Understanding the Yarowsky Algorithm

Steven Abney

August 11, 2003

1 Introduction

Bootstrapping, or semi-supervised learning, has become an important topic in computational linguistics. For many language-processing tasks, there is an abundance of unlabeled data, but labeled data is lacking and too expensive to create in large quantities, making bootstrapping techniques desirable.

The Yarowsky algorithm [5] was one of the first bootstrapping algorithms to become widely known in computational linguistics. An alternative algorithm, co-training [2], has subsequently become more popular, perhaps in part because it has proven amenable to theoretical analysis [4], in contrast to the Yarowsky algorithm, which is as yet mathematically poorly understood. The current paper aims to rectify that lack, increasing the attractiveness of the Yarowsky algorithm as an alternative to co-training. The Yarowsky algorithm does have the advantage of placing less restriction on the data sets it can be applied to. Co-training requires data attributes to be separable into two *views* that are conditionally independent given the target label; the Yarowsky algorithm makes no such assumption about its data.

In previous work, I did propose an assumption about the data called precision independence under which the Yarowsky algorithm could be shown effective [1]. That assumption is ultimately unsatisfactory, however, not only because it restricts the data sets on which the algorithm can be shown effective, but for additional internal reasons. A detailed discussion would take us too far afield here, but suffice it to say that precision independence is a property that it would be preferable not to assume, but rather to derive from more basic properties of a data set; and that closer empirical study shows that precision independence fails to be satisfied in some data sets on which the Yarowsky algorithm is effective.

This paper proposes a different approach. Instead of making assumptions about the data, it proposes viewing the Yarowsky algorithm as optimizing an objective function. We will show that several variants of the algorithm (though not the algorithm in precisely its original form) optimize either likelihood or a second, related objective function. To the extent that these variants capture the essence of the original algorithm, we have a better formal understanding of its effectiveness. Even if the variants are deemed to depart substantially from the original algorithm, we have at least obtained a family of new bootstrapping algorithms that are mathematically understood.

We consider two objective functions and three families of Yarowsky algorithm variants. The first objective function, H, is the negative of log likelihood; to minimize H is to maximize likelihood. The second objective function, K, is an upper bound on H that can in principle be reduced to zero, squeezing H to zero as well.

The first Yarowsky variant, called Y-1/DL-EM, combines the "outer loop" of the Yarowsky algorithm with a different "inner loop" based on the Expectation-Maximization (EM) algorithm. It is shown to decrease H in each iteration. The second variant, Y-1/DL-1, uses the same outer loop, but an inner loop that is closer to the original Yarowsky inner loop. It is shown to decrease K in each iteration. The third variant, YS, differs from the first two in that it does sequential update (adding a single rule in each iteration), rather than parallel update (updating all rules in each iteration). It optimizes K. Within each variant, there are subvariants that differ with respect to the update rule that is used to compute rule weights, yielding a total of seven different algorithms that optimize either H or K.

2 The Generic Yarowsky Algorithm

2.1 The Original Algorithm Y-0

The original Yarowsky algorithm, which we call Y-0, is given in table 1. It is an iterative algorithm. One begins with a "seed set" Λ_0 of labeled examples, and a set V_0 of unlabeled examples. At each iteration, a classifier is constructed from the labeled examples; then the classifier is applied to the unlabeled examples to create a new labeled set.

Formally, we assume a set X of examples and a series of partial functions $Y^{(t)}$ assigning labels to examples; t represents the iteration number. An unlabeled example x is one for which $Y^{(t)}(x) = \emptyset$. The set of labeled examples Λ_t and unlabeled examples V_t are defined as follows:

$$\Lambda_t \equiv \{x | Y^{(t)}(x) \neq \emptyset\}$$

$$\mathrm{V}_t \equiv \{x | Y^{(t)}(x) = \emptyset\}$$

It will also be useful to have a notation for the set of examples with label j:

$$Y_j^{(t)} \equiv \{x | Y^{(t)}(x) = j\}$$

Note that Λ_t is the disjoint union of the sets $Y_j^{(t)}$.

In each iteration, one uses a supervised learning algorithm to train a classifier on the labeled examples. Let us call this algorithm the *base learning algorithm*; it is a function from $(X, Y^{(t)})$ to a classifier π drawn from a space of classifiers Π. It is assumed that the classifier makes confidence-weighted predictions. That is, the classifier defines a scoring function $\pi(x, j)$, and the predicted label for example x is

$$\hat{y} \equiv \arg\max_j \pi(x, j) \tag{1}$$

(1) Given: $X, Y^{(0)}$
(2) For $t \in \{0, 1, \ldots\}$
  (2.1) Train classifier on $(\Lambda_t, Y^{(t)})$, where $\Lambda_t = \{x | Y^{(t)} \neq \emptyset\}$
    The result is $\pi^{(t+1)}$
  (2.2) For each example x:
    (2.2.1) Set $\hat{y} = \arg\max_j \pi_x^{(t+1)}(j)$
    (2.2.2) Set $Y^{(t+1)}(x) = \begin{cases} \hat{y} & \text{if } \pi_x^{(t+1)}(\hat{y}) > \zeta \\ \emptyset & \text{otherwise} \end{cases}$
  (2.3) If $Y^{(t+1)} = Y^{(t)}$, stop

Table 1: The generic Yarowsky algorithm (Y-0)

Ties are broken arbitrarily. Technically, we assume a fixed order over labels, and define the maximization to return the first label in the ordering, in case of a tie.

It will be convenient to assume that the scoring function is nonnegative and bounded, in which case we can normalize it to make $\pi(x, j)$ a conditional distribution over labels j for a given example x. Henceforward, we write $\pi_x(j)$ instead of $\pi(x, j)$, understanding π_x to be a probability distribution over labels j. We call this distribution the *prediction distribution* of the classifier on example x.

To complete an iteration of the Yarowsky algorithm, one recomputes labels for examples. Specifically, the label $\hat{y}$ is assigned to example x if the score $\pi_x(\hat{y})$ exceeds a threshold ζ, called the *labeling threshold*. The new labeled set Λ_{t+1} contains all examples for which $\pi_x(\hat{y}) > \zeta$.

The algorithm continues until convergence. The particular base learning algorithm that Yarowsky uses is deterministic, in the sense that the classifier induced is a deterministic function of the labeled data. Hence, the algorithm is known to have converged whenever the labeling remains unchanged.

There is one point on which the text of the original article is somewhat ambiguous. It is not clear whether the labeling in step (2.2) actually applies (a) to all examples, (b) to all originally unlabeled examples, or (c) only to examples that are unlabeled at the beginning of the current iteration. The text is never explicit, and different passages suggest different interpretations. My best guess is that (a) is the intended interpretation, and I have defined Y-0 accordingly.

Note that the algorithm as stated leaves the base learning algorithm unspecified. We can distinguish between the *generic Yarowsky algorithm* Y-0, for which the base learning algorithm is an open parameter, and the *specific Yarowsky algorithm*, which includes a specification of the base learner. Informally, we call the generic algorithm the "outer loop" and the base learner the "inner loop" of the specific Yarowsky algorithm. The base learner that Yarowsky assumes is a decision list induction algorithm. We postpone discussion of it until section 3.

2.2 An Objective Function

Machine learning algorithms are typically designed to optimize some objective function that represents a formal measure of performance. The maximum likelihood criterion is the most commonly used objective function. Suppose we have a set of examples Λ with labels $Y(x)$ for $x \in \Lambda$, and a parametric family of models π_θ such that $\pi(j|x;\theta)$ represents the probability of assigning label j to example x, according to the model. The likelihood of θ is the probability of the full data set according to the model, viewed as a function of θ, and the maximum likelihood criterion instructs us to choose the parameter settings $\hat{\theta}$ that maximize likelihood, or equivalently, log likelihood:

$$\begin{aligned} l(\theta) &= \log \prod_{x \in \Lambda} \pi(Y(x)|x;\theta) \\ &= \sum_{x \in \Lambda} \log \pi(Y(x)|x;\theta) \\ &= \sum_{x \in \Lambda} \sum_{j} [\![j = Y(x)]\!] \log \pi(j|x;\theta) \end{aligned}$$

(The notation "$[\![\Phi]\!]$" represents the truth value of the proposition Φ; it is 1 if Φ is true and 0 otherwise.)

Let us define:

$$\phi_x(j) = [\![j = Y(x)]\!]$$

Note that ϕ_x satisfies the formal requirements of a probability distribution over labels j: specifically, it is a point distribution with all its mass concentrated on $Y(x)$. We call it the *labeling distribution.* Now we can write:

$$\begin{aligned} l(\theta) &= \sum_{x \in \Lambda} \sum_{j} \phi_x(j) \log \pi(j|x;\theta) \\ &= -\sum_{x \in \Lambda} H(\phi_x \| \pi_x) \end{aligned} \tag{2}$$

In (2) we have written π_x for the distribution $\pi(\cdot|x;\theta)$, leaving the dependence on θ implicit. We have also used the nonstandard notation $H(p\|q)$ for what is sometimes called "cross entropy". It is easy to verify that:

$$H(p\|q) = H(p) + D(p\|q)$$

where $H(p)$ is the entropy of p and D is Kullbach-Leibler divergence. Note that, when p is a point distribution, $H(p) = 0$ and hence $H(p\|q) = D(p\|q)$. In particular:

$$l(\theta) = -\sum_{x \in \Lambda} D(\phi_x \| \pi_x) \tag{3}$$

Thus when, as here, ϕ_x is a point distribution, we can restate the maximum likelihood criterion as instructing us to choose the model that minimizes the

total divergence between the empirical labeling distributions ϕ_x and the model's prediction distributions π_x.

To extend $l(\theta)$ to unlabeled examples, we need only observe that unlabeled examples are ones about whose labels the data provide no information. Accordingly, we revise the definition of ϕ_x to treat unlabeled examples as ones whose labeling distribution is the maximally uncertain distribution, which is to say, the uniform distribution:

$$\phi_x(j) = \begin{cases} [\![j = Y(x)]\!] & \text{for } x \in \Lambda \\ \frac{1}{L} & \text{for } x \in \mathrm{V} \end{cases}$$

where L is the number of labels. Equivalently:

$$\phi_x(j) = [\![x \in Y_j]\!] + [\![x \in \mathrm{V}]\!]\frac{1}{L}$$

With this revision, however, the expressions (2) and (3) are no longer equivalent; we must use (2). Since $H(\phi_x \| \pi_x) = H(\phi_x) + D(\phi_x \| \pi_x)$, and $H(\phi_x)$ is minimized when x is labeled, minimizing $H(\phi_x \| \pi_x)$ forces one to label unlabeled examples. On labeled examples, $H(\phi_x \| \pi_x) = D(\phi_x \| \pi_x)$, and $D(\phi_x \| \pi_x)$ is minimized when the labels of examples agree with the predictions of the model.

In short, we adopt as objective function

$$H \equiv \sum_x H(\phi_x \| \pi_x) = -l(\phi, \theta)$$

We seek to minimize H.

2.3 The Modified Algorithm Y-1

We can show that a modified version of the Yarowsky algorithm finds a local minimum of H. Two modifications are necessary.

- The labeling function Y is recomputed in each iteration as before, but with the constraint that an example once labeled stays labeled. The label may change, but a labeled example cannot become unlabeled again.
- We fix the threshold ζ at $1/L$. As a result, the only examples that remain unlabeled after the labeling step are those for which π_x is the uniform distribution. The problem with an arbitrary threshold is that it prevents the algorithm from converging to a minimum of H. A threshold that gradually decreases to $1/L$ would also address the problem, but would complicate the analysis.

The modified algorithm, Y-1, is given in table 2.

To obtain a proof, it will be necessary to make an assumption about the supervised classifier $\pi^{(t+1)}$ induced by the base learner in step (2.1). A natural assumption is that the base learner chooses $\pi^{(t+1)}$ so as to minimize $\sum_{x \in \Lambda_t} D(\phi_x^{(t)} \| \pi_x^{(t+1)})$.

(1) Given: $X, Y^{(0)}$
(2) For $t \in \{0, 1, \ldots\}$
  (2.1) Train classifier on $(\Lambda_t, Y^{(t)})$; result is $\pi^{(t+1)}$
  (2.2) For each example x:
    (2.2.1) Set $\hat{y} = \arg\max_j \pi_x^{(t+1)}(j)$
    (2.2.2) Set $Y^{(t+1)}(x) = \begin{cases} \hat{y} & \text{if } x \in \Lambda_t \vee \pi_x^{(t+1)}(\hat{y}) > 1/L \\ \emptyset & \text{otherwise} \end{cases}$
  (2.3) If $Y^{(t+1)} = Y^{(t)}$, stop

Table 2: The modified generic Yarowsky algorithm (Y-1)

A weaker assumption will suffice, however. We assume that the base learner reduces divergence, if possible. That is, we assume:

$$\Delta D_\Lambda \equiv \sum_{x \in \Lambda_t} D(\phi_x^{(t)} \| \pi_x^{(t+1)}) - \sum_{x \in \Lambda_t} D(\phi_x^{(t)} \| \pi_x^{(t)}) \leq 0 \tag{4}$$

with equality only if there is no classifier $\pi^{(t+1)} \in \Pi$ that makes $\Delta D_\Lambda < 0$. Note that any learning algorithm that minimizes $\sum_{x \in \Lambda_t} D(\phi_x^{(t)} \| \pi_x^{(t+1)})$ satisfies the weaker assumption (4), inasmuch as the option of setting $\pi_x^{(t+1)} = \pi_x^{(t)}$ is always available.

We also consider a somewhat stronger assumption, namely, that the base learner reduces divergence over all examples, not just labeled examples:

$$\Delta D_X \equiv \sum_{x \in X} D(\phi_x^{(t)} \| \pi_x^{(t+1)}) - \sum_{x \in X} D(\phi_x^{(t)} \| \pi_x^{(t)}) \leq 0 \tag{5}$$

We can now state the main theorem of this section.

Theorem 1 *If the base learning algorithm satisfies (4) or (5), algorithm Y-1 decreases H at each iteration until it reaches a critical point of H.*

There is a lemma that we require in order to prove the theorem.

Lemma 1 *For all distributions p*

$$H(p) \geq \log \frac{1}{\max_j p(j)}$$

with equality iff p is the uniform distribution.

Proof. By definition, for all k:

$$\begin{aligned} p(k) &\leq \max_j p(j) \\ \log \frac{1}{p(k)} &\geq \log \frac{1}{\max_j p(j)} \end{aligned}$$

Since this is true for all k, it is true if we take the expectation with respect to p:

$$\begin{aligned} \sum_k p(k) \log \frac{1}{p(k)} &\geq \sum_k p(k) \log \frac{1}{\max_j p(j)} \\ H(p) &\geq \log \frac{1}{\max_j p(j)} \end{aligned}$$

We have equality only if $p(k) = \max_j p(j)$ for all k, that is, only if p is the uniform distribution. □

We now prove the theorem.

Proof of theorem 1. The algorithm produces a sequence of labelings $\phi^{(0)}, \phi^{(1)}, \ldots$ and a sequence of classifiers $\pi^{(1)}, \pi^{(2)}, \ldots$. The classifier $\pi^{(t+1)}$ is trained on $\phi^{(t)}$, and the labeling $\phi^{(t+1)}$ is created using $\pi^{(t+1)}$.

Recall that:

$$H = \sum_x [H(\phi_x) + D(\phi_x \| \pi_x)]$$

In the training step (2.1), we hold ϕ fixed and change π, and in the labeling step (2.2), we hold π fixed and change ϕ. We will show that the training step minimizes H as a function of π, and the labeling step minimizes H as a function of ϕ except on examples where it is at a critical point of H. Hence, H is non-increasing in each iteration of the algorithm, and is strictly decreasing unless $(\phi^{(t)}, \pi^{(t)})$ is a critical point of H.

Let us consider the labeling step first. In this step, π is held constant, but ϕ (possibly) changes, and we have:

$$\Delta H = \sum_x \Delta H(x)$$

where

$$\Delta H(x) \equiv H(\phi_x^{(t+1)} \| \pi_x^{(t+1)}) - H(\phi_x^{(t)} \| \pi_x^{(t+1)})$$

We can show that ΔH is nonpositive if we can show that $\Delta H(x)$ is nonpositive for all x.

We can guarantee that $\Delta H(x) \leq 0$ if $\phi^{(t+1)}$ minimizes $H(p \| \pi_x^{(t+1)})$ viewed as a function of p. By definition:

$$H(p \| \pi_x^{(t+1)}) = \sum_j p_j \log \frac{1}{\pi_x^{(t+1)}(j)}$$

We wish to find the distribution p that minimizes $H(p \| \pi_x^{(t+1)})$. Clearly, we accomplish that by placing all the mass of p in p_{j*} where $j*$ minimizes $-\log \pi_x^{(t+1)}(j)$. If there is more than one minimizer, $H(p \| \pi_x^{(t+1)})$ is minimized by any distribution p that distributes all its mass among the minimizers of $-\log \pi_x^{(t+1)}(j)$.

Observe further that

$$\begin{aligned} \arg\min_j \log \frac{1}{\pi_x^{(t+1)}(j)} &= \arg\max_j \pi_x^{(t+1)}(j) \\ &= \hat{y} \end{aligned}$$

That is, we minimize $H(p\|\pi_x^{(t+1)})$ by setting $p_j = [\![j = \hat{y}]\!]$, which is to say, by labeling x as predicted by $\pi^{(t+1)}$. That is how algorithm Y-1 defines $\phi_x^{(t+1)}$ for all examples $x \in \Lambda_{t+1}$.

Note that $\phi_x^{(t+1)}$ does not minimize $H(p\|\pi_x^{(t+1)})$ for examples $x \in \mathrm{V}_{t+1}$, that is, for examples x that remain unlabeled at $t+1$. However, in algorithm Y-1, any example that is unlabeled at $t+1$ is necessarily also unlabeled at t, so for any such example, $\Delta H(x) = 0$. Hence, if any label changes in the labeling step, H decreases, and if no label changes, H remains unchanged; in either case, H does not increase.

We can show further that even for examples $x \in \mathrm{V}_{t+1}$, the labeling distribution $\phi_x^{(t+1)}$ assigned by Y-1 represents a critical point of H. For any example $x \in \mathrm{V}_{t+1}$, the prediction distribution $\pi_x^{(t+1)}$ is the uniform distribution (otherwise Y-1 would have labeled x). Hence the divergence between $\phi^{(t+1)}$ and $\pi^{(t+1)}$ is zero, and thus at a minimum. It would be possible to decrease $H(\phi_x^{(t+1)}\|\pi_x^{(t+1)})$ by decreasing $H(\phi_x^{(t+1)})$ at the cost of an increase in $D(\phi_x^{(t+1)}\|\pi_x^{(t+1)})$, but all directions of motion (all ways of selecting labels to receive increased probability mass) are equally good. That is to say, the gradient of H is zero; we are at a critical point.

Essentially, we have reached a saddle point. We have minimized H with respect to ϕ_{xj} along those dimensions with a non-zero gradient. Along the remaining dimensions, we are actually at a local maximum, but without a gradient to choose a direction of descent.

Now let us consider the training step (2.1). In this step, ϕ is held constant, so the change in H is equal to the change in D—recall that $H(\phi\|\pi) = H(\phi) + D(\phi\|\pi)$. By the hypothesis of the theorem, there are two cases: either the base learner satisfies (4) or (5). If it satisfies (5), the base learner minimizes D as a function of π, hence it follows immediately that it minimizes H as a function of π.

Suppose instead that the base learner satisfies (4). We can express H as

$$H = \sum_x H(\phi_x) + \sum_{x \in \Lambda_t} D(\phi_x\|\pi_x) + \sum_{x \in \mathrm{V}_t} D(\phi_x\|\pi_x)$$

In the training step, the first term remains constant. The second term decreases, by hypothesis. But the third term may increase. However, we can show that any increase in the third term is more than offset in the labeling step.

Consider an arbitrary example in V_t. Since it is unlabeled at time t, we know that $\phi_x^{(t)}$ is the uniform distribution u:

$$u(j) = \frac{1}{L}$$

Moreover, $\pi_x^{(t)}$ must also be the uniform distribution; otherwise example x would have been labeled in a previous iteration. Therefore the value of $H(x) = H(\phi_x \| \pi_x)$ at the beginning of iteration t is H_0:

$$H_0 = \sum_j \phi_x^{(t)}(j) \log \frac{1}{\pi_x^{(t)}(j)} = \sum_j u(j) \log \frac{1}{u(j)} = H(u)$$

After the training step, the value is H_1:

$$H_1 = \sum_j \phi_x^{(t)}(j) \log \frac{1}{\pi_x^{(t+1)}(j)}$$

If π_x remains unchanged in the training step, then the new distribution $\pi_x^{(t+1)}$, like the old one, is the uniform distribution, and the example remains unlabeled. Hence there is no change in H, and in particular, H is non-increasing, as desired. On the other hand, if π_x does change, then the new distribution $\pi_x^{(t+1)}$ is non-uniform, and the example is labeled in the labeling step. Hence the value of $H(x)$ at the end of the iteration, after the labeling step, is H_2:

$$H_2 = \sum_j \phi_x^{(t+1)}(j) \log \frac{1}{\pi_x^{(t+1)}(j)} = \log \frac{1}{\pi_x^{(t+1)}(\hat{y})}$$

By lemma 1, $H_2 < H(u)$; hence $H_2 < H_0$.

As we observed above, $H_1 > H_0$, but if we consider the change overall, we find that the increase in the training step is more than offset in the labeling step:

$$\Delta H(x) = H_2 - H_1 + H_1 - H_0 < 0$$

□

3 The Specific Yarowsky Algorithm

3.1 The Original Decision List Induction Algorithm DL-0

When one speaks of the Yarowsky algorithm, one often has in mind not just the generic algorithm Y-0 (or Y-1), but an algorithm whose specification includes the particular choice of base learning algorithm made by Yarowsky. Specifically, Yarowsky's base learner constructs a decision list, that is, a list of rules of form $f \to j$, where f is a feature and j is a label, with score θ_{fj}. A rule $f \to j$ matches example x if x possesses the feature f. The label predicted for a given example x is the label of the highest-scoring rule that matches x.

Yarowsky uses smoothed precision for rule scoring. As the name suggests, smoothed precision $\tilde{q}_f(j)$ is a smoothed version of (raw) precision $q_f(j)$, which is the probability that rule $f \to j$ is correct given that it matches:

$$q_f(j) \equiv \begin{cases} |X_f Y_j| / |X_f \Lambda| & \text{if } |X_f \Lambda| > 0 \\ 1/L & \text{otherwise} \end{cases} \tag{6}$$

where X_f is the set of examples that possess feature f. Note that the product notation "$X_f Y_j$" here represents the intersection of the sets X_f and Y_j.

Smoothed precision $\tilde{q}(j|f;\epsilon)$ is defined as follows:

$$\tilde{q}(j|f;\epsilon) \equiv \frac{|X_f Y_j| + \epsilon}{|X_f \Lambda| + L\epsilon} \tag{7}$$

We also write $\tilde{q}_f(j)$ when ϵ is clear from context.

Yarowsky's update rule uses smoothed precision:

$$\theta_{fj} = \tilde{q}_f(j)$$

We will also consider raw precision as an alternative: $\theta_{fj} = q_f(j)$. Both raw and smoothed precision have the properties of a conditional probability distribution. Generally, we view θ_{fj} as a conditional distribution over labels j for a fixed feature f.

We can define

$$\pi_x(j) \propto \arg\max_{f \in F_x} \theta_{fj} \tag{8}$$

where F_x is the set of features of x. (Note that $f \in F_x$ if and only if $x \in X_f$.) Since the classifier's prediction for x is defined, in equation (1), to be the label that maximizes $\pi_x(j)$, definition (8) implies that the classifier's prediction is the label of the highest-scoring rule matching x, as desired.

We have written "$\propto$" in (8) rather than "=" because maximizing θ_{fj} across $f \in F_x$ for each label j will not in general yield a probability distribution over labels – though the scores will be positive and bounded, hence normalizable. Considering only the final predicted label $\hat{y}$ for a given example x, the normalization will have no effect, inasmuch as all scores θ_{fj} being compared will be scaled in the same way.

As characterized by Yarowsky, a decision list contains only those rules $f \to j$ whose score $\tilde{q}_f(j)$ exceeds the labeling threshold ζ. This can be seen purely as an efficiency measure. Including rules whose score falls below the labeling threshold will have no effect on the classifier's predictions, as the threshold will be applied when the classifier is applied to examples. For this reason, we do not prune the list. That is, we represent a decision list as a set of parameters $\{\theta_{fj}\}$, one for every possible rule $f \to j$ in the cross product of the set of features and the set of labels.

The decision-list induction algorithm used by Yarowsky is summarized in table 3; we call it DL-0. Note that the step labeled "(*)" is not actually a step of the induction algorithm, but rather specifies how the decision list is used to compute a prediction distribution π_x for a given example x.

Unfortunately, we cannot show that the algorithm DL-0 satisfies equation (4). A large part of the problem is the definition of π given in (8). As mentioned just above, the parameters θ_{fj} can be thought of as defining a prediction distribution $\theta_f(j)$ over labels j for each feature f. Hence the equation (8) specifies how the prediction distributions θ_f for the features of example x are to be combined to yield a prediction distribution π_x for x. Analysis of the algorithm

(0) Given: a fixed value for $\epsilon > 0$
Initialize $C[f, j] = 0$, $Z[f] = 0$ for all f, j
(1) For each example-label pair (x, j)
(1.1) Increment $C[f, j]$, $Z[f]$
(2) For each feature f and label j
(2.1) Set $\theta_{fj} = \frac{C[f,j]+\epsilon}{Z[f]+L\epsilon}$
(*) Define $\pi_x(j) \propto \arg\max_{f \in F_x} \theta_{fj}$

Table 3: The decision list induction algorithm DL-0

becomes more manageable if, instead of combining distributions by maximizing θ_f across $f \in F_x$ as in equation 8, we take a mixture of the θ_f:

$$\pi_x(j) = \frac{1}{m} \sum_{f \in F_x} \theta_{fj} \tag{9}$$

Here $m = |F_x|$ is the number of features that x possesses; for sake of simplicity, we assume that all examples have the same number of features. Since, for each f, θ_f is a probability distribution, any convex combination of distributions θ_f is also a distribution, including in particular π_x.

The two definitions for $\pi_x(j)$, (8) and (9), will often have the same mode $\hat{y}$, but that is guaranteed only in the rather severely restricted case of two features and two labels. Under definition (8), the prediction is determined entirely by the strongest θ_f, whereas definition (9) permits a bloc of weaker θ_f to outvote the strongest one. Yarowsky explicitly wished to avoid the possibility of such interactions. Nonetheless, we henceforth assume definition (9).

3.2 The Decision List Induction Algorithm DL-EM

Even adopting (9) in place of (8), we cannot show that DL-0 satisfies (4). We will be able to prove something weaker about some close variants of DL-0, but before we do so, we present an alternative decision list induction algorithm that does satisfy (4). The alternative algorithm, DL-EM, is a special case of the Expectation-Maximization (EM) algorithm. We consider two versions of the algorithm: DL-EM-Λ and DL-EM-X. DL-EM-Λ is somewhat simpler, and we describe it first.

Unlike DL-0, which constructs a classifier "from scratch," DL-EM seeks to improve on a previous classifier. (In the context of the Yarowsky algorithm, the previous classifier is the one from the previous iteration.) We write θ_{fj}^{old} for the parameters and π_x^{old} for the prediction distributions of the previous classifier.

Conceptually, DL-EM-Λ considers the label j assigned to an example x to be generated by choosing a feature $f \in F_x$ and then assigning the label j according to the feature's prediction distribution $\theta_f(j)$. The choice of feature f is a hidden variable. The degree to which an example labeled j is imputed to feature f is

(0) Initialize $C[f, j] = 0$ for all f, j
(1) For each example x labeled j
  (1.1) Let $Z = \sum_{g \in F_x} \theta^{\text{old}}_{gj}$
  (1.2) For each $f \in F_x$, increment $C[f, j]$ by $\frac{1}{Z}\theta^{\text{old}}_{fj}$
(2) For each feature f
  (2.1) Let $Z = \sum_j C[f, j]$
  (2.2) For each label j, set $\theta_{fj} = \frac{1}{Z}C[f, j]$

Table 4: DL-EM-Λ decision list induction algorithm

determined by the old distribution:

$$\pi^{\text{old}}(f|x, j) = \frac{[\![f \in F_x]\!]\theta^{\text{old}}_{fj}}{\sum_g [\![g \in F_x]\!]\theta^{\text{old}}_{gj}} = \frac{[\![f \in F_x]\!]\frac{1}{m}\theta^{\text{old}}_{fj}}{\pi^{\text{old}}_x(j)}$$

One can think of $\pi^{\text{old}}(f|x, j)$ either as the posterior probability that feature f was responsible for the label j, or as the portion of the labeled example (x, j) that is imputed to feature f. We also write $\pi^{\text{old}}_{xj}(f)$ as a synonym for $\pi^{\text{old}}(f|x, j)$. The new estimate θ_{fj} is obtained by summing imputed occurrences of (f, j) and normalizing across labels:

$$\theta_{fj} = \frac{\sum_{x \in Y_j} \pi^{\text{old}}(f|x, j)}{\sum_k \sum_{x \in Y_k} \pi^{\text{old}}(f|x, k)}$$

The algorithm is summarized in table 4.

The second version of the algorithm, DL-EM-X is summarized in table 5. It is like DL-EM-Λ, except that it uses the update rule:

$$\theta_{fj} = \frac{\sum_{x \in Y_j} \pi^{\text{old}}(f|x, j) + \frac{1}{L}\sum_{x \in \text{V}} \pi^{\text{old}}(f|x, j)}{\sum_k \left[\sum_{x \in Y_k} \pi^{\text{old}}(f|x, k) + \frac{1}{L}\sum_{x \in \text{V}} \pi^{\text{old}}(f|x, k)\right]} \tag{10}$$

The update rule (10) includes unlabeled examples as well as labeled examples. Conceptually, it divides each unlabeled example equally among the labels, then divides the resulting fractional labeled example among the example's features.

We note that both variants of the DL-EM algorithm constitute a single iteration of an EM-like algorithm. A single iteration suffices to prove the following theorem, though multiple iterations would also be effective.

Theorem 2 *The classifier produced by the DL-EM-Λ algorithm satisfies equation (4), and the classifier produced by the DL-EM-X algorithm satisfies equation (5).*

Combining theorems 1 and 2 yields:

Corollary 1 *The Yarowsky algorithm Y-1, using DL-EM-Λ or DL-EM-X as its base learning algorithm, decreases H at each iteration until it reaches a critical point of H.*

(0) Initialize $C[f,j]=0$ and $U[f,j]=0$, for all f, j
(1) For each example x labeled j
  (1.1) Let $Z=\sum_{g\in F_x}\theta^{\text{old}}_{gj}$
  (1.2) For each $f\in F_x$, increment $C[f,j]$ by $\frac{1}{Z}\theta^{\text{old}}_{fj}$
(2) For each unlabeled example x
  (2.1) Let $Z=\sum_{g\in F_x}\theta^{\text{old}}_{gj}$
  (2.2) For each $f\in F_x$, increment $U[f,j]$ by $\frac{1}{Z}\theta^{\text{old}}_{fj}$
(3) For each feature f
  (3.1) Let $Z=\sum_j(C[f,j]+\frac{1}{L}U[f,j])$
  (3.2) For each label j, set $\theta_{fj}=\frac{1}{Z}\left(C[f,j]+\frac{1}{L}U[f,j]\right)$

Table 5: DL-EM-X decision list induction algorithm

Proof of theorem 2. Let θ^{old} represent the parameter values at the beginning of the call to DL-EM, let θ represent a family of free variables that we will optimize, and let π^{old} and π be the corresponding prediction distributions. The labeling distribution ϕ is fixed. For any set of examples α, let ΔD_α be the change in $\sum_{x\in\alpha}D(\phi_x\|\pi_x)$ resulting from the change in θ. We are obviously particularly interested in the two cases where α is the set of all examples X (for DL-EM-X) and where α is the set of labeled examples Λ (for DL-EM-Λ). In either case, $\Delta D_\alpha\le 0$.

Observe that:

$$\begin{aligned}
-\Delta D_\alpha &= \sum_{x\in\alpha}\left[D(\phi_x\|\pi^{\text{old}}_x)-D(\phi_x\|\pi_x)\right]\\
&= \sum_{x\in\alpha}\left[H(\phi_x\|\pi^{\text{old}}_x)-H(\phi_x)-H(\phi_x\|\pi_x)+H(\phi_x)\right]\\
&= \sum_{x\in\alpha}\sum_j\phi_x(j)\left[\log\pi_x(j)-\log\pi^{\text{old}}_x(j)\right] \qquad (11)
\end{aligned}$$

The EM algorithm is based on the fact that divergence is non-negative, and strictly positive if the distributions compared are not identical.

$$\begin{aligned}
0 &\le \sum_j\sum_{x\in\alpha}\phi_x(j)D(\pi^{\text{old}}_{xj}\|\pi_{xj})\\
&= \sum_j\sum_{x\in\alpha}\phi_x(j)\sum_{f\in F_x}\pi^{\text{old}}_{xj}(f)\log\frac{\pi^{\text{old}}_{xj}(f)}{\pi_{xj}(f)}\\
&= \sum_j\sum_{x\in\alpha}\phi_x(j)\sum_{f\in F_x}\pi^{\text{old}}_{xj}(f)\log\left(\frac{\theta^{\text{old}}_{fj}}{\pi^{\text{old}}_x(j)}\cdot\frac{\pi_x(j)}{\theta_{fj}}\right)
\end{aligned}$$

Which yields the inequality:

$$\sum_j\sum_{x\in\alpha}\phi_x(j)\left[\log\pi_x(j)-\log\pi^{\text{old}}_x(j)\right]$$

$$\geq \sum_j \sum_{x \in \alpha} \phi_x(j) \sum_{f \in F_x} \pi_{xj}^{\text{old}}(f) \left[\log \theta_{fj} - \log \theta_{fj}^{\text{old}} \right]$$

By (11), this can be written as:

$$-\Delta D_\alpha \geq \sum_j \sum_{x \in \alpha} \phi_x(j) \sum_{f \in F_x} \pi_{xj}^{\text{old}}(f) \left[\log \theta_{fj} - \log \theta_{fj}^{\text{old}} \right] \tag{12}$$

Since θ_{fj}^{old} is constant, by maximizing

$$\sum_j \sum_{x \in \alpha} \phi_x(j) \sum_{f \in F_x} \pi_{xj}^{\text{old}}(f) \log \theta_{fj} \tag{13}$$

we maximize a lower bound on $-\Delta D_\alpha$. It is easy to see that $-\Delta D_\alpha$ is bounded above by 0: we simply set $\theta_{fj} = \theta_{fj}^{\text{old}}$. Since divergence is zero only if the two distributions are identical, we have strict inequality in (12) unless the best choice for θ is θ^{old}, in which case no choice of θ makes $\Delta D_\alpha < 0$.

It remains to show that DL-EM computes the parameter set θ that maximizes (13). We wish to maximize (13) under the constraints that the values $\{\theta_{fj}\}$ for fixed f sum to unity across choices of j. We express the constraints in the form:

$$C_f = 0$$

where

$$C_f \equiv \sum_j \theta_{fj} - 1$$

We seek a solution to the family of equations:

$$\frac{\partial}{\partial \theta_{fj}} \sum_k \sum_{x \in \alpha} \phi_x(k) \sum_{g \in F_x} \pi_{xk}^{\text{old}}(g) \log \theta_{gk} = \lambda_f \frac{\partial C_f}{\partial \theta_{fj}} \tag{14}$$

We derive an expression for the derivative on the left hand side:

$$\frac{\partial}{\partial \theta_{fj}} \sum_k \sum_{x \in \alpha} \phi_x(k) \sum_{g \in F_x} \pi_{xk}^{\text{old}}(g) \log \theta_{gk} = \sum_{x \in X_f \alpha} \phi_x(j) \pi_{xj}^{\text{old}}(f) \frac{1}{\theta_{fj}}$$

Similarly for the right hand side:

$$\frac{\partial C_f}{\partial \theta_{fj}} = 1$$

Substituting these into equation (14):

$$\sum_{x \in X_f \alpha} \phi_x(j) \pi_{xj}^{\text{old}}(f) \frac{1}{\theta_{fj}} = \lambda_f$$

$$\theta_{fj} = \sum_{x \in X_f \alpha} \phi_x(j) \pi_{xj}^{\text{old}}(f) \frac{1}{\lambda_f} \tag{15}$$

Using the constraint $C_f = 0$ and solving for λ_f:

$$\sum_j \sum_{x \in X_f \alpha} \phi_x(j) \pi_{xj}^{\text{old}}(f) \frac{1}{\lambda_f} - 1 = 0$$

$$\lambda_f = \sum_j \sum_{x \in X_f \alpha} \phi_x(j) \pi_{xj}^{\text{old}}(f)$$

Substituting back into (15):

$$\theta_{fj} = \frac{\sum_{x \in X_f \alpha} \phi_x(j) \pi_{xj}^{\text{old}}(f)}{\sum_k \sum_{x \in X_f \alpha} \phi_x(k) \pi_{xk}^{\text{old}}(f)} \tag{16}$$

If we consider the case where α is the set of all examples, and expand out ϕ_x in (16), we obtain:

$$\theta_{fj} = \frac{1}{Z} \left[\sum_{x \in X_f Y_j} \pi_{xj}^{\text{old}}(f) + \frac{1}{L} \sum_{x \in X_f \mathrm{V}} \pi_{xj}^{\text{old}}(f) \right]$$

where Z normalizes θ_f. It is not hard to see that this is the update rule that DL-EM-X computes, using the intermediate values:

$$\begin{aligned} C[f,j] &= \sum_{x \in X_f Y_j} \pi_{xj}^{\text{old}}(f) \\ U[f,j] &= \sum_{x \in X_f \mathrm{V}} \pi_{xj}^{\text{old}}(f) \end{aligned}$$

If we consider the case where α is the set of labeled examples, and expand out ϕ_x in (16), we obtain:

$$\theta_{fj} = \frac{1}{Z} \sum_{x \in X_f Y_j} \pi_{xj}^{\text{old}}(f)$$

This is the update rule that DL-EM-Λ computes. Thus we see that DL-EM-X reduces D_X, and DL-EM-Λ reduces D_Λ. □

We note in closing that DL-EM-X can be simplified when used with algorithm Y-1, inasmuch as it is known that $\theta_{fj} = 1/L$ for all (f, j) where $f \in F_x$ for some $x \in \mathrm{V}$. Then the expression for $U[f, j]$ simplifies as follows:

$$\begin{aligned} & \sum_{x \in X_f \mathrm{V}} \pi_{xj}^{\text{old}}(f) \\ &= \sum_{x \in X_f \mathrm{V}} \left[\frac{1/L}{\sum_{g \in F_x} 1/L} \right] \\ &= \frac{|X_f \mathrm{V}|}{m} \end{aligned}$$

The dependence on j disappears, so we can replace $U[f, j]$ with $U[f]$ in algorithm DL-EM-X, delete step (2.1), and replace step (2.2) with the statement "For each $f \in F_x$, increment $U[f]$ by $1/m$".

3.3 The Objective Function K

The decision list induction algorithm DL-EM is substantially different from the one that Yarowsky uses, DL-0. We are not able to show that DL-0 optimizes the objective function H. However, in this section, we define a different (but related) objective function K, and in the next section, we show that modified versions of DL-0 optimize it.

The objective function K is an upper bound on H. Observe that:

$$\begin{aligned} H &= -\sum_x \sum_j \phi_{xj} \log \sum_{g \in F_x} \frac{1}{m} \theta_{gj} \\ &\leq -\sum_x \sum_j \phi_{xj} \sum_{g \in F_x} \frac{1}{m} \log \theta_{gj} \\ &= \frac{1}{m} \sum_x \sum_{g \in F_x} H(\phi_x \| \theta_g) \end{aligned}$$

We define:

$$K \equiv \sum_x \sum_{g \in F_x} H(\phi_x \| \theta_g) \tag{17}$$

By minimizing K, we minimize an upper bound on H. Moreover, it is in principle possible to reduce K to zero. Since $H(\phi_x \| \theta_g) = H(\phi_x) + D(\phi_x \| \theta_g)$, K is reduced to zero if all examples are labeled, each feature concentrates its prediction distribution in a single label, and the label of every example agrees with the prediction of every feature it possesses. In this limiting case, any minimizer of K is also a minimizer of H.

We hasten to add a proviso: it is not possible to reduce K to zero for all data sets. The following provides a necessary and sufficient condition. Consider an undirected bipartite graph G whose nodes are examples and features. There is an edge between example x and feature f just in case f is a feature of x. Define examples x_1 and x_2 to be neighbors if they both belong to the same connected component of G. K is reducible to zero if and only if all pairs of neighbors in Λ_0 are labeled consistently.

3.4 Algorithm DL-1

We consider two variants of DL-0, called DL-1-R and DL-1-VS. They differ from DL-0 in two ways. First, the DL-1 algorithms assume the "mean" definition of π_x given in equation (9) rather than the "max" definition of equation (8). This is not actually a difference in the induction algorithm itself, but in the way the decision list is used to construct a prediction distribution π_x.

Second, the DL-1 algorithms use update rules that differ from the smoothed precision of DL-0. DL-1-R (table 6) uses raw precision instead of smoothed precision. DL-1-VS (table 7) uses smoothed precision, but unlike DL-0, DL-1-VS does not used a fixed smoothing constant ϵ; rather ϵ varies from feature to feature.

(0) Initialize $C[f,j] = 0$, $Z[f] = 0$ for all f, j
(1) For each example-label pair (x, j)
  (1.1) For each feature $f \in F_x$, increment $C[f,j]$, $Z[f]$
(2) For each feature f and label j
  (2.1) Set $\theta_{fj} = \frac{C[f,j]}{Z[f]}$
(*) Define $\pi_x(j) = \frac{1}{m} \sum_{f \in F_x} \theta_{fj}$

Table 6: The decision list induction algorithm DL-1-R

(0) Initialize $C[f,j] = 0$, $Z[f] = 0$, $U[f] = 0$ for all f, j
(1) For each example-label pair (x, j)
  (1.1) For each feature $f \in F_x$, increment $C[f,j]$, $Z[f]$
(2) For each unlabeled example x
  (2.1) For each feature $f \in F_x$, increment $U[f]$
(3) For each feature f and label j
  (3.1) Set $\epsilon = U[f]/L$
  (3.2) Set $\theta_{fj} = \frac{C[f,j]+\epsilon}{Z[f]+U[f]}$
(*) Define $\pi_x(j) = \frac{1}{m} \sum_{f \in F_x} \theta_{fj}$

Table 7: The decision list induction algorithm DL-1-VS

The value of ϵ used by DL-1-VS can be expressed in another way that will prove useful. Let us define:

$$p(\Lambda|f) \equiv \frac{|X_f \Lambda|}{|X_f|}$$

$$p(\mathrm{V}|f) \equiv \frac{|X_f \mathrm{V}|}{|X_f|}$$

Lemma 2 *The parameter values $\{\theta_{fj}\}$ computed by DL-1-VS can be expressed as:*

$$\theta_{fj} = p(\Lambda|f) q_f(j) + p(\mathrm{V}|f) u(j) \tag{18}$$

Proof. First we show that smoothed precision can be expressed as a convex combination of raw precision (6) and the uniform distribution. Define $\delta = \epsilon/|X_f \Lambda|$. Then:

$$\begin{aligned} \tilde{q}_f(j) &= \frac{|X_f Y_j| + \epsilon}{|X_f \Lambda| + L\epsilon} \\ &= \frac{|X_f Y_j|/|X_f \Lambda| + \delta}{1 + L\delta} \\ &= \frac{1}{1 + L\delta} q_f(j) + \frac{L\delta}{1 + L\delta} \cdot \frac{\delta}{L\delta} \end{aligned}$$

$$= \frac{1}{1+L\delta} q_f(j) + \frac{L\delta}{1+L\delta} u(j) \tag{19}$$

Now we show that the mixing coefficient $1/(1+L\delta)$ of (19) is the same as the mixing coefficient $p(\Lambda|f)$ of the lemma, when $\epsilon = |X_f \mathrm{V}|/L$ as in step (3.1) of DL-1-VS.

$$\begin{aligned} \epsilon &= \frac{|X_f \mathrm{V}|}{L} \\ &= \frac{|X_f \Lambda|}{L} \cdot \frac{p(\mathrm{V}|f)}{p(\Lambda|f)} \\ L\delta &= \frac{1}{p(\Lambda|f)} - 1 \\ \frac{1}{1+L\delta} &= p(\Lambda|f) \end{aligned}$$

□

The main theorem of this section (corollary 2) is that the specific Yarowsky algorithm Y-1/DL-1 decreases K in each iteration until it reaches a critical point. It is stated as a corollary of two theorems. The first (theorem 3) shows that DL-1 minimizes K as a function of θ, holding ϕ constant, and the second (theorem 4) shows that Y-1 decreases K as a function of ϕ, holding θ constant. More precisely, DL-1-R minimizes K over labeled examples Λ, and DL-1-VS minimizes K over all examples X. Either is sufficient for Y-1 to be effective.

Theorem 3 *DL-1 minimizes K as a function of θ holding ϕ constant. Specifically, DL-1-R minimizes K over labeled examples Λ, and DL-1-VS minimizes K over all examples X.*

Proof. We wish to minimize K as a function of θ under the constraints:

$$C_f \equiv \sum_j \theta_{fj} - 1 = 0$$

for each f. As before, to minimize K under the constraints $C_f = 0$, we express the gradient of K as a linear combination of the gradients of the constraints, and solve the resulting system of equations:

$$\frac{\partial K}{\partial \theta_{fj}} = \lambda_f \frac{\partial C_f}{\partial \theta_{fj}} \tag{20}$$

First we derive expressions for the derivatives of C_f and K. The variable α represents the set of examples over which we are minimizing K.

$$\frac{\partial C_f}{\partial \theta_{fj}} = 1$$

$$\begin{aligned} \frac{\partial K}{\partial \theta_{fj}} &= -\frac{\partial}{\partial \theta_{fj}} \sum_{x \in \alpha} \sum_{g \in F_x} \sum_k \phi_{xk} \log \theta_{gk} \\ &= -\sum_{x \in X_f \alpha} \phi_{xj} \frac{1}{\theta_{fj}} \end{aligned}$$

We substitute these expressions into (20) and solve for θ_{fj}:

$$-\sum_{x \in X_f \alpha} \phi_{xj} \frac{1}{\theta_{fj}} = \lambda_f$$

$$\theta_{fj} = -\sum_{x \in X_f \alpha} \phi_{xj} / \lambda_f$$

Substituting the latter expression into the equation for $C_f = 0$ and solving for f:

$$\sum_j \left(-\sum_{x \in X_f \alpha} \phi_{xj} / \lambda_f \right) = 1$$

$$-|X_f \alpha| = \lambda_f$$

Substituting this back into the expression for θ_{fj}:

$$\theta_{fj} = \frac{1}{|X_f \alpha|} \sum_{x \in X_f \alpha} \phi_{xj} \tag{21}$$

If $\alpha = \Lambda$, we have:

$$\begin{aligned} \theta_{fj} &= \frac{1}{|X_f \Lambda|} \sum_{x \in X_f \Lambda} [\![x \in Y_j]\!] \\ &= q_f(j) \end{aligned}$$

This is the update computed by DL-1-R, showing that DL-1-R computes the parameter values $\{\theta_{fj}\}$ that minimize K over the labeled examples Λ.

If $\alpha = X$, we have:

$$\begin{aligned} \theta_{fj} &= \frac{1}{|X_f|} \sum_{x \in X_f \Lambda} [\![x \in Y_j]\!] + \frac{1}{|X_f|} \sum_{x \in X_f \mathrm{V}} \frac{1}{L} \\ &= \frac{|X_f \Lambda|}{|X_f|} \cdot \frac{|X_f Y_j|}{|X_f \Lambda|} + \frac{|X_f \mathrm{V}|}{|X_f|} \cdot \frac{1}{L} \\ &= p(\Lambda | f) \cdot q_f(j) + p(\mathrm{V} | f) \cdot u(j) \end{aligned}$$

By lemma 2, this is the update computed by DL-1-VS, hence DL-1-VS minimizes K over the complete set of examples X. □

Theorem 4 *If the base learner decreases K over X or over Λ, where the prediction distribution is computed as*

$$\pi_x(j) = \frac{1}{m} \sum_{f \in F_x} \theta_{fj}$$

then algorithm Y-1 decreases K at each iteration until it reaches a critical point, considering K as a function of ϕ with θ held constant.

The proof has the same structure as the proof of theorem 1, so we give only a sketch here. We minimize K as a function of ϕ by minimizing it for each example separately.

$$\begin{aligned} K(x) &= \sum_{g \in F_x} H(\phi_x \| \theta_g) \\ &= \sum_j \phi_{xj} \sum_{g \in F_x} \log \frac{1}{\theta_{gj}} \end{aligned}$$

To minimize $K(x)$, we choose ϕ_{xj} so as to concentrate all mass in

$$\arg\min_j \sum_{g \in F_x} \log \frac{1}{\theta_{gj}} = \arg\max_j \pi_x(j)$$

This is the labeling rule used by Y-1.

If the base learner minimizes over Λ only, rather than X, it can be shown that any increase in K on unlabeled examples is compensated for in the labeling step, as in the proof of theorem 1. □

Corollary 2 *The specific Yarowsky algorithms Y-1/DL-1-R and Y-1/DL-1-VS decrease K at each iteration until they reach a critical point.*

4 Sequential Algorithms

4.1 The Family YS

The Yarowsky algorithm variants we have considered to now do "parallel" updates in the sense that the parameters $\{\theta_{fj}\}$ are completely recomputed at each iteration. In this section, we consider a family YS of "sequential" variants of the Yarowsky algorithm, in which a single feature is selected for update at each iteration. The YS algorithms resemble the "Yarowsky-Cautious" algorithm of Collins & Singer [3], though they differ from the Yarowsky-Cautious algorithm in that they update a single feature in each iteration, rather than a small set of features, as in Yarowsky-Cautious. The YS algorithms are intended to be as close to the Y-1/DL-1 algorithm as is consonant with single-feature updates. The YS algorithms differ from each other, and from Y-1/DL-1, in the choice of update rule. An interesting range of update rules work in the sequential setting. In particular, smoothed precision with fixed ϵ, as in the original algorithm Y-0/DL-0, works in the sequential setting, though with a proviso that will be spelled out later.

At each iteration there is a set of selected features S_t, beginning with an initial set S_0. One feature is selected in each iteration. A feature once selected remains in the selected set. It is permissible for a feature to be selected more than once; this permits us to continue reducing K even after all features have been selected. In short, there is a sequence of selected features $f_0, f_1, \ldots$, and

$$S_{t+1} = S_t \cup \{f_t\}$$

(0) Given: S, θ with $\theta_{gj} = 1/L$ for $g \notin S$
(1) For each example x possessing a feature in S, set $Y(x) = \hat{y}$
(2) Loop:
  (2.1) Choose a feature f such that $X_f \Lambda \neq \emptyset$ and $\theta_f \neq q_f$
    If there is none, stop
  (2.2) Add f to S
  (2.2) For each label j, set $\theta_{fj} = \text{update}(f, j)$
  (2.3) For each example x possessing a feature in S, set $Y(x) = \hat{y}$

Table 8: The sequential algorithm YS

The algorithm is also provided initially with a set of parameters $\{\theta_{fj}^{(0)}\}$, with $\theta_{fj}^{(0)} = 1/L$ if $f \notin S_0$. In iteration t, the only parameters θ_{gj} that change are those with $g = f_t$, where f_t is the selected feature at iteration t. That is:

$$\theta_{gj}^{(t+1)} = \theta_{gj}^{(t)} \quad \text{for } g \neq f_t$$

It follows that, for all t:

$$\theta_{gj}^{(t)} = \frac{1}{L} \quad \text{for } g \notin S_t$$

In each iteration, one chooses a feature f_t and computes (or recomputes) the prediction distribution θ_{f_t} for f_t. Then labels are recomputed. An example x is assigned label $\hat{y}$ if any feature of x belongs to S_{t+1}. In particular, all previously labeled examples continue to be labeled (though their labels may change), and any unlabeled examples possessing feature f_t become labeled.

The algorithm is summarized in table 8. It is actually an algorithm schema; the definition for "update" needs to be supplied. We consider three different update functions: one that uses raw precision as prediction distribution, one that uses smoothed precision, and one that goes in the opposite direction, using what we might call "peaked precision". As we have seen, smoothed precision can be expressed as a mixture of raw precision and the uniform (i.e., maximum-entropy) distribution (19). Peaked precision $\hat{q}(f)$ mixes in a certain amount of the point (i.e., minimum-entropy) distribution that has all its mass on the label that maximizes raw precision:

$$\hat{q}_f(j) \equiv p(\Lambda_t|f) q_f(j) + p(\mathrm{V}_t|f) [\![j = j\dagger]\!]$$

where:

$$j\dagger \equiv \arg\max_j q_f(j)$$

Note that peaked precision involves a variable amount of "peaking"; the mixing parameters depend on the relative proportions of labeled and unlabeled examples. Note also that $j\dagger$ is a function of f, though we do not explicitly represent that dependence.

The three instantiations of algorithm YS that we consider are:

YS-P ("peaked")	$\theta_{fj} = \hat{q}_f(j)$
YS-R ("raw")	$\theta_{fj} = q_f(j)$
YS-FS ("fixed smoothing")	$\theta_{fj} = \tilde{q}_f(j)$

We will show that the first two algorithms reduce K in each iteration. We will show that the third algorithm, YS-FS, reduces K in iterations in which f_t is a new feature, not previously selected. Unfortunately, we are unable to show that YS-FS reduces K when f_t is a previously-selected feature. This suggests using a mixed algorithm in which smoothed precision is used for new features but raw or peaked precision is used for previously-selected features.

A final issue with the algorithm schema YS concerns the selection of features in step (2.1). The schema as stated does not specify which feature is to be selected. In essence, the manner in which rules are selected does not matter, as long as one selects rules that have room for improvement, in the sense that the current prediction distribution θ_f differs from raw precision q_f. (The justification for this choice is given in corollary 4 below.) The theorems of the following sections show that K decreases in each iteration, so long as any such rule can be found.

One could choose greedily by choosing the feature that maximizes gain G (22), though in the next section we give lower bounds for G that are rather more easily computed (theorems 5 and 6). There is in addition one case of particular interest that we return to in section 4.6.

4.2 Gain

From this point on, we consider a single iteration of the YS algorithm, and discard the variable t. We write θ^{old} and ϕ^{old} for the parameter-set and labeling at the beginning of the iteration, and we write simply θ and ϕ for the new parameter-set and new labeling. The set Λ (resp., V) represents the examples that are labeled (resp., unlabeled) at the beginning of the iteration. The selected feature is f.

We wish to choose a prediction distribution for f so as to guarantee that K decreases in each iteration. The *gain* in the current iteration is:

$$G = \sum_x \sum_{g \in F_x} \left[H(\phi_x^{\text{old}} \| \theta_g^{\text{old}}) - H(\phi_x \| \theta_g) \right] \tag{22}$$

Gain is the negative change in K; it is positive when K decreases.

In considering the reduction in K from $(\phi^{\text{old}}, \theta^{\text{old}})$ to (ϕ, θ), it will be convenient to consider the following intermediate values.

$$\begin{aligned} K_0 &= \sum_x \sum_{g \in F_x} H(\phi_x^{\text{old}} \| \theta_g^{\text{old}}) \\ K_1 &= \sum_x \sum_{g \in F_x} H(\psi_x \| \theta_g^{\text{old}}) \\ K_2 &= \sum_x \sum_{g \in F_x} H(\psi_x \| \theta_g) \end{aligned}$$

$$K_3 \quad = \quad \sum_x \sum_{g \in F_x} H(\phi_x \| \theta_g)$$

where

$$\psi_{xj} = \begin{cases} [\![j = j*]\!] & \text{if } x \in X_f \mathrm{V} \\ \phi_{xj}^{\text{old}} & \text{otherwise} \end{cases}$$

and

$$j* \equiv \arg\max_j \theta_{fj}$$

One should note that:

- θ_f is the new prediction distribution for the candidate f; $\theta_{gj} = \theta_{gj}^{\text{old}}$ for $g \neq f$.
- ϕ is the new label distribution, after relabeling. It is defined as:

$$\phi_{xj} = \begin{cases} [\![j = \hat{y}(x)]\!] & \text{if } x \in \Lambda \cup X_f \\ \frac{1}{L} & \text{otherwise} \end{cases} \tag{23}$$

- For $x \in X_f \mathrm{V}$, the only selected feature at $t+1$ is f, hence $j* = \hat{y}$ for such examples. It follows that ψ and ϕ agree on examples in $X_f \mathrm{V}$. They also agree on examples that are unlabeled at $t+1$, assigning them the uniform label distribution. If ψ and ϕ differ, it is only on old labeled examples (Λ) that need to be relabeled, given the addition of f.

The gain G can be represented as the sum of three intermediate gains, corresponding to the intermediate values just defined:

$$G = G_{\mathrm{V}} + G_\theta + G_\Lambda \tag{24}$$

where:

$$\begin{aligned} G_{\mathrm{V}} &= K_0 - K_1 \\ G_\theta &= K_1 - K_2 \\ G_\Lambda &= K_2 - K_3 \end{aligned}$$

The gain G_{V} intuitively represents the gain that is attributable to labeling previously unlabeled examples in accordance with the predictions of θ. The gain G_θ represents the gain that is attributable to changing the values θ_{fj} where f is the selected feature. The gain G_Λ represents the gain that is attributable to changing the labels of previously labeled examples to make labels agree with the predictions of the new model θ. The gain G_θ corresponds to step 2.2 of algorithm YS, in which θ is changed but ϕ is held constant; and the combined G_{V} and G_Λ gains correspond to step 2.3 of algorithm YS, in which ϕ is changed while holding θ constant.

In the remainder of this section, we derive two lower bounds for G. In following sections, we show that the updates YS-P, YS-R, and YS-FS guarantee that the lower bounds given below are non-negative, hence that G is non-negative.

Lemma 3 $G_{\mathrm{V}} = 0$

Proof. We show that K remains unchanged if we substitute ψ for ϕ^{old} in K_0. The only property of ψ that we need is that it agrees with ϕ^{old} on previously labeled examples.

Since $\psi_x = \phi_x^{\text{old}}$ for $x \in \Lambda$, we need only consider examples in V. Since these examples are unlabeled at the beginning of the iteration, none of their features have been selected, hence $\theta_{gj}^{\text{old}} = 1/L$ for all their features g. Hence:

$$\begin{aligned} K_1 &= -\sum_{x \in \mathrm{V}} \sum_{g \in F_x} \sum_j \psi_{xj} \log \theta_{fj}^{\text{old}} \\ &= -\sum_{x \in \mathrm{V}} \sum_{g \in F_x} \left[\sum_j \psi_{xj} \right] \log \frac{1}{L} \\ &= -\sum_{x \in \mathrm{V}} \sum_{g \in F_x} \left[\sum_j \phi_{xj}^{\text{old}} \right] \log \frac{1}{L} \\ &= -\sum_{x \in \mathrm{V}} \sum_{g \in F_x} \sum_j \phi_{xj}^{\text{old}} \log \theta_{fj}^{\text{old}} \quad = K_0 \end{aligned}$$

This shows that $K_0 = K_1$, hence that $G_{\mathrm{V}} = 0$. □

Lemma 4 $G_\Lambda \geq 0$

We must show that relabeling old labeled examples – that is, setting $\phi_x(j) = [\![j = \hat{y}]\!]$ for $x \in \Lambda$ – does not increase K. The proof has the same structure as the proof of theorem 1, and is omitted.

Lemma 5 *G_θ is equal to*

$$|X_f \Lambda| \left[H(q_f \| \theta_f^{old}) - H(q_f \| \theta_f) \right] + |X_f \mathrm{V}| \left[\log L - \log \frac{1}{\theta_{fj*}} \right] \tag{25}$$

Proof. By definition, $G_\theta = K_1 - K_2$, and K_1 and K_2 are identical everywhere except on examples in X_f. Hence:

$$G_\theta = \sum_{x \in X_f} \sum_{g \in F_x} \left[H(\psi_x \| \theta_g^{\text{old}}) - H(\psi_x \| \theta_g) \right]$$

We divide this sum into three partial sums:

$$\begin{aligned} G_\theta &= A + B + C \qquad (26) \\ A &= \sum_{x \in X_f \Lambda} \left[H(\psi_x \| \theta_f^{\text{old}}) - H(\psi_x \| \theta_f) \right] \\ B &= \sum_{x \in X_f \mathrm{V}} \left[H(\psi_x \| \theta_f^{\text{old}}) - H(\psi_x \| \theta_f) \right] \\ C &= \sum_{x \in X_f} \sum_{g \neq f \in F_x} \left[H(\psi_x \| \theta_g^{\text{old}}) - H(\psi_x \| \theta_g) \right] \end{aligned}$$

We consider each partial sum separately.

$$\begin{aligned}
A &= \sum_{x \in X_f \Lambda} \left[H(\psi_x \| \theta_f^{\text{old}}) - H(\psi_x \| \theta_f) \right] \\
&= -\sum_{x \in X_f \Lambda} \sum_k \psi_{xk} \left[\log \theta_{fk}^{\text{old}} - \log \theta_{fk} \right] \\
&= -\sum_{x \in X_f \Lambda} \sum_k [\![x \in Y_k]\!] \left[\log \theta_{fk}^{\text{old}} - \log \theta_{fk} \right] \\
&= -\sum_k |X_f Y_k| \left[\log \theta_{fk}^{\text{old}} - \log \theta_{fk} \right] \\
&= -|X_f \Lambda| \sum_k q_f(k) \left[\log \theta_{fk}^{\text{old}} - \log \theta_{fk} \right] \\
&= |X_f \Lambda| \left[H(q_f \| \theta_f^{\text{old}}) - H(q_f \| \theta_f) \right] \qquad (27)
\end{aligned}$$

$$\begin{aligned}
B &= \sum_{x \in X_f \text{V}} \left[H(\psi_x \| \theta_f^{\text{old}}) - H(\psi_x \| \theta_f) \right] \\
&= -\sum_{x \in X_f \text{V}} \sum_k \psi_{xk} \left[\log \theta_{fk}^{\text{old}} - \log \theta_{fk} \right] \\
&= -\sum_{x \in X_f \text{V}} \sum_k [\![k = j*]\!] \left[\log \theta_{fk}^{\text{old}} - \log \theta_{fk} \right] \\
&= |X_f \text{V}| \left[\log \frac{1}{\theta_{fj*}^{\text{old}}} - \log \frac{1}{\theta_{fj*}} \right] \\
&= |X_f \text{V}| \left[\log L - \log \frac{1}{\theta_{fj*}} \right] \qquad (28)
\end{aligned}$$

The justification for the last step is a bit subtle. If f is a new feature, not previously selected, then $\theta_{fk}^{\text{old}} = 1/L$ for all k, and the substitution is valid. On the other hand, if f is a previously selected feature, then $|X_f \text{V}| = 0$, and even though the substitution of $1/L$ for $\theta_{fj*}^{\text{old}}$ may not be valid, it is innocuous.

$$\begin{aligned}
C &= \sum_{x \in X_f} \sum_{g \neq f \in F_x} \left[H(\psi_x \| \theta_g^{\text{old}}) - H(\psi_x \| \theta_g) \right] \\
&= \sum_{x \in X_f} \sum_{g \neq f \in F_x} \left[H(\psi_x \| \theta_g^{\text{old}}) - H(\psi_x \| \theta_g^{\text{old}}) \right] \\
&= 0 \qquad (29)
\end{aligned}$$

Combining (26), (27), (28) and (29) yields the lemma. □

Theorem 5 *G is bounded below by (25).*

Proof. Combining (24) with lemmas 3, 4, and 5. □

Theorem 6 *G is bounded below by*

$$|X_f\Lambda|\left[H(q_f\|\theta_f^{old}) - H(q_f\|\theta_f)\right]$$

Proof. The theorem follows immediately from theorem 5 if we can show that

$$\log L - \log\frac{1}{\theta_{fj*}} \geq 0$$

Observe first that $\log L = H(u)$. By lemma 1, we know that

$$\begin{aligned} H(u) - \log\frac{1}{\theta_{fj*}} &\geq H(u) - H(\theta_f) \\ &\geq 0 \end{aligned}$$

The latter following because the uniform distribution maximizes entropy. □

Corollary 3 *G is bounded below by*

$$|X_f\Lambda|\left[D(q_f\|\theta_f^{old}) - D(q_f\|\theta_f)\right]$$

Proof. Immediate from theorem 6 and the fact that:

$$\begin{aligned} H(q_f\|\theta_f^{\text{old}}) - H(q_f\|\theta_f) &= H(q_f) + D(q_f\|\theta_f^{\text{old}}) - H(q_f) - D(q_f\|\theta_f) \\ &= D(q_f\|\theta_f^{\text{old}}) - D(q_f\|\theta_f) \end{aligned}$$

□

Corollary 4 *If $\theta_f^{old} \neq q_f$, then there is a choice of θ_f that yields strictly positive gain.*

Proof. If $\theta_f^{\text{old}} \neq q_f$, then

$$D(q_f\|\theta_f^{\text{old}}) > 0$$

Setting $\theta_f = q_f$ has the result that

$$|X_f\Lambda|\left[D(q_f\|\theta_f^{\text{old}}) - D(q_f\|\theta_f)\right] = |X_f\Lambda|D(q_f\|\theta_f^{\text{old}}) > 0$$

Hence $G > 0$ by corollary 3. □

4.3 Algorithm YS-P

We now use the results of the previous section to show that the algorithm YS-P is correct in the sense that it reduces K in every iteration.

Theorem 7 *In each iteration of algorithm YS-P, K decreases.*

Proof. We wish to show that $G > 0$. By theorem 5, that is true if the expression (25) is positive. By corollary 4, there exist choices for θ_f that make (25) positive, hence in particular we guarantee $G > 0$ by maximizing (25). We maximize (25) by minimizing:

$$|X_f\Lambda|H(q_f\|\theta_f) + |X_f\mathrm{V}|\log\frac{1}{\theta_{fj*}} \tag{30}$$

Since

$$H(q_f\|\theta_f) = H(q_f) + D(q_f\|\theta_f)$$

we minimize (30) by minimizing:

$$|X_f\Lambda|D(q_f\|\theta_f) + |X_f\mathrm{V}|\log\frac{1}{\theta_{fj*}} \tag{31}$$

Both terms are nonnegative. The first term is 0 if $\theta_f = q_f$. The second term is 0 for any distribution that concentrates all its mass in a single label $j*$; it is symmetric in all choices of $j*$, and decreases monotonically as θ_{fj*} approaches 1. Hence, the minimum of (31) will have $j*$ equal to the mode of q_f, though it may be more peaked than q_f, at the cost of an increase in the first term, but offset by a decrease in the second term.

Recall that $j\dagger = \arg\max_j q_f(j)$. By the reasoning of the previous paragraph, we know that $j\dagger = j*$ at the minimum of (31). Hence we can minimize (31) by minimizing:

$$|X_f\Lambda|D(q_f\|\theta_f) - |X_f\mathrm{V}|\sum_k [\![k = j\dagger]\!]\log\theta_{fk} \tag{32}$$

We compute the gradient:

$$\begin{aligned}
&\frac{\partial}{\partial\theta_{fj}}\left[|X_f\Lambda|D(q_f\|\theta_f) - |X_f\mathrm{V}|\sum_k [\![k = j\dagger]\!]\log\theta_{fk}\right] \\
&= \frac{\partial}{\partial\theta_{fj}}\left[|X_f\Lambda|H(q_f\|\theta_f) - |X_f\Lambda|H(q_f) - |X_f\mathrm{V}|\sum_k [\![k = j\dagger]\!]\log\theta_{fk}\right] \\
&= \frac{\partial}{\partial\theta_{fj}}|X_f\Lambda|H(q_f\|\theta_f) - \frac{\partial}{\partial\theta_{fj}}|X_f\mathrm{V}|\sum_k [\![k = j\dagger]\!]\log\theta_{fk} \\
&= -|X_f\Lambda|\frac{\partial}{\partial\theta_{fj}}\sum_k q_f(k)\log\theta_{fk} - |X_f\mathrm{V}|\frac{\partial}{\partial\theta_{fj}}\sum_k [\![k = j\dagger]\!]\log\theta_{fk} \\
&= -|X_f\Lambda|\frac{\partial}{\partial\theta_{fj}}q_f(j)\log\theta_{fj} - |X_f\mathrm{V}|\frac{\partial}{\partial\theta_{fj}}[\![j = j\dagger]\!]\log\theta_{fj} \\
&= -|X_f\Lambda|q_f(j)\frac{1}{\theta_{fj}} - |X_f\mathrm{V}|[\![j = j\dagger]\!]\frac{1}{\theta_{fj}}
\end{aligned} \tag{33}$$

As before, the derivative of the constraint $C_f = 0$ is 1, and we minimize (32) under the constraint by solving:

$$-|X_f\Lambda|q_f(j)\frac{1}{\theta_{fj}} - |X_f\mathrm{V}|[\![j = j\dagger]\!]\frac{1}{\theta_{fj}} = \lambda$$

$$\theta_{fj} = \left(-|X_f\Lambda|q_f(j) - |X_f\mathrm{V}|[\![j = j\dagger]\!]\right)/\lambda \tag{34}$$

Substituting into the constraint:

$$\sum_j \left(-|X_f\Lambda|q_f(j) - |X_f\mathrm{V}|[\![j = j\dagger]\!]\right)/\lambda = 1$$

$$-|X_f\Lambda| - |X_f\mathrm{V}| = \lambda$$

$$-|X_f| = \lambda$$

Substituting this back into (34):

$$\theta_{fj} = p(\Lambda|f)q_f(j) + p(\mathrm{V}|f)[\![j = j\dagger]\!] \tag{35}$$

That is, the maximizing solution is peaked precision, which is the update rule for YS-P. □

4.4 Algorithm YS-R

We now show that YS-R also decreases K in each iteration. In fact, it has essentially already been proven.

Theorem 8 *Algorithm YS-R decreases K in each iteration.*

Proof. In the proof of corollary 4, we showed that the choice

$$\theta_f = q_f$$

yields strictly positive gain. This is the update rule used by YS-R. □

4.5 Algorithm YS-FS

The original Yarowsky algorithm YS-0/DL-0 used smoothed precision with fixed ϵ as update rule. We have been unsuccessful at justifying this choice of update rule in general. However, we are able at least to show that it does decrease K when the selected feature is a new feature, not previously selected.

Theorem 9 *Algorithm YS-FS has positive gain in each iteration in which the selected feature has not been previously selected.*

Proof. By theorem 6, gain is positive if

$$H(q_f\|\theta_f^{\mathrm{old}}) > H(q_f\|\theta_f) \tag{36}$$

By the assumption that the selected feature f has not been previously selected, θ_f^{old} is the uniform distribution u, and the left-hand side of (36) is equal to $H(q_f\|u)$. It is easy to verify that $H(p\|u) = H(u)$ for any distribution p, hence

the left-hand side of (36) is equal to $H(u)$. Further, YS-FS uses smoothed precision as update rule, $\theta_f = \tilde{q}_f$, so (36) can be rewritten as:

$$H(u) > H(q_f \| \tilde{q}_f)$$

This condition does not hold trivially, inasmuch as cross entropy, like divergence, is unbounded. But we can show that it holds in this particular case.

We derive an expression for $H(q_f \| \tilde{q}_f)$.

$$\begin{aligned}
H(q_f \| \tilde{q}_f) &= -\sum_j q_f(j) \log \tilde{q}_f(j) \\
&= -\sum_j q_f(j) \log \left[\frac{1}{1+L\epsilon} q_f(j) + \frac{L\epsilon}{1+L\epsilon} u(j) \right] \\
&\leq -\sum_j q_f(j) \left[\frac{1}{1+L\epsilon} \log q_f(j) + \frac{L\epsilon}{1+L\epsilon} \log u(j) \right] \\
&= \frac{1}{1+L\epsilon} H(q_f) + \frac{L\epsilon}{1+L\epsilon} H(q_f \| u) \\
&= \frac{1}{1+L\epsilon} H(q_f) + \frac{L\epsilon}{1+L\epsilon} H(u) \qquad (37)
\end{aligned}$$

Observe that:

$$H(u) > \frac{1}{1+L\epsilon} H(q_f) + \frac{L\epsilon}{1+L\epsilon} H(u) \qquad (38)$$

iff

$$\left[1 - \frac{L\epsilon}{1+L\epsilon} \right] H(u) > \frac{1}{1+L\epsilon} H(q_f)$$

iff

$$H(u) > H(q_f)$$

We know that $H(u) \geq H(q_f)$ because the uniform distribution maximizes entropy. We know that the inequality is strict by the following reasoning. Since f is a new feature, $\theta_f^{\text{old}} = u$. Because of the restriction on step (2.1) in algorithm YS, $\theta_f^{\text{old}} \neq q_f$, hence $q_f \neq u$, and $H(u)$ is strictly greater than $H(q_f)$.

Hence (38) is true, and combining (38) with (37), we have shown (36) to be true, proving the theorem. □

4.6 Gain Revisited

We mentioned earlier that, in step (2.1) of algorithm YS, one might choose to select the feature that maximizes one of the lower bounds on G (theorems 5 and 6).

Of particular interest is theorem 6; it shows that:

$$G \geq |X_f \Lambda| \left[H(q_f \| \theta_f^{\text{old}}) - H(q_f \| \theta_f) \right]$$

If we use raw precision as update rule (that is, algorithm YS-R), and consider only new features, for which $\theta_f^{\text{old}} = u$, this becomes:

$$\begin{aligned} G &\geq |X_f \Lambda| \left[H(q_f \| u) - H(q_f \| q_f) \right] \\ &= |X_f \Lambda| \left[H(u) - H(q_f) \right] \\ &= |X_f \Lambda| \left[\log L - H(q_f) \right] \end{aligned} \tag{39}$$

Observe that, for any distribution p:

$$\begin{aligned} H(p) &= \sum_j p_j \log \frac{1}{p_j} \\ &= p_{j*} \log \frac{1}{p_{j*}} + \sum_{j \neq j*} p_j \log \frac{1}{p_j} \quad \text{where } j* = \arg\max_j p_j \\ &\leq p_{j*} \log \frac{1}{p_{j*}} \\ &\leq \log \frac{1}{p_{j*}} \end{aligned} \tag{40}$$

Combining (39) and (40):

$$G \geq |X_f \Lambda| \left[\log L + \log q_f(j\dagger) \right]$$

where $j\dagger$, as defined earlier, is the label that maximizes q_f. It follows that we maximize a lower bound on G by choosing the feature that maximizes $q_f(j\dagger)$. That is, we maximize a lower bound on G by choosing the rule $f \rightarrow j$ with maximal score $q_f(j)$. (By definition, $q_f(j\dagger) > q_f(j)$ for all $j \neq j\dagger$.) This is the strategy for rule selection that is used by Collins & Singer [3].

5 Minimization of Feature Entropy

At the beginning of the paper, the co-training algorithm was mentioned as an alternative to the Yarowsky algorithm. There is in fact a connection between co-training and the Yarowsky algorithm. In the original co-training paper [2] it was suggested that the algorithm be understood as seeking to maximize agreement on unlabeled data between classifiers trained on two different "views" of the data. Subsequent work [4] has proven a direct connection between classifier error and such cross-view agreement on unlabeled data.

In the current context, there is also justification for pursuing agreement on unlabeled data. However, the Yarowsky algorithm makes no assumption of the existence of two conditionally-independent views of the data. Rather, there is a motivation for seeking agreement on unlabeled data between arbitrary pairs of features.

Recall that our original objective function, H, can be expressed as the sum of an entropy term and a divergence term:

$$H = \sum_x \left[H(\phi_x) + D(\phi_x \| \pi_x) \right]$$

As $D(\phi_x \| \pi_x)$ becomes small and $H(\pi_x)$ becomes small, $H(\phi_x)$ necessarily also becomes small; hence we can limit H by limiting $H(\pi_x)$ and $D(\phi_x \| \pi_x)$. Intuitively, we wish to reduce the uncertainty of the model's predictions, while also improving the fit between the model's predictions and the known labels.

Let us focus now on the uncertainty of the model's predictions.

$$\begin{aligned}
H(\pi_x) &= -\sum_j \pi_x(j) \log \pi_x(j) \\
&= -\sum_j \pi_x(j) \log \left[\sum_{g \in F_x} \frac{1}{m} \theta_{gj} \right] \\
&\leq -\sum_j \pi_x(j) \sum_{g \in F_x} \frac{1}{m} \log \theta_{gj} \\
&= -\sum_j \left(\sum_{f \in F_x} \frac{1}{m} \theta_{fj} \right) \sum_{g \in F_x} \frac{1}{m} \log \theta_{gj} \\
&= -\frac{1}{m^2} \sum_{f \in F_x} \sum_{g \in F_x} \sum_j \theta_{fj} \log \theta_{gj} \\
&= \frac{1}{m^2} \sum_{f \in F_x} \sum_{g \in F_x} H(\theta_f \| \theta_g) \\
&= \frac{1}{m^2} \sum_{f \in F_x} \sum_{g \in F_x} [H(\theta_f) + D(\theta_f \| \theta_g)] \\
&= \frac{1}{m} \sum_{f \in F_x} H(\theta_f) + \frac{1}{m^2} \sum_{f \in F_x} \sum_{g \in F_x} D(\theta_f \| \theta_g) \qquad (41)
\end{aligned}$$

In other words, by decreasing the uncertainty of the prediction distributions of individual features, and simultaneously increasing the agreement among features (that is, decreasing their pairwise divergence), we decrease an upper bound on $H(\pi_x)$. This motivates inter-feature agreement without recourse to an assumption of independent views.

6 Conclusion

In this paper, we have presented a number of variants of the Yarowsky algorithm, and we have shown that they optimize natural objective functions. We considered first the modified generic Yarowsky algorithm Y-1 and showed that it minimizes the objective function H (which is equivalent to maximizing likelihood), provided that its base learner reduces H.

We then considered three families of specific Yarowsky-like algorithms. The Y-1/DL-EM algorithms (Y-1/DL-EM-Λ and Y-1/DL-EM-X) minimize H, but have the disadvantage that the DL-EM base learner has no similarity to Yarowsky's original base learner. A much better approximation to Yarowsky's original base

learner is provided by DL-1, and the Y-1/DL-1 algorithms (Y-1/DL-1-R and Y-1/DL-1-VS) were shown to minimize the objective function K, an upper bound for H. Finally, the YS algorithms (YS-P, YS-R, and YS-FS) are sequential variants, reminiscent of the Yarowsky-Cautious algorithm of Collins & Singer; we showed that they minimize K.

To the extent that these algorithms capture the essence of the original Yarowsky algorithm, they provide a formal understanding of Yarowsky's approach. Even if they are deemed to diverge too much from the original to cast light on its workings, they at least represent a new family of bootstrapping algorithms with solid mathematical foundations.

References

[1] Steven Abney. Bootstrapping. *Proceedings of ACL.* 2002.

[2] A. Blum and T. Mitchell. Combining labeled and unlabeled data with co-training. *COLT: Proceedings of the Workshop on Computational Learning Theory.* Morgan Kaufmann Publishers. 1998.

[3] Michael Collins and Yoram Singer. Unsupervised Models for Named Entity Classification. *Proceedings of EMNLP.* 1999.

[4] Sanjoy Dasgupta, Michael Littman, and David McAllester. "PAC Generalization Bounds for Co-training." *Proceedings of NIPS.* 2001.

[5] David Yarowsky. Unsupervised word sense disambiguation rivaling supervised methods. *Proceedings of the 33rd Annual Meeting of the Association for Computational Linguistics,* 189–196. 1995.