

Analysis of Stopping Active Learning based on Stabilizing Predictions

Michael Bloodgood

Center for Advanced Study of Language
University of Maryland
College Park, MD 20740
meb@umd.edu

John Grothendieck

Raytheon BBN Technologies
9861 Broken Land Parkway, Suite 400
Columbia, MD 21046
jgrothen@bbn.com

Abstract

Within the natural language processing (NLP) community, active learning has been widely investigated and applied in order to alleviate the annotation bottleneck faced by developers of new NLP systems and technologies. This paper presents the first theoretical analysis of stopping active learning based on stabilizing predictions (SP). The analysis has revealed three elements that are central to the success of the SP method: (1) bounds on Cohen’s Kappa agreement between successively trained models impose bounds on differences in F-measure performance of the models; (2) since the stop set does not have to be labeled, it can be made large in practice, helping to guarantee that the results transfer to previously unseen streams of examples at test/application time; and (3) good (low variance) sample estimates of Kappa between successive models can be obtained. Proofs of relationships between the level of Kappa agreement and the difference in performance between consecutive models are presented. Specifically, if the Kappa agreement between two models exceeds a threshold T (where $T > 0$), then the difference in F-measure performance between those models is bounded above by $\frac{4(1-T)}{T}$ in all cases. If precision of the positive conjunction of the models is assumed to be p , then the bound can be tightened to $\frac{4(1-T)}{(p+1)T}$.

1 Introduction

Active learning (AL), also called *query learning* and *selective sampling*, is an approach to reduce the costs of creating training data that has received considerable interest (e.g., (Argamon-Engelson

and Dagan, 1999; Baldrige and Osborne, 2008; Bloodgood and Vijay-Shanker, 2009b; Bloodgood and Callison-Burch, 2010; Hachey et al., 2005; Haertel et al., 2008; Haffari and Sarkar, 2009; Hwa, 2000; Lewis and Gale, 1994; Sassano, 2002; Settles and Craven, 2008; Shen et al., 2004; Thompson et al., 1999; Tomanek et al., 2007; Zhu and Hovy, 2007)).

Within the NLP community, active learning has been widely investigated and applied in order to alleviate the annotation bottleneck faced by developers of new NLP systems and technologies. The main idea is that by judiciously selecting which examples to have labeled, annotation effort will be focused on the most helpful examples and less annotation effort will be required to achieve given levels of performance than if a passive learning policy had been used.

Historically, the problem of developing methods for detecting when to stop AL was tabled for future work and the research literature was focused on how to select which examples to have labeled and analyzing the selection methods (Cohn et al., 1996; Seung et al., 1992; Freund et al., 1997; Roy and McCallum, 2001). However, to realize the savings in annotation effort that AL enables, we must have a method for knowing when to stop the annotation process. The challenge is that if we stop too early while useful generalizations are still being made, then we can wind up with a model that performs poorly, but if we stop too late after all the useful generalizations are made, then human annotation effort is wasted and the benefits of using active learning are lost.

Recently research has begun to develop methods for stopping AL (Schohn and Cohn, 2000; Ertekin et al., 2007b; Ertekin et al., 2007a; Zhu and Hovy, 2007; Laws and Schütze, 2008; Zhu et al., 2008a; Zhu et al., 2008b; Vlachos, 2008; Bloodgood, 2009; Bloodgood and Vijay-Shanker, 2009a; Ghayoomi, 2010). The methods are all

heuristics based on estimates of model confidence, error, or stability. Although these heuristic methods have appealing intuitions and have had experimental success on a small handful of tasks and datasets, the methods are not widely usable in practice yet because our community’s understanding of the stopping methods remains too coarse and inexact. Pushing forward on understanding the mechanics of stopping at a more exact level is therefore crucial for achieving the design of widely usable effective stopping criteria.

Bloodgood and Vijay-Shanker (2009a) introduce the terminology *aggressive* and *conservative* to describe the behavior of stopping methods¹ and conduct an empirical evaluation of the different published stopping methods on several datasets. While most stopping methods tend to behave conservatively, stopping based on stabilizing predictions computed via inter-model Kappa agreement has been shown to be consistently aggressive without losing performance (in terms of F-Measure²) in several published empirical tests. This method stops when the Kappa agreement between consecutively learned models during AL exceeds a threshold for three consecutive iterations of AL. Although this is an intuitive heuristic that has performed well in published experimental results, there has not been any theoretical analysis of the method.

The current paper presents the first theoretical analysis of stopping based on stabilizing predictions. The analysis helps to explain at a deeper and more exact level *why* the method works as it does. The results of the analysis help to characterize classes of problems where the method can be expected to work well and where (unmodified) it will not be expected to work as well. The theory is suggestive of modifications to improve the robustness of the stopping method for certain classes of problems. And perhaps most important, the approach that we use in our analysis provides an enabling framework for more precise analysis of stopping criteria and possibly other parts of the active learning decision space.

In addition, the information presented in this pa-

¹Aggressive methods stop sooner, aggressively trying to reduce unnecessary annotations while conservative methods are careful not to risk losing model performance, even if it means annotating many more examples than were necessary.

²For the rest of this paper, we will use F-measure to denote F1-measure, that is, the balanced harmonic mean of precision and recall, which is a standard metric used to evaluate NLP systems.

per is useful for works that consider switching between different active learning strategies and operating regions such as (Baram et al., 2004; Dönmez et al., 2007; Roth and Small, 2008). Knowing when to switch strategies, for example, is similar to the stopping problem and is another setting where detailed understanding of the variance of stabilization estimates and their link to performance ramifications is useful. More exact understanding of the mechanics of stopping is also useful for applications of co-training (Blum and Mitchell, 1998), and agreement-based co-training (Clark et al., 2003) in particular. Finally, the proofs of the Theorems regarding the relationships between Cohen’s Kappa statistic and F-measure may be of broader use in works that consider inter-annotator agreement and its ramifications for performance appraisals, a topic that has been of longstanding interest in computational linguistics (Carletta, 1996; Artstein and Poesio, 2008).

In the next section we summarize the stabilizing predictions (SP) stopping method. Section 3 analyzes SP and Section 4 concludes.

2 Stopping Active Learning based on Stabilizing Predictions

The intuition behind the SP method is that the models learned during AL can be applied to a large representative set of unlabeled data called a *stop set* and when consecutively learned models have high agreement on their predictions for classifying the examples in the stop set, this indicates that it is time to stop (Bloodgood and Vijay-Shanker, 2009a; Bloodgood, 2009). The active learning stopping strategy explicitly examined in (Bloodgood and Vijay-Shanker, 2009a) (after the general form is discussed) is to calculate Cohen’s Kappa agreement statistic between consecutive rounds of active learning and stop once it is above 0.99 for three consecutive calculations.

Since the Kappa statistic is an important aspect of this method, we now discuss some background regarding measuring agreement in general, and Cohen’s Kappa in particular. Measurement of agreement between human annotators has received significant attention and in that context, the drawbacks of using percentage agreement have been recognized (Artstein and Poesio, 2008). Alternative metrics have been proposed that take chance agreement into account. Artstein and Poesio (2008) survey several agreement metrics. Most

of the agreement metrics they discuss are of the form:

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

where A_o = observed agreement, and A_e = agreement expected by chance. The different metrics differ in how they compute A_e . All the instances of usage of an agreement metric in this article will have two categories and two coders. The two categories are “+1” and “-1” and the two coders are the two consecutive models for which agreement is being measured.

Cohen’s Kappa statistic³ (Cohen, 1960) measures agreement expected by chance by modeling each coder (in our case model) with a separate distribution governing their likelihood of assigning a particular category. Formally, Kappa is defined by Equation 1 with A_e computed as follows:

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

where each c_i is one of the coders (in our case, models), and $P(k|c_i)$ is the probability that coder (model) c_i labels an instance as being in category k . Kappa estimates the $P(k|c_i)$ in Equation 2 based on the proportion of observed instances that coder (model) c_i labeled as being in category k .

3 Analysis

This section analyzes the SP stopping method. Section 3.1 analyzes the variance of the estimator of Kappa that SP uses and in particular the relationship of this variance to specific aspects of the operationalization of SP, such as the stop set size. Section 3.2 analyzes relationships between the Kappa agreement between two models and the difference in F-measure between those two models.

3.1 Variance of Kappa Estimator

SP bases its decision to stop on the information contained in the contingency tables between the classifications of models learned at consecutive iterations during AL. In determining whether to stop at iteration t , the classifications of the current model M_t are compared with the classifications of the previous model M_{t-1} . Table 1 shows the population parameters for these two models, where:

³We note that there are other agreement measures (beyond Cohen’s Kappa) which could also be applicable to stopping based on stabilizing predictions, but an analysis of these is outside the scope of the current paper.

M_{t-1}	M_t		Total
	+	-	
+	π_{++}	π_{+-}	π_{+}
-	π_{-+}	π_{--}	π_{-}
Total	$\pi_{\cdot+}$	$\pi_{\cdot-}$	1

Table 1: Contingency table population probabilities for M_t (model learned at iteration t) and M_{t-1} (model learned at iteration $t-1$).

population probability π_{ij} for $i, j \in \{+, -\}$ is the probability of an example being placed in category i by model M_{t-1} and category j by model M_t ; population probability $\pi_{\cdot j}$ for $j \in \{+, -\}$ is the probability of an example being placed in category j by model M_t ; and population probability π_i for $i \in \{+, -\}$ is the probability of an example being placed in category i by model M_{t-1} . The actual probability of agreement is $\pi_o = \pi_{++} + \pi_{--}$. As indicated in Equation 2, Kappa models the probability of agreement expected due to chance by assuming that classifications are made independently. Hence, the probability of agreement expected by chance in terms of the population probabilities is $\pi_e = \pi_{\cdot+} \cdot \pi_{+} + \pi_{\cdot-} \cdot \pi_{-}$. From the definition of Kappa (see Equation 1), we then have that the Kappa parameter K in terms of the population probabilities is given by

$$K = \frac{\pi_o - \pi_e}{1 - \pi_e}. \quad (3)$$

For practical applications we will not know the true population probabilities and we will have to resort to using sample estimates. The SP method uses a stop set of size n for deriving its estimates. Table 2 shows the contingency table counts for the classifications of models M_t and M_{t-1} on a sample of size n . The population probabilities π_{ij} can be estimated by the relative frequencies p_{ij} for $i, j \in \{+, -, \cdot\}$, where: $p_{++} = a/n$; $p_{+-} = b/n$; $p_{-+} = c/n$; $p_{--} = d/n$; $p_{\cdot+} = (a+b)/n$; $p_{\cdot-} = (c+d)/n$; $p_{+} = (a+c)/n$; and $p_{-} = (b+d)/n$. Let $p_o = p_{++} + p_{--}$, the observed proportion of agreement and let $p_e = p_{\cdot+} \cdot p_{+} + p_{\cdot-} \cdot p_{-}$, the proportion of agreement expected by chance if we assume that M_t and M_{t-1} make their classifications independently. Then the Kappa measure of agreement K between M_t and M_{t-1} (see Equation 3) is estimated by

$$\hat{K} = \frac{p_o - p_e}{1 - p_e}. \quad (4)$$

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

Table 2: Contingency table counts for M_t (model learned at iteration t) and M_{t-1} (model learned at iteration t-1).

Using the delta method, as described in (Bishop et al., 1975), Fleiss et al. (1969) derived an estimator of the large-sample variance of \hat{K} . According to Hale and Fleiss (1993), the estimator simplifies to

$$\begin{aligned}
\text{Var}(\hat{K}) &= \frac{1}{n(1-p_e)^2} \times \\
&\left\{ \sum_{i \in \{+, -\}} p_{ii} [1 - 4\bar{p}_i(1 - \hat{K})] \right. \\
&\quad - (\hat{K} - p_e(1 - \hat{K}))^2 + (1 - \hat{K})^2 \times \\
&\quad \left. \sum_{i,j \in \{+, -\}} p_{ij} [2(\bar{p}_i + \bar{p}_j) - (p_{i.} + p_{.j})]^2 \right\}, \tag{5}
\end{aligned}$$

where $\bar{p}_i = (p_{i.} + p_{.i})/2$. From Equation 5, we can see that the variance of our estimate of Kappa is inversely proportional to the size of the stop set we use.

Bloodgood and Vijay-Shanker (2009a) used a stop set of size 2000 for each of their datasets. Although this worked well in the results they reported, we do not believe that 2000 is a fixed size that will work well for all tasks and datasets where the SP method could be used. Table 3 shows the variances of \hat{K} computed using Equation 5 at the points at which SP stopped AL for each of the datasets⁴ from (Bloodgood and Vijay-Shanker, 2009a).

These variances indicate that the size of 2000 was typically sufficient to get tight estimates of Kappa, helping to illuminate the empirical success of the SP method on these datasets. More generally, the SP method can be augmented with a variance check: if the variance of estimated Kappa at a potential stopping point exceeds some desired

⁴We note that each of the datasets was set up as a binary classification task (or multiple binary classification tasks). Further details and descriptions of each of the datasets can be found in (Bloodgood and Vijay-Shanker, 2009a).

threshold, then the stop set size can be increased as needed to reduce the variance.

Looking at Equation 5 again, one can note that when p_e is relatively close to 1, the variance of \hat{K} can be expected to get quite large. In these situations, users of SP should expect to have to use larger stop set sizes and in extreme conditions, SP may not be an advisable method to use.

3.2 Relationship between Kappa agreement and change in performance between models

Heretofore, the published literature contained only informal explanations of why stabilizing predictions is expected to work well as a stopping method (along with empirical tests demonstrating successful operation on a handful of tasks and datasets). In the remainder of this section we describe the mathematical foundations for stopping methods based on stabilizing predictions. In particular, we will prove that even in the worst possible case, if the Kappa agreement between two subsequently learned models is greater than a threshold T , then it must be the case that the change in performance between these two models is bounded above by $\frac{4(1-T)}{T}$. We then go on to prove additional Theorems that tighten this bound when assumptions are made about model precision.

Lemma 3.1 *Suppose F-measure F and Kappa K are computed from the same contingency table of counts, such as the one given in Table 2. Suppose $ad - bc \geq 0$. Then $F \geq K$.*

Proof By definition, in terms of the contingency table counts,

$$K = \frac{2ad - 2bc}{(a+b)(b+d) + (a+c)(c+d)} \tag{6}$$

and

$$F = \frac{2a}{2a + b + c}. \tag{7}$$

Rewriting F so that it will have the same numerator as K , we have:

$$F = F \left(\frac{d - \frac{bc}{a}}{d - \frac{bc}{a}} \right) \tag{8}$$

$$= \left(\frac{2a}{2a + b + c} \right) \left(\frac{d - \frac{bc}{a}}{d - \frac{bc}{a}} \right) \tag{9}$$

$$= \frac{2ad - 2bc}{2ad + bd + cd - 2bc - \frac{b^2c + bc^2}{a}}. \tag{10}$$

Task-Dataset	Variance of \hat{K}
NER-DNA (10-fold CV)	0.0000223
NER-cellType (10-fold CV)	0.0000211
NER-protein (10-fold CV)	0.0000074
Reuters (10 Categories)	0.0000298
20 Newsgroups (20 Categories)	0.0000739
WebKB Student (10-fold CV)	0.0000137
WebKB Project (10-fold CV)	0.0000190
WebKB Faculty (10-fold CV)	0.0000115
WebKB Course (10-fold CV)	0.0000179
TC-spamassassin (10-fold CV)	0.0000042
TC-TREC-SPAM (10-fold CV)	0.0000043
Average (macro-avg)	0.0000209

Table 3: Estimates of the variance of \hat{K} . For each dataset, the estimate of the variance of \hat{K} is computed (using Equation 5) from the contingency table at the point at which SP stopped AL and the average of all the variances (across all folds of CV) is displayed. The last row contains the macro-average of the average variances for all the datasets.

We can see that the expression for F in Equation 10 has the same numerator as K in Equation 6 but the denominator of K in Equation 6 is \geq the denominator of F in Equation 10. Therefore, $F \geq K$. ■

Theorem 3.2 *Let M_t be the model learned at iteration t of active learning and M_{t-1} be the model learned at iteration $t - 1$. Let K_t be the estimate of Kappa agreement between the classifications of M_t and M_{t-1} on the examples in the stop set. Let \tilde{F}_t be the F-measure between the classifications of M_t and truth on the stop set. Let \tilde{F}_{t-1} be the F-measure between the classifications of M_{t-1} and truth on the stop set. Let $\Delta\tilde{F}_t$ be $\tilde{F}_t - \tilde{F}_{t-1}$. Suppose $T > 0$. Then $K_t > T \Rightarrow |\Delta\tilde{F}_t| \leq \frac{4(1-T)}{T}$.*

Proof Suppose M_t , M_{t-1} , K_t , \tilde{F}_t , \tilde{F}_{t-1} , $\Delta\tilde{F}_t$, and T are defined as stated in the statement of Theorem 3.2. Let F_t be the F-measure between the classifications of M_t and M_{t-1} on the examples in the stop set. Let Table 2 show the contingency table counts for M_t versus M_{t-1} on the examples in the stop set. Then, from their definitions, we have $K_t = \frac{2(ad-bc)}{(a+b)(b+d)+(a+c)(c+d)}$ and $F_t = \frac{2a}{2a+b+c}$. There exist true labels for the examples in the stop set, which we don't know since the stop set is unlabeled, but nonetheless must exist. We use the truth on the stop set to split Table 2 into two subtables of counts, one table for all the examples that are truly positive and one table for all the examples that are truly negative. Table 4

M_{t-1}	M_t		Total
	+	-	
+	a_1	b_1	$a_1 + b_1$
-	c_1	d_1	$c_1 + d_1$
Total	$a_1 + c_1$	$b_1 + d_1$	n_1

Table 4: Contingency table counts for M_t (model learned at iteration t) versus M_{t-1} (model learned at iteration t-1) for only the examples in the stop set that have truth = +1.

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

Table 5: Contingency table counts for M_t (model learned at iteration t) versus M_{t-1} (model learned at iteration t-1) for only the examples in the stop set that have truth = -1.

shows the contingency table for M_t versus M_{t-1} for all of the examples in the stop set that have true labels of +1 and Table 5 shows the contingency table for M_t versus M_{t-1} for all of the examples in the stop set that have true labels of -1.

From Tables 2, 4, and 5 one can see that a is the number of examples in the stop set that both M_t and M_{t-1} classified as positive. Furthermore, out of these a examples, a_1 of them truly are pos-

Truth	M_t		Total
	+	-	
+	$a_1 + c_1$	$b_1 + d_1$	n_1
-	$a_{-1} + c_{-1}$	$b_{-1} + d_{-1}$	n_{-1}
Total	$a + c$	$b + d$	n

Table 6: Contingency table counts for M_t (model learned at iteration t) versus truth. (Derived from Tables 4 and 5)

Truth	M_{t-1}		Total
	+	-	
+	$a_1 + b_1$	$c_1 + d_1$	n_1
-	$a_{-1} + b_{-1}$	$c_{-1} + d_{-1}$	n_{-1}
Total	$a + b$	$c + d$	n

Table 7: Contingency table counts for M_{t-1} (model learned at iteration t-1) versus truth. (Derived from Tables 4 and 5)

itive and a_{-1} of them truly are negative. Similar explanations hold for the other counts. Also, from Tables 2, 4, and 5, one can see that the equalities $a = a_1 + a_{-1}$, $b = b_1 + b_{-1}$, $c = c_1 + c_{-1}$, and $d = d_1 + d_{-1}$ all hold. The contingency tables for M_t versus truth and M_{t-1} versus truth can be derived from Tables 4 and 5. For convenience, Table 6 shows the contingency table for M_t versus truth and Table 7 shows the contingency table for M_{t-1} versus truth. Suppose that $K_t > T$. This implies, by Lemma 3.1⁵, that $F_t > T$. This implies that

$$\frac{2a}{2a+b+c} > T \quad (11)$$

$$\Rightarrow 2a > (2a + b + c)T \quad (12)$$

$$\Rightarrow 2a(1 - T) > (b + c)T \quad (13)$$

$$\Rightarrow b + c < \frac{2a(1-T)}{T}. \quad (14)$$

Note that Equations 12 and 14 are justified since $2a + b + c > 0$ and $T > 0$, respectively.

From Table 6 we can see that $\tilde{F}_t = \frac{2(a_1+c_1)}{2(a_1+c_1)+b_1+d_1+a_{-1}+c_{-1}}$; from Table 7 we can see that $\tilde{F}_{t-1} = \frac{2(a_1+b_1)}{2(a_1+b_1)+c_1+d_1+a_{-1}+b_{-1}}$. For notational convenience, let: $g = 2(a_1 + c_1) + b_1 + d_1 + a_{-1} + c_{-1}$; and $h = 2(a_1 + b_1) + c_1 + d_1 + a_{-1} + b_{-1}$.

⁵Note that the condition $ad - bc \geq 0$ of Lemma 3.1 is met since $K_t > T$ and $T > 0$ imply $K_t > 0$, which in turn implies $ad - bc > 0$.

It follows that

$$\Delta F_t = \frac{2(a_1 + c_1)}{g} - \frac{2(a_1 + b_1)}{h} \quad (15)$$

$$= \frac{(2a_1 + 2c_1)h - (2a_1 + 2b_1)g}{gh} \quad (16)$$

For notational convenience, let: $x = 2(a_1c_1 + a_1b_{-1} + c_1^2 + c_1d_1 + c_1a_{-1} + c_1b_{-1})$; and $y = 2(a_1b_1 + a_1c_{-1} + b_1^2 + b_1d_1 + b_1a_{-1} + b_1c_{-1})$. Then picking up from Equation 16, it follows that

$$\Delta F_t = \frac{x - y}{gh} \quad (17)$$

$$= \frac{2[u_1 + c_1u_2 - b_1u_3]}{gh}, \quad (18)$$

where $u_1 = a_1c_1 - a_1b_1 + a_1b_{-1} - a_1c_{-1}$, $u_2 = c_1 + d_1 + a_{-1} + b_{-1}$, and $u_3 = b_1 + d_1 + a_{-1} + c_{-1}$.

For notational convenience, let: $d_A = c_1 - b_1$ and $d_B = c_{-1} - b_{-1}$. Then it follows that

$$\Delta F_t = \frac{2u_4}{gh}, \quad (19)$$

where: $u_4 = a_1(d_A - d_B) + d_A(d_1 + a_{-1} + b_1 + c_1) + c_1b_{-1} - b_1c_{-1}$.

Noting that $g = h + d_A + d_B$, we have

$$\Delta F_t = \frac{2u_4}{h(h + d_A + d_B)}. \quad (20)$$

Noting that $2u_4 = 2[d_A(a_1 + b_1 + c_1 + d_1 + a_{-1} + b_{-1}) - d_B(a_1 + b_1)]$ and letting $u_5 = a_1 + b_1 + c_1 + d_1 + a_{-1} + b_{-1}$, we have

$$\Delta F_t = \frac{2[d_Au_5 - d_B(a_1 + b_1)]}{h(h + d_A + d_B)}. \quad (21)$$

Therefore,

$$|\Delta F_t| \leq 2 \left(\left| \frac{d_Au_5}{h(h + d_A + d_B)} \right| + \left| \frac{d_B(a_1 + b_1)}{h(h + d_A + d_B)} \right| \right) \quad (22)$$

Recall that $b + c = b_1 + b_{-1} + c_1 + c_{-1}$. Then observe that the following three inequalities hold: $b + c \geq d_A$; $b + c \geq d_B$; and $h(h + d_A + d_B) > 0$. Therefore,

$$|\Delta F_t| \leq \frac{2(b+c)[2a_1+2b_1+c_1+d_1+a_{-1}+b_{-1}]}{h(h+d_A+d_B)} \quad (23)$$

$$= \frac{2(b+c)h}{h(h+d_A+d_B)} \quad (24)$$

$$= \frac{2(b+c)}{h+d_A+d_B} \quad (25)$$

$$\leq \frac{2(2a)(1-T)}{T(h+d_A+d_B)} \quad (26)$$

$$= \left(\frac{4(1-T)}{T} \right) \left(\frac{a}{h+d_A+d_B} \right). \quad (27)$$

Observe that $h + d_A + d_B = 2a_1 + b_1 + 2c_1 + d_1 + a_{-1} + c_{-1}$. Therefore, $\frac{a}{h+d_A+d_B} \leq 1$. Therefore, we have

$$|\Delta F_t| \leq \frac{4(1-T)}{T}. \quad \blacksquare \quad (28)$$

Note that in deriving Inequality 26, we used the previously derived Inequality 14. Also, the proof of Theorem 3.2 assumes a worst possible case in the sense that all examples where the classifications of M_t and M_{t-1} differ are assumed to have truth values that all serve to maximize one model's F-measure and minimize the other model's F-measure so as to maximize $|\Delta F_t|$ as much as possible. A resulting limitation is that the bound is loose in many cases. It may be possible to derive tighter bounds, perhaps by easing off to an expected case instead of a worst case and/or by making additional assumptions.⁶

Taking this possibility up, we now prove tighter bounds when assumptions about the precision of the models M_t and M_{t-1} are made. Consider that in the proof of Theorem 3.2 when transitioning from Equality 27 to Inequality 28, we used the fact that $\frac{a}{h+d_A+d_B} \leq 1$. Note that $\frac{a}{h+d_A+d_B} = \frac{a}{2a_1+b_1+2c_1+d_1+a_{-1}+c_{-1}}$, from which one sees that $\frac{a}{h+d_A+d_B} = 1$ only if all of a_1, b_1, c_1, d_1 and c_{-1} are all zero. This is a pathological case. In many practically important classes of cases to consider, $\frac{a}{h+d_A+d_B}$ will be strictly less than 1, and often substantially less than 1. The following two Theorems prove tighter bounds on $|\Delta F_t|$ than Theorem 3.2 by utilizing this insight.

Theorem 3.3 *Suppose $M_t, M_{t-1}, K_t, \tilde{F}_t, \tilde{F}_{t-1}, \Delta F_t$, and T are defined as stated in the statement of Theorem 3.2. Let the contingency tables be defined as they were in the proof of Theorem 3.3. Let $M_{PositiveConjunction}$ be a model that only classifies an example as positive if both models M_t and M_{t-1} classify the example as positive. Suppose that $M_{PositiveConjunction}$ has perfect precision on the stop set, or in other words that every single example from the stop set that both M_t and M_{t-1} classify as positive is truthfully positive (i.e., $a_{-1} = 0$). Then $K_t > T \Rightarrow |\Delta F_t| \leq \frac{2(1-T)}{T}$.*

Proof The proof of Theorem 3.2 holds exactly as it is up until Equality 27. Now, using the additional assumption that $a_{-1} = 0$, we have

⁶If one is planning to undertake this challenge, we would suggest further consideration of Inequalities 22, 23, 26, and 28 as a possible starting point.

$\frac{a}{h+d_A+d_B} \leq \frac{1}{2}$. Therefore, we have

$$|\Delta F_t| \leq \frac{2(1-T)}{T}. \quad \blacksquare \quad (29)$$

Theorem 3.3 is a special case (in the limit) of a more general Theorem. Before stating and proving the more general Theorem, we prove a Lemma that will be helpful in making the proof of the general Theorem clearer.

Lemma 3.4 *Let f, d_A, d_B and contingency table counts be defined as they were in the proof of Theorem 3.2. Suppose $a_1 = xa_{-1}$. Then $\frac{a}{h+d_A+d_B} \leq \frac{x+1}{2x+1}$.*

Proof $a_1 = xa_{-1}$ by hypothesis. $a = a_1 + a_{-1}$ by definition of contingency table counts. Hence, $a = (x+1)a_{-1}$. Therefore,

$$\begin{aligned} \frac{a}{h+d_A+d_B} &\leq \frac{(x+1)a_{-1}}{2xa_{-1}+a_{-1}} & (30) \\ &= \frac{(x+1)a_{-1}}{(2x+1)a_{-1}} \\ &= \frac{x+1}{2x+1}. \quad \blacksquare \end{aligned}$$

The following Theorem generalizes Theorem 3.3 to cases when $M_{PositiveConjunction}$ has precision p in $(0, 1)$.⁷

Theorem 3.5 *Suppose $M_t, M_{t-1}, K_t, \tilde{F}_t, \tilde{F}_{t-1}, \Delta F_t$, and T are defined as stated in the statement of Theorem 3.2. Let the contingency tables be defined as they were in the proof of Theorem 3.2. Let $M_{PositiveConjunction}$ be a model that only classifies an example as positive if both models M_t and M_{t-1} classify the example as positive. Suppose that $M_{PositiveConjunction}$ has precision p on the stop set. Then $K_t > T \Rightarrow |\Delta F_t| \leq \frac{4(1-T)}{(p+1)T}$.*

Proof The proof of Theorem 3.2 holds exactly as it is up until Equality 27. $M_{PositiveConjunction}$ has precision p on the stop set $\Rightarrow p = \frac{a_1}{a_1+a_{-1}}$. Solving for a_1 in terms of a_{-1} we have $a_1 = \frac{p}{1-p}a_{-1}$. Therefore, applying Lemma 3.4 with $x = \frac{p}{1-p}$, we have $\frac{a}{h+d_A+d_B} \leq \frac{\frac{p}{1-p}+1}{\frac{2p}{1-p}+1}$. Therefore we have

$$|\Delta F_t| \leq 4 \left(\frac{\frac{p}{1-p}+1}{\frac{2p}{1-p}+1} \right) \frac{(1-T)}{T} \quad (31)$$

$$= \frac{4(1-T)}{(p+1)T}. \quad \blacksquare \quad (32)$$

⁷The case when $p = 0$ is handled by Theorem 3.2 and the case when $p = 1$ is handled by Theorem 3.3.

Precision	$\frac{1}{p+1}$ (to 3 decimal places)
50%	0.667
80%	0.556
90%	0.526
95%	0.513
98%	0.505
99%	0.503
99.9%	0.500

Table 8: Values of the scaling factor from Theorem 3.5 for different precision values.

The scaling factor $\frac{1}{p+1}$ in Theorem 3.5 shows how the precision of the conjunctive model affects the bound. Theorem 3.2 had the scaling factor implicitly set to 1 in order to handle the pathological case where the positive conjunctive model has precision = 0. In Theorem 3.3, where the positive conjunctive model has precision = 1 on the examples in the stop set, the scaling factor is set to 1/2. Theorem 3.5 generalizes the scaling factor so that it is a function of the precision of the positive conjunctive model. For convenience, Table 8 shows the scaling factor values for a few different precision values.

The bounds in Theorems 3.2, 3.3, and 3.5 all bound the difference in performance *on the stop set* of two consecutively learned models M_t and M_{t-1} . An issue to consider is how connected the difference in performance on the stop set is to the difference in performance on a stream of application examples generated according to the population probabilities. Taking up this issue, consider that the proof of Theorems 3.2, 3.3, and 3.5 would hold as it is if we had used sample proportions instead of sample counts (this can be seen by simply dividing every count by n , the size of the stop set). Since the stop set is unbiased (selected at random from the population), as n approaches infinity, the sample proportions will approach the population probabilities and the difference between the difference in performance between M_t and M_{t-1} on the stop set and on a stream of application examples generated according to the population probabilities will approach zero.

4 Conclusions

To date, the work on stopping criteria has been dominated by heuristics based on intuitions and experimental success on a small handful of tasks

and datasets. But the methods are not widely usable in practice yet because our community’s understanding of the stopping methods remains too inexact. Pushing forward on understanding the mechanics of stopping at a more exact level is therefore crucial for achieving the design of widely usable effective stopping criteria.

This paper presented the first theoretical analysis of stopping based on stabilizing predictions. The analysis revealed three elements that are central to the SP method’s success: (1) the sample estimates of Kappa have low variance; (2) Kappa has tight connections with differences in F-measure; and (3) since the stop set doesn’t have to be labeled, it can be arbitrarily large, helping to guarantee that the results transfer to previously unseen streams of examples at test/application time.

We presented proofs of relationships between the level of Kappa agreement and the difference in performance between consecutive models. Specifically, if the Kappa agreement between two models is at least T , then the difference in F-measure performance between those models is bounded above by $\frac{4(1-T)}{T}$. If precision of the positive conjunction of the models is assumed to be p , then the bound can be tightened to $\frac{4(1-T)}{(p+1)T}$.

The setup and methodology of the proofs can serve as a launching pad for many further investigations, including: analyses of stopping; works that consider switching between different active learning strategies and operating regions; and works that consider stopping co-training, and especially agreement-based co-training. Finally, the relationships that have been exposed between the Kappa statistic and F-measure may be of broader use in works that consider inter-annotator agreement and its interplay with system evaluation, a topic that has been of long-standing interest.

References

- Shlomo Argamon-Engelson and Ido Dagan. 1999. Committee-based sample selection for probabilistic classifiers. *Journal of Artificial Intelligence Research (JAIR)*, 11:335–360.
- Ron Artstein and Massimo Poesio. 2008. Inter-coder agreement for computational linguistics. *Computational Linguistics*, 34(4):555–596.
- Jason Baldridge and Miles Osborne. 2008. Active learning and logarithmic opinion pools for hpsg parse selection. *Nat. Lang. Eng.*, 14(2):191–222.

- Yoram Baram, Ran El-Yaniv, and Kobi Luz. 2004. Online choice of active learning algorithms. *Journal of Machine Learning Research*, 5:255–291, March.
- Yvonne M. Bishop, Stephen E. Fienberg, and Paul W. Holland. 1975. *Discrete Multivariate Analysis: Theory and Practice*. MIT Press, Cambridge, MA.
- Michael Bloodgood and Chris Callison-Burch. 2010. 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*, pages 854–864, Uppsala, Sweden, July. Association for Computational Linguistics.
- Michael Bloodgood and K Vijay-Shanker. 2009a. 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)*, pages 39–47, Boulder, Colorado, June. Association for Computational Linguistics.
- Michael Bloodgood and K Vijay-Shanker. 2009b. 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*, pages 137–140, Boulder, Colorado, June. Association for Computational Linguistics.
- Michael Bloodgood. 2009. *Active learning with support vector machines for imbalanced datasets and a method for stopping active learning based on stabilizing predictions*. Ph.D. thesis, University of Delaware, Newark, DE, USA.
- Avrim Blum and Tom Mitchell. 1998. Combining labeled and unlabeled data with co-training. In *COLT'98: Proceedings of the eleventh annual conference on Computational learning theory*, pages 92–100, New York, NY, USA. ACM.
- J. Carletta. 1996. Assessing agreement on classification tasks: The kappa statistic. *Computational linguistics*, 22(2):249–254.
- Stephen Clark, James Curran, and Miles Osborne. 2003. Bootstrapping pos-taggers using unlabelled data. In Walter Daelemans and Miles Osborne, editors, *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 49–55.
- J. Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20:37–46.
- David A. Cohn, Zoubin Ghahramani, and Michael I. Jordan. 1996. Active learning with statistical models. *Journal of Artificial Intelligence Research*, 4:129–145.
- Meryem Pinar Dönmez, Jaime G. Carbonell, and Paul N. Bennett. 2007. Dual strategy active learning. In Joost N. Kok, Jacek Koronacki, Ramon López de Mántaras, Stan Matwin, Dunja Mladenic, and Andrzej Skowron, editors, *Machine Learning: ECML 2007, 18th European Conference on Machine Learning, Warsaw, Poland, September 17-21, 2007, Proceedings*, volume 4701 of *Lecture Notes in Computer Science*, pages 116–127. Springer.
- Seyda Ertekin, Jian Huang, Léon Bottou, and C. Lee Giles. 2007a. Learning on the border: active learning in imbalanced data classification. In Mário J. Silva, Alberto H. F. Laender, Ricardo A. Baeza-Yates, Deborah L. McGuinness, Bjørn Olstad, Øystein Haug Olsen, and André O. Falcão, editors, *Proceedings of the Sixteenth ACM Conference on Information and Knowledge Management, CIKM 2007, Lisbon, Portugal, November 6-10, 2007*, pages 127–136. ACM.
- Seyda Ertekin, Jian Huang, and C. Lee Giles. 2007b. Active learning for class imbalance problem. In Wessel Kraaij, Arjen P. de Vries, Charles L. A. Clarke, Norbert Fuhr, and Noriko Kando, editors, *SIGIR 2007: Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Amsterdam, The Netherlands, July 23-27, 2007*, pages 823–824. ACM.
- Joseph L. Fleiss, Jacob Cohen, and B. S. Everitt. 1969. Large sample standard errors of kappa and weighted kappa. *Psychological Bulletin*, 72(5):323 – 327.
- Yoav Freund, H. Sebastian Seung, Eli Shamir, and Naf-tali Tishby. 1997. Selective sampling using the query by committee algorithm. *Machine Learning*, 28:133–168.
- Masood Ghayoomi. 2010. 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*, pages 1–9, Los Angeles, California, June. Association for Computational Linguistics.
- Ben Hachey, Beatrice Alex, and Markus Becker. 2005. Investigating the effects of selective sampling on the annotation task. In *Proceedings of the Ninth Conference on Computational Natural Language Learning (CoNLL-2005)*, pages 144–151, Ann Arbor, Michigan, June. Association for Computational Linguistics.
- Robbie Haertel, Eric Ringger, Kevin Seppi, James Carroll, and Peter McClanahan. 2008. Assessing the costs of sampling methods in active learning for annotation. In *Proceedings of ACL-08: HLT, Short Papers*, pages 65–68, Columbus, Ohio, June. Association for Computational Linguistics.
- Gholamreza Haffari and Anoop Sarkar. 2009. Active learning for multilingual statistical machine translation. In *Proceedings of the Joint Conference of*

- the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP, pages 181–189, Suntec, Singapore, August. Association for Computational Linguistics.
- Cecilia A. Hale and Joseph L. Fleiss. 1993. Interval estimation under two study designs for kappa with binary classifications. *Biometrics*, 49(2):pp. 523–534.
- Rebecca Hwa. 2000. Sample selection for statistical grammar induction. In Hinrich Schütze and Keh-Yih Su, editors, *Proceedings of the 2000 Joint SIGDAT Conference on Empirical Methods in Natural Language Processing*, pages 45–53. Association for Computational Linguistics, Somerset, New Jersey.
- Florian Laws and Hinrich Schütze. 2008. Stopping criteria for active learning of named entity recognition. In *Proceedings of the 22nd International Conference on Computational Linguistics (Coling 2008)*, pages 465–472, Manchester, UK, August. Coling 2008 Organizing Committee.
- David D. Lewis and William A. Gale. 1994. 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*, pages 3–12, New York, NY, USA. Springer-Verlag New York, Inc.
- D. Roth and K. Small. 2008. Active learning for pipeline models. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 683–688.
- Nicholas Roy and Andrew McCallum. 2001. Toward optimal active learning through sampling estimation of error reduction. In *Proceedings of the 18th International Conference on Machine Learning*, pages 441–448. Morgan Kaufmann.
- Manabu Sassano. 2002. An empirical study of active learning with support vector machines for japanese word segmentation. In *ACL '02: Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, pages 505–512, Morristown, NJ, USA. Association for Computational Linguistics.
- Greg Schohn and David Cohn. 2000. Less is more: Active learning with support vector machines. In *Proc. 17th International Conf. on Machine Learning*, pages 839–846. Morgan Kaufmann, San Francisco, CA.
- Burr Settles and Mark Craven. 2008. An analysis of active learning strategies for sequence labeling tasks. In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 1070–1079, Honolulu, Hawaii, October. Association for Computational Linguistics.
- H. S. Seung, M. Opper, and H. Sompolinsky. 1992. Query by committee. In *COLT '92: Proceedings of the fifth annual workshop on Computational learning theory*, pages 287–294, New York, NY, USA. ACM.
- Dan Shen, Jie Zhang, Jian Su, Guodong Zhou, and Chew-Lim Tan. 2004. Multi-criteria-based active learning for named entity recognition. In *Proceedings of the 42nd Meeting of the Association for Computational Linguistics (ACL'04), Main Volume*, pages 589–596, Barcelona, Spain, July.
- Cynthia A. Thompson, Mary Elaine Califf, and Raymond J. Mooney. 1999. Active learning for natural language parsing and information extraction. In *Proceedings of the 16th International Conference on Machine Learning*, pages 406–414. Morgan Kaufmann, San Francisco, CA.
- Katrin Tomanek, Joachim Wermter, and Udo Hahn. 2007. An approach to text corpus construction which cuts annotation costs and maintains reusability of annotated data. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 486–495.
- Andreas Vlachos. 2008. A stopping criterion for active learning. *Computer Speech and Language*, 22(3):295–312.
- Jingbo Zhu and Eduard Hovy. 2007. 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)*, pages 783–790.
- Jingbo Zhu, Huizhen Wang, and Eduard Hovy. 2008a. Learning a stopping criterion for active learning for word sense disambiguation and text classification. In *IJCNLP*.
- Jingbo Zhu, Huizhen Wang, and Eduard Hovy. 2008b. Multi-criteria-based strategy to stop active learning for data annotation. In *Proceedings of the 22nd International Conference on Computational Linguistics (Coling 2008)*, pages 1129–1136, Manchester, UK, August. Coling 2008 Organizing Committee.