost with active learning. The model needs to correctly identify the batch that will increase the performance of the model by the most and correctly identify the optimal time to end the active learning process.

In previous papers, there has been extensive research on what is the optimal batch to select. Lewis and Gale proposed a method called uncertainty sampling where examples are chosen where the model is least certain of classification [6]. Schohn and Cohn have developed an uncertainty sampling method for support vector machines, which can be called closest-to-the-hyperplane, where examples for the batch are chosen based on the distance from the hyperplane [13]. Tang, Luo, and Roukos developed an uncertainty sampling method based on the entropy formula [14]. Another technique, called query by committee, involves selecting the examples based on how much a set of models disagree on the output of an example [15]. However, closest-to-hyperplane selection has been found to be superior to query-by-committee [16]; hence, we use closest-to-hyperplane selection in our experiments. Konyushkova, Raphael, and Fua propose a sampling method that selects the batch that reduces the expected error of a model [17]. The focus of this paper will be on the stopping method.

B. Stopping Method

A stopping method essentially tells the model when to end the process of active learning. Using the terminology introduced in [18], a stopping method is considered *aggressive* if a stopping method prioritizes minimizing annotation costs over achieving maximum performance. Hence an aggressive stopping method will stop relatively earlier. A stopping method is considered *conservative* if a stopping method prioritizes maximizing performance over minimizing annotation costs. Hence a conservative method will stop relatively later.

The benefits of an effective stopping method are two-fold. First, an effective stopping method can limit annotation cost ideally without compromising performance. This can be economically beneficial in many applications. Second, there are reported instances where a model trained on some data is actually superior to a model trained on all the data [13].

Evaluating a stopping criterion can often be a subjective issue that is determined by many factors. This is because users may prefer a conservative or aggressive stopping method based on the application. In a cost-sensitive project where performance maximization isn't critical, an aggressive stopping method may be preferred. In a domain where performance is the priority, a conservative stopping method may be preferred.

II. RELATED WORK

Various stopping methods have been proposed over the years. In 2000, Schohn and Cohn recommended to end the active learning process once all the examples within the margins of an SVM model have been labeled [13]. The reasoning behind this stopping method was that the examples most likely to change the hyperplane of a support vector machine lie within the margins. Once the margin has been exhausted, all the useful examples have been labeled and active learning can be stopped.

In 2007 and 2008, Zhu, Wang, and Hovy proposed several stopping methods. In 2007, Zhu and Hovy proposed two stopping methods (max-confidence and min-error) [19]. Max-confidence stops when the entropy on each selected unlabeled example is less than a threshold. Min-error stops when the model's predictions on the unlabeled batch exceeds a certain performance. In 2008, Zhu, Wang, and Hovy proposed a new strategy called the minimum expected error strategy that stops when the total expected error on the unlabeled pool is less than a certain threshold [20]. Zhu, Wang, and Hovy also proposed a combination of the previous strategies with the classification change stopping method which stops when no change of classification of an unlabeled example happens on consecutive iterations [21].

In 2008, Laws and Schutze proposed two stopping methods called performance convergence and uncertainty convergence [22]. Performance convergence stops when the estimation of accuracy of the unlabeled pool converges so that the gradient of performance estimates is below a certain threshold. Uncertainty convergence stops when the uncertainty of the last selected instance of the batch has an uncertainty measure below a minimum threshold. For support vector machines and other learners that do not provide probabilities of classification, only uncertainty convergence can be used.

In 2008, Vlachos proposed a stopping criterion that stops when the confidence on a set of examples consistently drops [23]. The logic in this stopping method is that as you add more labeled data to the model, the confidence of the model will initially rise, but then decline after the performance stabilized. However, in experiments, it has been found that this assumption does not always hold.

In 2009, Bloodgood and Vijay-Shanker proposed a stopping method called stabilizing predictions [18]. This stopping method takes a set of data that does not have to be labeled called a stop set and classifies each example in the stop set as positive or negative at each iteration of active learning. If the classifications of the examples in the stop set begin to stabilize so that the agreement between the classifications of successively learned models across several iterations is above a threshold as measured by Cohen's Kappa [24], then the active learning process ends. This stopping method is easy-to-implement, applicable with any base learner, and the number of iterations and the Kappa threshold can be altered depending on how conservative or aggressive one chooses to be.

In 2010, Ghayoomi proposed a stopping criterion called the

extended variance model [25]. The extended variance model stops once the variance of the confidence of the unlabeled pool decreases by a minimum threshold over a certain number of iterations. However, in our experiments with text classification (see Section IV-C), this stopping method does not stop.

In 2013, Bloodgood and Grothendieck were able to prove that when Cohen's Kappa is greater than or equal to some threshold T , then the change of F measure is less than or equal to $\frac{4(1-T)}{T}$ [26]. This means that stabilizing predictions provides an intuition of how much performance will change on the stop set. Most recently, Beatty and others [27] were able to show that stopping methods based on unlabeled data such as the stabilizing predictions method of [18], [26] are more effective than stopping methods based on labeled data.

III. PREDICTED CHANGE OF F

As stated earlier, determining how effective a stopping method is can be subjective as one does not know how aggressive or conservative the user wants to be. Therefore, an ideal stopping method should accomplish the following goals:

- 1) The user should have the ability to determine how aggressive or conservative the stopping method should be based on a given set of parameters.
- 2) The stopping method should provide a sense of intuition on how much performance was gained or lost on the previous iteration so that the user can effectively determine if the marginal benefit of the change in performance will outweigh the cost of annotating another batch.
- 3) The stopping method should be able to work with various learners.

Hypothetically, one could calculate the change in performance by creating a separate test set of examples with the labels provided called the validation set and calculate the change in performance from this set. However, this idea brings about two problems. First, if the validation set is really small, there are no guarantees that this validation set will be an accurate representation of the population. Second, if the validation set is really large, the annotation costs for providing the labels will be really high which defeats the whole purpose of active learning. Therefore, another method is preferred that can determine when to stop without requiring any additional labeled data.

In 2009, Bloodgood and Vijay-Shanker proposed a stopping method that stops when the predictions at different iterations stabilize [18]. This method uses a set of data that does not need to be labeled called the stop set. At each iteration, the model classifies each example in the stop set as positive or negative. After three consecutive iterations, if the measurement of agreement between the three iterations as measured by the Kappa value is above a threshold, active learning ends. In other words, once the predictions on the stop set stabilize, no more examples are annotated. The advantage of this approach is that it is applicable to any base learner, offers the user the flexibility to adjust the threshold, and was shown to have stable strong performance over many datasets.

TABLE I

A CONTINGENCY TABLE OF LABELED EXAMPLES AT THE MOST RECENT ITERATION AND THE PREVIOUS ITERATION.

		M_t		Total
		+	-	
M_{t-1}	+	a	b	a + b
	-	c	d	c + d
Total		a + c	b + d	n

TABLE II

A CONTINGENCY TABLE OF LABELED EXAMPLES THAT ARE TRULY POSITIVE AT THE MOST RECENT ITERATION AND THE PREVIOUS ITERATION.

		M_t		Total
		+	-	
M_{t-1}	+	a_+	b_+	$a_+ + b_+$
	-	c_+	d_+	$c_+ + d_+$
Total		$a_+ + c_+$	$b_+ + d_+$	n_+

In 2013, Bloodgood and Grothendieck proved that when the Kappa value is greater than or equal to any threshold T , then the change of F measure is less than or equal to $\frac{4(1-T)}{T}$ [26]. This theorem provides an intuition of how much performance will change on the stop set, which is valuable information to the user. F measure was used over accuracy because F measure is commonly used in text classification and other information extraction tasks. F measure is defined as the harmonic mean of precision and recall. Precision is the percentage of system-proposed positives that are truly positive. Recall is the percentage of truly positive examples that the system proposed as positive.

The result from [26] only provides an upper bound. In practice, the difference in performance could be much smaller. It would be an advantage to the user to provide a more accurate estimate of the change of F measure. To do this, a more direct conversion to change of performance is needed.

In 2013, Bloodgood and Grothendieck described a contingency table that compares the predictions made by M_t and M_{t-1} where M_t is the model that was learned at the most recent iteration and M_{t-1} is the model that was learned at the previous iteration[26]. As seen in Table I, when M_t and M_{t-1} both label the example as positive, that example is added to the count of a . When M_t and M_{t-1} both label the example as negative, that example is added to the count of d . When M_t labels the example as positive and M_{t-1} labels the example as negative, that example is added to the count of c . When M_t labels the example as negative and M_{t-1} labels the example as positive, that example is added to the count of b .

However, we only know the system-proposed labels of examples in the stop set. Therefore, Bloodgood and Grothendieck described the following two tables (Table II and Table III) which show the truly positive and negative examples of a , b , c , and d . Let a_+ be the number of examples contributing to the count of a that are truly positives and let a_- be the examples contributing to the count of a that are truly negative. Let the same apply for b , c , and d .

The contingency tables of M_t versus ground truth and

TABLE III

A CONTINGENCY TABLE OF LABELED EXAMPLES THAT ARE TRULY NEGATIVE AT THE MOST RECENT ITERATION AND THE PREVIOUS ITERATION.

		M_t		Total
		+	-	
M_{t-1}	+	a_-	b_-	$a_- + b_-$
	-	c_-	d_-	$c_- + d_-$
Total		$a_- + c_-$	$b_- + d_-$	n_-

TABLE IV

A CONTINGENCY TABLE OF LABELED EXAMPLES AT THE MOST RECENT ITERATION VERSUS THE GROUND TRUTH.

		M_t		Total
		+	-	
<i>Truth</i>	+	$a_+ + c_+$	$b_+ + d_+$	n_+
	-	$a_- + c_-$	$b_- + d_-$	n_-
Total		$a + c$	$b + d$	n

M_{t-1} versus ground truth are required if the true change of F performance needs to be calculated. Table IV shows M_t versus ground truth and Table V shows M_{t-1} versus ground truth.

To calculate the F measure of the model from the contingency counts, the formula is:

$$F = \frac{2a}{2a + b + c} \quad (1)$$

The formula for the change in F measure in terms of a_+ , a_- , b_+ , b_- , c_+ , c_- , d_+ , and d_- is:

$$\Delta F = \frac{2(a_+ + c_+)}{2(a_+ + c_+) + b_+ + d_+ + a_- + c_-} - \frac{2(a_+ + b_+)}{2(a_+ + b_+) + c_+ + d_+ + a_- + b_-} \quad (2)$$

However, one does not know the true values of a_+ , a_- , b_+ , b_- , c_+ , c_- , d_+ , and d_- . Therefore, a method is needed to estimate the values of a_+ , a_- , b_+ , b_- , c_+ , c_- , d_+ , and d_- . For estimating the values, the following assumption will be made: the model at each iteration will get progressively better as it is trained on more labeled data. Therefore, on all cases where M_{t-1} and M_t disagree, we will assume that M_t is correct. In other words, we assume $a_+ = a$, $a_- = 0$, $b_+ = 0$, $b_- = b$, $c_+ = c$, $c_- = 0$, $d_+ = 0$, and $d_- = d$. In Section IV-B, we evaluate the effectiveness of this assumption.

The predicted change of F can now be described as:

$$\Delta \hat{F} = \frac{2(a + c)}{2(a + c)} - \frac{2a}{2a + c + b} = 1 - \frac{2a}{2a + b + c} \quad (3)$$

When $\Delta \hat{F} < \epsilon$ where ϵ is a predefined threshold over k windows, the active learning process stops. Or in other words:

$$\text{StoppingMethod} = \begin{cases} \text{Stop,} & \text{if } \Delta \hat{F} < \epsilon \\ & \text{over } k \text{ windows} \\ \text{Continue,} & \text{otherwise.} \end{cases}$$

TABLE V
A CONTINGENCY TABLE OF LABELED EXAMPLES AT THE PREVIOUS
ITERATION VERSUS THE GROUND TRUTH.

		M_{t-1}		Total
		+	-	
$Truth$	+	$a_+ + b_+$	$c_+ + d_+$	n_+
	-	$a_- + b_-$	$c_- + d_-$	n_-
Total		$a + b$	$c + d$	n

This stopping method fulfills all three goals of an effective stopping method. First, the user has the ability to determine how aggressive or conservative the stopping method should be based on the parameters k and ϵ . Second, the stopping method provides an estimate of how much the performance of the model will change at each iteration by $\Delta \hat{F}$. These experiments will be shown in Section IV-C. Third, the stopping method is able to work with any learner as only the predictions on a stop set of data are needed.

IV. EXPERIMENTS

Four experiments will be run regarding predicted change of F:

- 1) First, we will do an analysis of the assumption that M_t is more accurate than M_{t-1} by looking at what proportion of a is truly positive (ideally approximately 1), what proportion of b is truly positive (ideally approximately 0), what proportion of c is truly positive (ideally approximately 1), and what proportion of d is truly positive (ideally approximately 0).
- 2) Second, the predicted change of F measure will be compared to the change of performance on a separate labeled test set. In this way, the accuracy of this new measurement can be evaluated and will also be compared to the threshold defined by the Kappa variable. This will also be compared to the upper bound proven by Bloodgood and Grothendieck [26].
- 3) Third, predicted change of F will be compared to previous stopping methods to evaluate how aggressive/conservative the stopping method is.
- 4) Fourth, an analysis on the parameters will be done to investigate how the change in ϵ and k will change stopping behavior.

A. Experimental Setup

Predicted change of F was evaluated in the domain of binary text classification. The datasets used were 20Newsgroups, Reuters, and WebKB. 20Newsgroups contains 20 different categories of articles such as atheism, space, or baseball. The average for 20Newsgroups was done by averaging the 20 categories. For Reuters, 10 categories were used that included topics such as corn, money, and wheat. This is consistent with past practices [28], [29]. The average was calculated by averaging the ten categories. The four categories, students, faculty, project, and course, were used for WebKB. This is also consistent with past practices [30], [21], [20]. Ten folds

were used, and the average was calculated for each category. All datasets were downloaded July 13, 2017.

The base learner used for these experiments was a support vector machine. A support vector machine defines a hyperplane that divides examples into two categories (+1 or -1). The decision function is defined in terms of $f(x) = \text{sgn}(w * x + b)$. Support vector machines were used because they have good performance in the field of text classification and a lot of past analysis of stopping methods has been done with support vector machines.

The dataset from 20Newsgroup includes 11314 articles of 20 different categories. The dataset from Reuters contains 7780 articles of 10 different categories. The dataset from WebKB contains 7445 articles of four different categories. We used a batch size of 0.5% of the size of the initial unlabeled pool and a stop set size of 50% of the size of the initial unlabeled pool for all the datasets. A larger batch size can increase computational tractability, however, it also degrades the effectiveness of stopping methods [31]. This is why we used a batch size of 0.5% of the initial unlabeled pool in all our experiments.

In our experiments, the selection algorithm selects the batch of examples that are closest to the hyperplane [32], [7], [33], [13]. Recently this selection algorithm was found to be superior to query by committee selection [16]. All parameters used for other stopping methods were as recommended in previous papers [13], [21], [22], [23], [18], [25], [26].

B. Analysis of Assumption

In Section III, the assumption was made that in areas of disagreement, the more recent iteration will generally be superior because it is trained on more data. Therefore, $a_+ = a$, $a_- = 0$, $b_+ = 0$, $b_- = b$, $c_+ = c$, $c_- = 0$, $d_+ = 0$, and $d_- = d$. However, it is fair to assume that not all of the examples labeled by consecutive iterations as positive (a) will have a ground-truth value of positive. Likewise, this holds true for b , c , and d .

Therefore, the goal of this experiment is to find out what proportion of examples labeled a are truly positive, what proportion of examples labeled b are truly positive, what proportion of examples labeled c are truly positive, and what proportion of examples labeled d are truly positive.

Let $P(+|a)$ be the conditional probability that an example has a ground truth value of positive given an example is labeled by M_t and M_{t-1} as positive. Let $P(+|b)$ be the conditional probability that an example has a ground truth value of positive given an example is labeled by M_t as negative and by M_{t-1} as positive. Let $P(+|c)$ be the conditional probability that an example has a ground truth value of positive given an example is labeled by M_t as positive and by M_{t-1} as negative. Let $P(+|d)$ be the conditional probability that an example has a ground truth value of positive given an example is labeled by M_t and M_{t-1} as negative.

If the assumptions were perfect, $P(+|a)$ would be 1; $P(+|b)$ would be 0; $P(+|c)$ would be 1; $P(+|d)$ would be 0.

TABLE VI

THIS TABLE SHOWS THE PROPORTION OF EXAMPLES LABELED a THAT ARE TRULY POSITIVE, THE PROPORTION OF EXAMPLES LABELED b THAT ARE TRULY POSITIVE, THE PROPORTION OF EXAMPLES LABELED c THAT ARE TRULY POSITIVE, AND THE PROPORTION OF EXAMPLES LABELED d THAT ARE TRULY POSITIVE FOR ALL DATASETS. EACH CONDITIONAL PROBABILITY IS CALCULATED AT DIFFERENT STOPPING POINTS FOR WHEN $\epsilon = 0.01$, $\epsilon = 0.05$ AND $\epsilon = 0.1$.

Dataset	ϵ	$P(+ a)$	$P(+ b)$	$P(+ c)$	$P(+ d)$
20Newsgroups	0.01	99.4%	3.70%	97.3%	0.497%
Reuters	0.01	89.0%	20.9%	64.9%	0.354%
WebKB - faculty	0.01	96.4%	12.7%	88.8%	1.17%
WebKB - project	0.01	98.7%	0.00%	100%	1.38%
WebKB - student	0.01	93.5%	17.5%	83.0%	2.25%
WebKB - course	0.01	95.3%	13.0%	93.0%	1.07%
20Newsgroups	0.05	98.0%	15.0%	74.9%	0.808%
Reuters	0.05	86.9%	29.4%	48.0%	1.01%
WebKB - faculty	0.05	91.9%	27.4%	58.2%	1.84%
WebKB - project	0.05	95.1%	9.45%	73.8%	2.11%
WebKB - student	0.05	87.4%	30.8%	52.6%	3.59%
WebKB - course	0.05	91.3%	32.0%	60.7%	1.73%
20Newsgroups	0.10	96.8%	17.3%	69.6%	1.15%
Reuters	0.10	86.2%	19.1%	61.7%	1.23%
WebKB - faculty	0.10	89.0%	27.4%	58.1%	2.42%
WebKB - project	0.10	91.4%	22.7%	62.5%	2.61%
WebKB - student	0.10	86.0%	39.1%	55.1%	5.01%
WebKB - course	0.10	89.1%	34.7%	60.1%	2.37%

TABLE VII

THIS TABLE SHOWS THE ABSOLUTE VALUE OF THE DIFFERENCE BETWEEN THE PREDICTED CHANGE OF F AND THE ACTUAL CHANGE OF F ON THE TEST SET. FOR WEBKB, THE AVERAGE OF THE 10 FOLDS WAS USED.

Dataset	Iterations							
	2	3	4	5	6	7	8	9
20Newsgroups	0.649	0.219	0.134	0.112	0.089	0.072	0.051	0.041
Reuters	0.249	0.228	0.182	0.117	0.095	0.061	0.047	0.022
WebKB - course	0.373	0.176	0.166	0.098	0.084	0.067	0.055	0.049
WebKB - faculty	0.325	0.224	0.155	0.109	0.069	0.086	0.054	0.039
WebKB - project	0.709	0.489	0.342	0.282	0.184	0.133	0.106	0.087
WebKB - student	0.407	0.220	0.148	0.140	0.111	0.088	0.074	0.090

The values of $P(+|a)$, $P(+|b)$, $P(+|c)$, and $P(+|d)$ at different stopping points for each dataset are shown in Table VI.

Several observations can be drawn from this experiment. First, when ϵ is 0.01, the assumptions are very accurate. However, as ϵ increases, the assumptions become slightly less accurate with regards to $P(+|b)$ and $P(+|c)$. However, these results are overall consistent with the assumption that $a_+ = a$, $a_- = 0$, $b_+ = 0$, $b_- = b$, $c_+ = c$, $c_- = 0$, $d_+ = 0$, and $d_- = d$.

C. Accuracy of Predicting Change of F Measure

The goal of the experiment reported in this section is to see how well predicted change of F predicts the true change of F on the test set. The metric used to evaluate this is the absolute value of the difference between the predicted change of F and the change of F measure on the test set. More formally, this can be referred to as: $D = |\text{predicted change of F} - \text{change of F on test set}|$. Predicted change of F was compared to the upper bound of change of F proven in [26]. The two tables are shown as Table VII and Table VIII.

There are a few observations that can be made from looking at the results. First, when there were few annotations labeled, the predicted change of F was inaccurate. However, as the

TABLE VIII

THIS TABLE SHOWS THE ABSOLUTE DIFFERENCE BETWEEN THE UPPER BOUND FOR CHANGE OF F ON THE STOP SET AS PROVEN IN [26] AND THE TRUE CHANGE OF F ON THE VALIDATION SET. FOR WEBKB, THE AVERAGE OF THE 10 FOLDS WAS USED.

Dataset	Iterations							
	2	3	4	5	6	7	8	9
20Newsgroups	1.00	1.00	1.00	0.851	0.578	0.424	0.299	0.193
Reuters	1.00	1.00	1.00	1.00	1.00	0.513	0.272	0.139
WebKB - course	1.00	1.00	1.00	0.655	0.527	0.392	0.297	0.243
WebKB - faculty	1.00	1.00	1.00	0.806	0.464	0.420	0.315	0.243
WebKB - project	1.00	1.00	1.00	1.00	1.00	0.932	0.719	0.586
WebKB - student	1.00	1.00	1.00	1.00	0.748	0.586	0.493	0.510

number of annotations increased, the accuracy of the predicted change of F became more accurate. As the predicted change of F converges to 0, the more accurate the method predicts the true change of F. Second, the predicted change of F always provided a more accurate prediction of change of F than if one were to use the upper bound of the change of F as an estimate of the change of F at every iteration on all datasets.

D. Analysis of Aggressiveness of Stopping Method

The goal of the experiment reported in this section is to see how our new stopping method compares with previous stopping methods.

We define the optimal ordered pair ($optimalA$, $optimalP$) where $optimalP$ is the maximum performance over the entire learning process minus ϵ_2 . The purpose of ϵ_2 was so that a small increase at the end of the training wouldn't count that iteration as an ideal stopping point. For all our experiments, $\epsilon_2 = 0.01$.

$optimalA$ is the number of annotations at the earliest iteration that achieves a minimum performance of $optimalP$. The metrics used to evaluate the aggressiveness of a stopping method are:

Percentage of performance achieved:

$$P = \frac{\text{performance of model at stopping point}}{optimalP}$$

and

Percentage of annotations used:

$$A = \frac{\text{number of annotations used at stopping point}}{optimalA}$$

The predicted change of F with $\epsilon = 0.005$ and $k = 1$ (PCF[0.005]) was compared to margin exhaustion (ME) [13], drop in confidence (DC) [23], confidence convergence (CC) [22], min-error with classification change (MECC) [21], stabilizing predictions (SP) [18], and the extended variance model (EVM) [25]. Note that all parameters for all methods were kept the same as conducted in the experiments of Bloodgood and Vijay-Shanker [18]. Table IX and Table X show these results.

When looking at the percentage of performance achieved, predicted change of F typically scored very high and achieved at least 90% of optimal performance on most datasets. When looking at each dataset individually, predicted change of F achieved at least 95% of optimal performance on all datasets

TABLE IX

THIS TABLE SHOWS THE AVERAGE PERCENTAGE OF *optimalP* ACHIEVED AMONG THE DIFFERENT STOPPING METHODS. FOR 20 NEWSGROUPS AND REUTERS, THE AVERAGE AMONG ALL THE CATEGORIES WERE USED. FOR WEBKB, THE AVERAGE OF THE 10 FOLDS WAS USED.

	ME	DC	CC	MECC	SP	EVM	PCF (0.005)
20Newsgroups	0.982	0.586	0.879	0.973	0.976	0.987	0.973
Reuters	0.969	0.974	0.843	0.971	0.955	0.967	0.958
WebKB - course	0.979	0.872	0.886	0.982	0.968	0.943	0.975
WebKB - faculty	0.972	0.895	0.923	0.973	0.975	0.954	0.975
WebKB - project	0.957	0.556	0.672	0.725	0.891	0.911	0.906
WebKB - student	0.980	0.877	0.831	0.991	0.978	0.971	0.976

TABLE X

THIS TABLE SHOWS THE AVERAGE PERCENTAGE OF *optimalA* ACHIEVED AMONG THE DIFFERENT STOPPING METHODS. FOR 20NEWSGROUPS AND REUTERS, THE AVERAGE AMONG ALL THE CATEGORIES WERE USED. FOR WEBKB, THE AVERAGE OF THE 10 FOLDS WAS USED.

	ME	DC	CC	MECC	SP	EVM	PCF - 0.005
20Newsgroups	2.05	0.32	0.81	0.90	0.98	13.36	0.91
Reuters	1.70	6.17	0.54	1.02	1.05	14.96	1.00
WebKB - course	2.51	11.94	0.46	1.65	1.46	15.20	1.36
WebKB - faculty	2.18	3.81	0.47	1.67	1.32	12.82	1.24
WebKB - project	1.07	0.79	0.29	0.54	0.72	5.77	0.67
WebKB - student	1.91	3.24	0.22	1.36	1.06	7.17	1.05

except WebKB - projects. It is important to note that most of the other stopping methods also struggled to achieve optimal performance on this dataset and category.

When looking at the percentage of annotations used, predicted change of F typically achieved a score very close to 1.00. This means that predicted change of F was relatively aggressive. It is also important to note that the parameters $\epsilon = 0.005$ and $k = 1$ can be adjusted to be more or less aggressive depending on the application. In the next subsection we investigate the impacts that the parameters have on stopping behavior.

E. Analysis of Parameters

One of the goals of this section is to analyze and see how a change in ϵ and the number of windows k will impact the aggressiveness of the stopping method. The intuition is that a decrease in ϵ and an increase in k would lead to a more conservative stopping method. However, to what degree would it be more conservative? To do this, the number of annotations at each stopping point was found for values of ϵ in the set $\{0.01, 0.02, 0.03, 0.04, 0.05, 0.06, 0.07, 0.08, \text{ and } 0.09\}$ and for different values of k in the set $\{1, 2, 3\}$. Then a regression was done with the model of: $\hat{y} = \beta_1\epsilon + \beta_2k + \beta_3$ where \hat{y} is the predicted number of annotations, ϵ is the epsilon parameter ranging from 0.01 to 0.09, k is the number of windows ranging from one to three, and β_1 , β_2 , and β_3 are associated parameters of the model. $\beta_1/100$ shows the predicted change in annotations when ϵ increases by 0.01. β_2 will show the predicted change in annotations when the number of windows increase by 1.

To conduct the experiment, the stopping point was found for each of the previously used categories of each dataset using different combinations of parameters with ϵ in the set $\{0.01, 0.02, 0.03, 0.04, 0.05, 0.06, 0.07, 0.08, \text{ and } 0.09\}$ and with k in the set $\{1, 2, 3\}$. For example, for the alt.atheism of 20Newsgroups, the number of annotations needed for stopping

TABLE XI

THIS TABLE SHOWS THE R^2 , $\beta_1/100$, β_2 , AND p -VALUES FOR β_1 AND β_2 FOR EACH DATASET DONE SEPARATELY AND ALL THE DATASETS TOGETHER. R^2 EXPLAINS WHAT PERCENTAGE OF THE VARIANCE OF ANNOTATIONS IS EXPLAINED FOR BY THE VARIATION OF THE INDEPENDENT VARIABLES OF ϵ AND k . $\beta_1/100$ SHOWS THE CHANGE IN ANNOTATIONS FOR EVERY INCREASE IN ϵ BY 0.01 FROM 0.01 TO 0.09. β_2 SHOWS THE CHANGE IN ANNOTATIONS FOR EVERY INCREASE IN THE NUMBER OF WINDOWS. LETTING $\alpha = 0.05$, IF THE p -VALUES FOR EACH β PARAMETER IS BELOW α , THEN THE β PARAMETER IS STATISTICALLY SIGNIFICANT.

Dataset	R^2	$\beta_1/100$	β_2	p_1 -val	p_2 -val
20News	44.7%	-421.36	680.81	6.90e-62	3.32e-12
Reuters	42.2%	-258.87	531.64	2.70e-26	6.15e-9
WebKB	55.1%	-727.32	645.61	6.69e-192	2.37e-10
ALL	44.5%	-572.98	639.39	1.46e-238	9.84e-21

to occur were recorded for $\epsilon = 0.01$ and $k = 1$, $\epsilon = 0.02$ and $k = 1$, $\epsilon = 0.03$ and $k = 1$ and so on until each combination of parameter values were used. Once the number of annotations for each category of each dataset was found, the data was trained on a linear regression model so that the independent variables were ϵ and k and the dependent variable was the number of annotations. The regressions were done for each dataset individually and then trained on all the datasets together. The batch size was also kept constant throughout the experiments at 0.05% of the training set.

The results of the experiment are shown in Table XI. Several conclusions can be drawn from Table XI. First, models that were made for specific datasets explained significantly more of the variation of annotations than when clustered all together. This shows evidence that the change of annotations is dependent on the dataset. Second, for every increase in ϵ by 0.01, the number of annotations decreased on average anywhere from 259 to 2218 depending on the dataset used. This was also statistically significant and confirms our intuition. When the window increased by one, the number of annotations on average needed increase from 531 to 1463. This too was also statistically significant for every dataset.

Another goal of this section was to see what type of relationship existed between the parameters and the difference between the actual change of F and the predicted change of F. In Section IV-C, it was said that as the predicted change of F decreases, the prediction becomes more accurate. In other words, as the method becomes more conservative, the predicted change of F measure becomes more accurate. This observation will be tested by fitting a model of the form: $\hat{y} = \beta_1\epsilon + \beta_2k + \beta_3$ where \hat{y} is the predicted difference between the actual change of F and the predicted change of F, ϵ is the epsilon parameter ranging from 0.01 to 0.09, k is the number of windows ranging from one to three, and β_1 , β_2 , and β_3 are associated parameters of the model. $\beta_1/100$ shows the predicted change in annotations when ϵ increases by 0.01. β_2 will show the predicted change in annotations when the number of windows increase by 1.

The data was compiled the same way as the previous regression, except the dependent variable is now the difference between the actual change of F and the predicted change of F

TABLE XII

THIS TABLE SHOWS THE R^2 , $\beta_1/100$, β_2 , β_3 , AND p -VALUES FOR β_1 AND β_2 FOR EACH DATASET DONE SEPARATELY AND ALL THE DATASETS TOGETHER. R^2 EXPLAINS WHAT PERCENTAGE OF THE VARIANCE OF ANNOTATIONS IS EXPLAINED FOR BY THE VARIATION OF THE INDEPENDENT VARIABLES OF ϵ AND THE NUMBER OF WINDOWS. $\beta_1/100$ SHOWS THE CHANGE IN THE DIFFERENCE OF THE PREDICTED AND THE ACTUAL CHANGE OF F MEASURE FOR EVERY INCREASE IN ϵ BY 0.01 FROM 0.01 TO 0.09. β_2 SHOWS THE CHANGE IN THE DIFFERENCE OF THE PREDICTED AND THE ACTUAL CHANGE OF F MEASURE FOR EVERY INCREASE IN THE NUMBER OF WINDOWS. LETTING $\alpha = 0.05$, IF THE p -VALUES FOR EACH β PARAMETER IS BELOW α , THEN THE β PARAMETER IS STATISTICALLY SIGNIFICANT.

Dataset	R^2	$\beta_1/100$	β_2	p_1 -val	p_2 -val
20News	51.4%	3.56e-3	-6.13e-3	1.21e-71	2.86e-15
Reuters	43.7%	3.86e-3	-5.35e-3	0.00	5.42e-12
WebKB	82.5%	6.99e-3	-5.46e-3	0.00	6.03e-12
ALL	56.0%	5.57e-3	-5.63e-3	0.00	9.25e-22

rather than the number of annotations.

The results of the experiment are shown in Table XII. There are several conclusions that can be drawn from the following table. First, there is a statistically significant relationship between ϵ and k and the difference between the predicted and actual change of F measure. Second, the variance in ϵ and k explains more of the variance of the difference between the actual and predicted change of F measure than the number of annotations. Third, it can be said then that making ϵ smaller or increasing the number of windows can lead to a more accurate prediction.

V. CONCLUSION

A new stopping method was introduced called predicted change of F. This stopping method provides the user the flexibility to adjust the parameters ϵ and k so that the stopping method can be more or less aggressive. This stopping method also provides an intuition of how much the performance of the model is changing. Finally, this stopping method can be used with any base learner.

The intuition behind such a method is that we want to use a stop set that does not need to be labeled to predict what the change of performance is expected to be. In order to predict the change of performance, we rely on an assumption that M_t will be superior to M_{t-1} because the most recent iteration would have been trained on a larger training set than the previous iteration. This assumption was found in Section IV-B to be reasonable. This assumption allows the predicted change of F to be evaluated as $\Delta \hat{F} = \frac{2(a+c)}{2(a+c)} - \frac{2a}{2a+c+b} = 1 - \frac{2a}{2a+c+b}$. As found in Section IV-C, the predicted change of F was found to be closer to the true change of F on the stop set than the upper bound proven by Bloodgood and Grothendieck [26].

We investigated the relationship between the parameters ϵ and k and the number of annotations along with the accuracy of the predicted change of F measure. Ultimately, both parameters were found to be statistically significant. Furthermore, as expected, when ϵ decreased and the number of windows k increased, the stopping method became more conservative and more annotations were required. Furthermore, as ϵ decreased

and k increased, the prediction of the change of F measure became more accurate.

The predicted change of F was found to be very accurate as the number of annotations increased. However, at earlier iterations, the predicted change of F measure was unreliable. This is not a problem when using predicted change of F as a stopping method, as the method only stops when the predicted change of F converges to a small number. The new stopping method was found to be effective due to the fact that it provides flexibility through its parameters, provides estimates of how much performance is changing, and can be applied with any base learner without any significant computational cost.

ACKNOWLEDGMENT

This work was supported in part by The College of New Jersey Support of Scholarly Activities (SOSA) program and by The College of New Jersey Mentored Undergraduate Summer Experience (MUSE) 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. The cluster is funded by the National Science Foundation under grant number OAC-1828163.

REFERENCES

- [1] S. Hantke, Z. Zhang, and B. Schuller, "Towards intelligent crowdsourcing for audio data annotation: Integrating active learning in the real world," *Proc. Interspeech*, pp. 3951–3955, 2017.
- [2] 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>
- [3] S.-W. Lee, D. Zhang, M. Li, M. Zhou, and H.-C. Rim, "Translation model size reduction for hierarchical phrase-based statistical machine translation," in *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Jeju Island, Korea: Association for Computational Linguistics, July 2012, pp. 291–295. [Online]. Available: <http://www.aclweb.org/anthology/P12-2057>
- [4] F. Mairesse, M. Gasic, F. Jurcicek, S. Keizer, B. Thomson, K. Yu, and S. Young, "Phrase-based statistical language generation using graphical models and active learning," in *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*. Uppsala, Sweden: Association for Computational Linguistics, July 2010, pp. 1552–1561. [Online]. Available: <http://www.aclweb.org/anthology/P10-1157>
- [5] 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>
- [6] D. D. Lewis and W. A. Gale, "A sequential algorithm for training text classifiers," in *Proceedings of the 17th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, ser. SIGIR '94. New York, NY, USA: Springer-Verlag New York, Inc., 1994, pp. 3–12. [Online]. Available: <http://dl.acm.org/citation.cfm?id=188490.188495>
- [7] S. Tong and E. Chang, "Support vector machine active learning for image retrieval," in *MULTIMEDIA '01: Proceedings of the Ninth ACM International Conference on Multimedia*. New York, NY, USA: ACM, 2001, pp. 107–118.

- [8] A. Mishler, K. Wonus, W. Chambers, and M. Bloodgood, "Filtering tweets for social unrest," in *Proceedings of the 2017 IEEE 11th International Conference on Semantic Computing (ICSC)*. San Diego, CA, USA: IEEE, January 2017, pp. 17–23. [Online]. Available: <https://doi.org/10.1109/ICSC.2017.75>
- [9] S. C. H. Hoi, R. Jin, and M. R. Lyu, "Large-scale text categorization by batch mode active learning," in *Proceedings of the 15th International Conference on World Wide Web*, ser. WWW '06. New York, NY, USA: ACM, 2006, pp. 633–642. [Online]. Available: <http://doi.acm.org/10.1145/1135777.1135870>
- [10] M. Janik and K. J. Kochut, "Wikipedia in action: Ontological knowledge in text categorization," in *2008 IEEE International Conference on Semantic Computing*. IEEE, 2008, pp. 268–275.
- [11] M. Allahyari, K. J. Kochut, and M. Janik, "Ontology-based text classification into dynamically defined topics," in *Semantic Computing (ICSC), 2014 IEEE International Conference on*. IEEE, 2014, pp. 273–278.
- [12] M. Kanakaraj and R. M. R. Guddeti, "Performance analysis of ensemble methods on twitter sentiment analysis using nlp techniques," in *Semantic Computing (ICSC), 2015 IEEE International Conference on*. IEEE, 2015, pp. 169–170.
- [13] G. Schohn and D. Cohn, "Less is more: Active learning with support vector machines," in *Proceedings of the 17th International Conference on Machine Learning (ICML)*, 2000, pp. 839–846.
- [14] M. Tang, X. Luo, and S. Roukos, "Active learning for statistical natural language parsing," in *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, ser. ACL '02. Stroudsburg, PA, USA: Association for Computational Linguistics, 2002, pp. 120–127. [Online]. Available: <https://doi.org/10.3115/1073083.1073105>
- [15] H. S. Seung, M. Opper, and H. Sompolinsky, "Query by committee," in *COLT '92: Proceedings of the Fifth Annual Workshop on Computational Learning Theory*. Pittsburgh, Pennsylvania, United States: ACM, 1992, pp. 287–294.
- [16] 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>
- [17] K. Konyushkova, R. Sznitman, and P. Fua, "Learning active learning from data," in *Advances in Neural Information Processing Systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds. Curran Associates, Inc., 2017, pp. 4225–4235. [Online]. Available: <http://papers.nips.cc/paper/7010-learning-active-learning-from-data.pdf>
- [18] 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>
- [19] J. Zhu and E. Hovy, "Active learning for word sense disambiguation with methods for addressing the class imbalance problem," in *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*. Prague, Czech Republic: Association for Computational Linguistics, June 2007, pp. 783–790. [Online]. Available: <http://www.aclweb.org/anthology/D/D07/D07-1082>
- [20] J. Zhu, H. Wang, and E. Hovy, "Learning a stopping criterion for active learning for word sense disambiguation and text classification," in *Proceedings of the Third International Joint Conference on Natural Language Processing: Volume-1*, 2008.
- [21] —, "Multi-criteria-based strategy to stop active learning for data annotation," in *Proceedings of the 22nd International Conference on Computational Linguistics - Volume 1*, ser. COLING '08. Stroudsburg, PA, USA: Association for Computational Linguistics, 2008, pp. 1129–1136. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1599081.1599223>
- [22] F. Laws and H. Schütze, "Stopping criteria for active learning of named entity recognition," in *Proceedings of the 22nd International Conference on Computational Linguistics - Volume 1*, ser. COLING '08. Stroudsburg, PA, USA: Association for Computational Linguistics, 2008, pp. 465–472. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1599081.1599140>
- [23] A. Vlachos, "A stopping criterion for active learning," *Computer Speech and Language*, vol. 22, no. 3, pp. 295–312, Jul. 2008. [Online]. Available: <http://dx.doi.org/10.1016/j.csl.2007.12.001>
- [24] J. Cohen, "A coefficient of agreement for nominal scales," *Educational and Psychological Measurement*, vol. 20, pp. 37–46, 1960.
- [25] M. Ghayoomi, "Using variance as a stopping criterion for active learning of frame assignment," in *Proceedings of the NAACL HLT 2010 Workshop on Active Learning for Natural Language Processing*, ser. ALNLP '10. Stroudsburg, PA, USA: Association for Computational Linguistics, 2010, pp. 1–9. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1860625.1860626>
- [26] 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>
- [27] 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.
- [28] 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. Rouveilol, Eds., vol. 1398. Springer, 1998, pp. 137–142.
- [29] 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.
- [30] A. K. McCallum and K. Nigam, "Employing EM in pool-based active learning for text classification," in *Proceedings of ICML-98, 15th International Conference on Machine Learning*, J. W. Shavlik, Ed. Madison, US: Morgan Kaufmann Publishers, San Francisco, US, 1998, pp. 350–358.
- [31] 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>
- [32] 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, 2002.
- [33] C. Campbell, N. Cristianini, and A. J. Smola, "Query learning with large margin classifiers," in *Proceedings of the 17th International Conference on Machine Learning (ICML)*, 2000, pp. 111–118.