
Fast and Scalable Structural SVM with Slack Rescaling

Heejin Choi

Ofer Meshi

Nathan Srebro

Toyota Technological Institute at Chicago
6045 S. Kenwood Ave. Chicago, Illinois 60637 USA

Abstract

We present an efficient method for training slack-rescaled structural SVM. Although finding the most violating label in a margin-rescaled formulation is often easy since the target function decomposes with respect to the structure, this is not the case for a slack-rescaled formulation, and finding the most violated label might be very difficult. Our core contribution is an efficient method for finding the most-violating-label in a slack-rescaled formulation, given an oracle that returns the most-violating-label in a (slightly modified) margin-rescaled formulation. We show that our method enables accurate and scalable training for slack-rescaled SVMs, reducing runtime by an order of magnitude compared to previous approaches to slack-rescaled SVMs.

1 Introduction

Many problems in machine learning can be seen as structured output prediction tasks, where one would like to predict a set of labels with rich internal structure [1]. This general framework has proved useful for a wide range of applications from computer vision, natural language processing, computational biology, and others. In order to achieve high prediction accuracy, the parameters of structured predictors are learned from training data. One of the most effective and commonly used approaches for this supervised learning task is *Structural SVM*, a method that generalizes binary SVM to structured outputs [15]. Since the structured error is non-convex, Tsochantaridis et al. [15] propose to replace it with a convex surrogate loss function. They formulate two such surrogates, known as *margin* and *slack rescaling*.

While slack rescaling often produces more accurate pre-

dictors, margin rescaling has been far more popular due to its better computational requirements. In particular, both formulations require optimizing over the output space, but while margin rescaling preserves the structure of the score and error functions, the slack-rescaling does not. This results in harder inference problems during training. To address this challenge, Sarawagi and Gupta [11] propose a method to reduce the problem of slack rescaling to a series of modified margin rescaling problems. They show that their method outperforms margin rescaling in several domains. However, there are two main caveats in their approach. First, the optimization is only heuristic, that is, it is not guaranteed to solve the slack rescaling objective exactly. Second, their method is specific to the cutting plane training algorithm and does not easily extend to stochastic algorithms. More recently, Bauer et al. [2] proposed an elegant dynamic programming approach to the slack rescaling optimization problem. However, their formulation is restricted to sequence labeling and hamming error, and does not apply to more general structures.

In this paper we propose an efficient method for solving the optimization problem arising from the slack rescaling formulation. Similar to Sarawagi and Gupta [11] our method reduces finding the most violated label in slack rescaling to a series of margin rescaling problems. However, in contrast to their approach, our approach can be easily used with training algorithms like stochastic gradient descent (SGD) [10] and block Frank-Wolfe (FW) [7], which often scale much better than cutting plane. We first propose a very simple approach that minimizes an upper bound on the slack rescaling objective function, and only requires access to a margin rescaling oracle. This formulation is quite general and can be used with any error function and model structure, and many training algorithms such as cutting plane, SGD and FW. However, this method is not guaranteed to always find the most violating label. Indeed, we show that always finding the most violating label for a slack-rescaled formulation is impossible using only a margin rescaling oracle. To address this, we suggest using a modified oracle, that is typically as easy to implement as margin rescaling, and present a more sophisticated algorithm which solves the slack rescaling formulation exactly, and also enjoys good approximation guarantees after a small number of it-

erations. We demonstrate empirically that our algorithm outperforms existing baselines on several real-world applications, including hierarchical and multi-label classification.

2 Problem Formulation

In this section we review the basics of structured output prediction and describe the relevant training objectives. In structured output prediction the task is to map data instances x to a set of output labels $y \in \mathcal{Y}$. Structured SVMs use a linear discriminant mapping of the form $y(x; w) = \operatorname{argmax}_{y \in \mathcal{Y}} w^\top \phi(x, y)$, where $\phi(x, y) \in \mathbb{R}^d$ is a feature function relating input-output pairs, and $w \in \mathbb{R}^d$ is a corresponding vector of weights. Our interest is in the supervised learning setting, where w is learned from training data $\{x_i, y_i\}_{i=1}^n$ by minimizing the empirical risk. The prediction quality is measured by an error function $L(y, y_i) \geq 0$ which determines how bad it is to predict y when the ground-truth is in fact y_i .

Since optimizing $L(y, y_i)$ directly is hard due to its complicated dependence on w , several alternative formulations minimize a convex upper bound instead. Structural SVM is an elegant max-margin approach which uses a structured hinge loss surrogate [15, 14]. Two popular surrogates are margin and slack rescaling. In particular, denoting the model score by $f(y) = w^\top \phi(x, y)$ (we omit the dependence on x and w to simplify notation), the margin rescaling training objective is given by:

$$\begin{aligned} \min_{w, \xi} \quad & \frac{C}{2} \|w\|_2^2 + \frac{1}{n} \sum_i \xi_i \\ \text{s.t.} \quad & f(y_i) - f(y) \geq L(y, y_i) - \xi_i \quad \forall i, y \neq y_i \\ & \xi_i \geq 0 \quad \forall i \end{aligned} \quad (1)$$

where C is the regularization constant. Similarly, the slack rescaling formulation scales the slack variables by the error term:

$$\begin{aligned} \min_{w, \xi} \quad & \frac{C}{2} \|w\|_2^2 + \frac{1}{n} \sum_i \xi_i \\ \text{s.t.} \quad & f(y_i) - f(y) \geq 1 - \frac{\xi_i}{L(y, y_i)} \quad \forall i, y \neq y_i \\ & \xi_i \geq 0 \quad \forall i \end{aligned} \quad (2)$$

Intuitively, both formulations seek to find a w which assigns high scores to the ground-truth compared to the other possible labellings. When y is very different than the true y_i (L is large) then the difference in scores should also be larger. There is, however, an important difference between the two forms. In margin rescaling, high loss can occur for labellings with high error even though they are already classified correctly with a margin. This may divert training

from the interesting labellings where the classifier errs, especially when L can take large values, as in the common case of hamming error. In contrast, in slack rescaling labellings that are classified with a margin incur no loss. Another difference between the two formulations is that the slack rescaling loss is invariant to scaling of the error term, while in margin rescaling such scaling changes the meaning of the features ϕ .

In many cases it is easier to optimize an unconstrained problem. In our case it is straightforward to write (1) and (2) in an unconstrained form:

$$\min_w \frac{C}{2} \|w\|_2^2 + \frac{1}{n} \sum_i \max_{y \in \mathcal{Y}} (L(y, y_i) + f(y) - f(y_i)) \quad (3)$$

$$\min_w \frac{C}{2} \|w\|_2^2 + \frac{1}{n} \sum_i \max_{y \in \mathcal{Y}} L(y, y_i) (1 + f(y) - f(y_i)) \quad (4)$$

Most of the existing training algorithms for structural SVM require solving the maximization-over-labellings problems in (4) and (3):

$$\text{Margin rescaling} : \operatorname{argmax}_{y \in \mathcal{Y}} L(y, y_i) + f(y) - f(y_i) \quad (5)$$

$$\text{Slack rescaling} : \operatorname{argmax}_{y \in \mathcal{Y}} L(y, y_i) (1 + f(y) - f(y_i)) \quad (6)$$

To better understand the difference between margin and slack rescaling we focus on a single training instance i and define the functions: $h(y) = 1 + f(y) - f(y_i)$ and $g(y) = L(y, y_i)$. With these definitions we see that the maximization (5) for margin rescaling is $\max_{y \in \mathcal{Y}} h(y) + g(y)$, while the maximization (6) for slack rescaling is $\max_{y \in \mathcal{Y}} h(y)g(y)$. It is now obvious why margin rescaling is often easier. When the score and error functions h and g decompose into a sum of simpler functions, we can exploit that structure in order to solve the maximization efficiently [15, 14, 5]. In contrast, the slack rescaling score does not decompose even when both h and g do. What we show, then, is how to solve problems of the form $\max_y h(y)g(y)$, and thus the maximization (6), having access only to an oracle for additive problems of the form $\max_y h(y) + \lambda g(y)$.

That is, we assume that we have access to a procedure, referred to as the λ -oracle, which can efficiently solve the problem:

$$y_\lambda = \mathcal{O}(\lambda) = \operatorname{argmax}_{y \in \mathcal{Y}} \mathcal{L}_\lambda(y) \quad (7)$$

where $\mathcal{L}_\lambda(y) = h(y) + \lambda g(y)$. This problem is just a rescaling of (5). E.g., for linear responses it is obtained by scaling the weight vector by $1/\lambda$. If we can handle margin rescaling efficiently we can most likely implement the λ -oracle efficiently. This is also the oracle used by Sarawagi and Gupta [11]. In Section 4, we show how to obtain a solution to the slack-rescaling problem (6) using such a λ -oracle. Our method can be used as a subroutine in a variety of training algorithms, and we demonstrate that it is more

scalable than previous methods. However, we also show that this approach is limited, since no procedure that only has access to a λ -oracle can guarantee the quality of its solution, no matter how much time it is allowed to run.

Therefore, we propose an alternative procedure that can access a more powerful oracle, which we call the *constrained λ -oracle*:

$$y_{\lambda, \alpha, \beta} = \mathcal{O}_c(\lambda, \alpha, \beta) = \max_{y \in \mathcal{Y}, \alpha h(y) > g(y), \beta h(y) \leq g(y)} \mathcal{L}_\lambda(y), \quad (8)$$

where $\alpha, \beta \in \mathbb{R}$. This oracle is similar to the λ -oracle, but can additionally handle linear constraints on the values $h(y)$ and $g(y)$. In the sequel we show that in many interesting cases this oracle is not more computationally expensive than the basic one. For example, when the λ -oracle is implemented as a linear program (LP), the additional constraints are simply added to the LP formulation and do not complicate the problem significantly. Before presenting our algorithms for optimizing (6), we first review the training framework in the next section.

3 Optimization for Slack Rescaling

In this section we briefly survey cutting plane and stochastic gradient descent optimization for the slack rescaled objective (2) and (4). This will be helpful in understanding the difference between our approach and that of prior work on the slack rescaled objective by Sarawagi and Gupta [11].

The cutting plane algorithm was proposed for solving the structural SVM formulation in [15, 6]. This algorithm has also been used in previous work on slack rescaling optimization [11, 2]. The difficulty in optimizing (2) stems from the number of constraints, which is equal to the size of the output space \mathcal{Y} (for each training instance). The cutting plane method maintains a small set of constraints and solves the optimization problem only over that set. At each iteration the active set of constraints is augmented with new violated constraints, and it can be shown that not too many such constraints need to be added for a good solution to be found [6]. The main computational bottleneck here is to find a violating constraint at each iteration, which is challenging since it requires searching over the output space for some violating labeling y .

Relying on this framework, Sarawagi and Gupta [11] use the formulation in (2) and rewrite the constraints as:

$$1 + f(y) - f(y_i) - \frac{\xi_i}{L(y, y_i)} \leq 0 \quad \forall i, y \neq y_i \quad (9)$$

In our notation, to find a violated constraint they aim at solving the problem:

$$\operatorname{argmax}_{y \in \mathcal{Y}'} \left(h(y) - \frac{\xi_i}{g(y)} \right) \quad (10)$$

where $\mathcal{Y}' = \{y | y \in \mathcal{Y}, h(y) > 0, y \neq y_i\}$. They suggest to minimize a convex upper bound of (10) which stems from the convex conjugate function of $\frac{\xi_i}{g(y)}$:

$$\begin{aligned} \max_{y \in \mathcal{Y}'} h(y) - \frac{\xi_i}{g(y)} &= \max_{y \in \mathcal{Y}'} \min_{\lambda \geq 0} \left(h(y) + \lambda g(y) - 2\sqrt{\xi_i \lambda} \right) \\ &\leq \min_{\lambda \geq 0} \max_{y \in \mathcal{Y}'} F'(\lambda, y) = \min_{\lambda \geq 0} \max_{y \in \mathcal{Y}'} F'(\lambda, y) = \min_{\lambda \geq 0} F(\lambda) \end{aligned} \quad (11)$$

where $F(\lambda) = \max_{y \in \mathcal{Y}'} F'(\lambda, y) = \max_{y \in \mathcal{Y}'} h(y) + \lambda g(y) - 2\sqrt{\xi_i \lambda}$. Since $F(\lambda)$ is a convex function, (11) can be solved by a simple search method such as golden search over λ [11].

Although this approach is suitable for the cutting plane algorithm, unfortunately it cannot be easily extended to other training algorithms. In particular, $F'(\lambda, y)$ is defined in terms of slack variables ξ_i , which ties it to the constrained form (2). On the other hand, algorithms such as stochastic gradient descent (SGD) [10, 12], stochastic dual coordinate ascent (SDCA) [13], or block-coordinate Frank-Wolfe (FW) [7], all optimize the unconstrained objective form (4). These methods are typically preferable in the large scale setting, since they have very low per-iteration cost, handling a single example at a time, with the same overall iteration complexity as cutting plane methods. In contrast, the cutting plane algorithm considers the entire training set at each iteration, so the method does not scale well to large problems. Since our goal in this work is to handle large datasets, we would like to be able to use the stochastic methods mentioned above, working on the unconstrained formulation (4). The update in these algorithms requires solving the maximization problem (6), which is the goal of Section 4. Note that solving (6) also allows using a cutting plane method if desired.

4 Algorithms

In this section we present our main contribution, a framework for solving the maximization problem (6), which we write as:

$$\max_y \Phi(y) := \max_y h(y)g(y) \quad (12)$$

We describe two new algorithms to solve this problem using access to the λ -oracle, which have several advantages over previous approaches. However, we also show that any algorithm which uses only the λ -oracle cannot always recover an optimal solution. Therefore, in Section 4.5 we proposed an improved algorithm which requires access to an augmented λ -oracle that can also handle linear constraints.

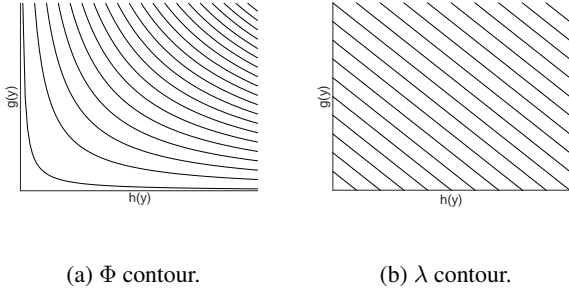


Figure 1: Contour of the two functions considered in \mathbb{R}^2 . Φ contour is the contour of the objective function, and λ contour is the contour used by the oracle.

4.1 Binary search

We first present a binary search algorithm similar to the one proposed by Sarawagi and Gupta [11], but with one main difference. Our algorithm can be easily used with training methods that optimize the unconstrained objective (4), and can therefore be used for SGD, SDCA and FW. The algorithm minimizes a convex upper bound on Φ without slack variable ξ_i . The algorithm is based on the following lemma (details and proofs are in Appendix A).

Lemma 1. Let $\bar{F}(\lambda) = \frac{1}{4} \max_{y \in \mathcal{Y}^+} (\frac{1}{\lambda} h(y) + \lambda g(y))^2$, then

$$\max_{y \in \mathcal{Y}} \Phi(y) \leq \min_{\lambda > 0} \bar{F}(\lambda)$$

and $\bar{F}(\lambda)$ is a convex function in λ .

Rather than minimizing this upper bound, we next present an algorithm that aims to optimize $\Phi(y)$ in a more direct manner, using a geometrical interpretation of mapping labels into \mathbb{R}^2 .

4.2 Geometrical Interpretation of λ -oracle search

To understand the problem better and motivate our approach, it is useful to consider the following geometrical interpretation of (12): we map each labeling y to a vector $\vec{y} = [h(y) \ g(y)] \in \mathbb{R}^2$. Let $\vec{\mathcal{Y}} = \{\vec{y} \in \mathbb{R}^2 | y \in \mathcal{Y}\}$ be the set of all mapped labels. The maximization (12) reduces to the problem: given a set of points $\vec{\mathcal{Y}} \subset \mathbb{R}^2$, maximize the product of their coordinates $y^* = \operatorname{argmax}_{\vec{y} \in \vec{\mathcal{Y}}} [\vec{y}]_1 \cdot [\vec{y}]_2$.

The contours of our objective function $\bar{\Phi}(\vec{y}) = [\vec{y}]_1 \cdot [\vec{y}]_2$ are then hyperbolas. We would like to maximize this function by repeatedly finding points that maximize linear objectives of the form $\vec{\mathcal{L}}_\lambda(\vec{y}) = [\vec{y}]_1 + \lambda[\vec{y}]_2$, whose contours form lines in the plane (see Figure 1). An example of mapping labels into \mathbb{R}^2 is shown in Appendix B.

The importance of the \mathbb{R}^2 mapping is that each y_λ revealed by the λ -oracle shows that the optimal y^* can only reside in a small slice of the plane. See figure 2.




Figure 2: Geometric interpretation of the λ -oracle: y^* must reside in the marked area between the upper bound S_λ and the lower bound C_λ . It follows that $h(y^*)$ and $g(y^*)$ reside in a simple segment \dot{H} and \dot{G} respectively.

Lemma 2. Let S_λ be a line through \vec{y}_λ and $\vec{z} = [\lambda[\vec{y}_\lambda]_2, \frac{1}{\lambda}[\vec{y}_\lambda]_1]$, and let $C_\lambda = \{\vec{y} \in \mathbb{R}^2 | [\vec{y}]_1 \cdot [\vec{y}]_2 = \bar{\Phi}(\vec{y}_\lambda)\}$ be the hyperbola through \vec{y}_λ . Then, y^* is on or below line S_λ , and y^* is on or above hyperbola C_λ .

Proof. If there exists a $\vec{y} \in \vec{\mathcal{Y}}$ which is above S_λ , it contradicts the fact that \vec{y}_λ is the argmax point for function $\vec{\mathcal{L}}_\lambda$. And the second argument follows from y^* being the argmax label w.r.t. $\bar{\Phi}$, and the area above C_λ corresponds to points whose $\bar{\Phi}$ value is greater than \vec{y}_λ . \square

It follows that $h(y^*)$ and $g(y^*)$ must each reside in a segment:

Lemma 3. Let $\dot{H} = [\min([\vec{y}_\lambda]_1, [\vec{z}]_1), \max([\vec{y}_\lambda]_1, [\vec{z}]_1)]$ and $\dot{G} = [\min([\vec{y}_\lambda]_2, [\vec{z}]_2), \max([\vec{y}_\lambda]_2, [\vec{z}]_2)]$. Then,

$$h(y^*) \in \dot{H}, \quad g(y^*) \in \dot{G}$$

Proof. This follows from the fact that S_λ and C_λ intersects at two points, \vec{y}_λ and \vec{z} , and the boundaries, S_λ and C_λ , are strictly decreasing functions in \mathbb{R}^2 . \square

4.3 Bisecting search

In this section, we propose a search algorithm which is based on the previous geometric interpretation. Similar to the binary search, our method also relies on the basic λ -oracle.

We first give an overview of the algorithm. We maintain a set of possible value ranges $\lambda^* = \operatorname{argmax}_{\lambda > 0} \Phi(y_\lambda)$, $h(\lambda^*)$, and $g(\lambda^*)$ as L , H , and G , respectively; all initialized as \mathbb{R} . First, for each y_λ returned by the oracle, we take an intersection of G and H with a segment of possible values of $h(y)$ and $g(y)$, respectively, using Lemmas 2 and 3.

Second, we reduce the space L of potential λ 's based on the following Lemma (proved in Appendix C).

Lemma 4. $h(y_\lambda)$ is a non-increasing function of λ , and $g(y_\lambda)$ is a non-decreasing function of λ .

Thus, we can discard $\{\lambda' | \lambda' > \lambda\}$ if $h(y_{\lambda'}^*) > h(y_\lambda)$ or $\{\lambda' | \lambda' < \lambda\}$ otherwise from L . Next, we pick $\lambda \in L$ in the middle, and query y_λ . The algorithm continues until at least one of L, H , and G is empty. This *bisecting search* procedure is given in Algorithm 1.

Similar to the binary search from the previous section, this algorithm can be used with training methods like SGD and SDCA, as well as the cutting-plane algorithm. However, this approach has several advantages compared to the binary search. First, the binary search needs explicit upper and lower bounds on λ , thus it has to search the entire λ space [11]. However, the bisecting search can directly start from any λ without an initial range, and for instance, this can be used to warm-start from the optimal λ in the previous iteration. Furthermore, we point out that since the search space of h and g is also bisected, the procedure can terminate early if either of them becomes empty.

Finally, in Appendix D we propose two improvements that can be applied to either the binary search or the bisecting search. Specifically, we first provide a simple stopping criterion that can be used to terminate the search when the current solution y_{λ_t} will not further improve. Second, we show how to obtain a bound on the suboptimality of the current solution, which can give some guarantee on its quality.

So far we have used the λ -oracle as a basic subroutine in our search algorithms. Unfortunately, as we show next, this approach is limited as we cannot guarantee finding the optimal solution y^* , even with unlimited number of calls to the λ -oracle. This is somewhat distressing since with unlimited computation we can find the optimum of (6) by enumerating all y 's.

4.4 Limitation of the λ -oracle

Until now, we used only the λ -oracle to search for Φ^* without directly accessing the functions h and g . We now show that this approach, searching Φ^* with only a λ -oracle, is very limited: even with an unlimited number of queries, the search cannot be exact and might return a trivial solution in the worst case (see Appendix E for proof).

Theorem 1. Let $\hat{H} = \max_y h(y)$ and $\hat{G} = \max_y g(y)$. For any $\epsilon > 0$, there exists a problem with 3 labels such that for any $\lambda \geq 0$, $\Phi(y_\lambda) = \min_{y \in \mathcal{Y}} \Phi(y) < \epsilon$, while $\Phi(y^*) = \frac{1}{4} \hat{H} \hat{G}$.

Theorem 1 shows that any search algorithm that can access the function only through λ -oracle, including the method of Sarawagi and Gupta [11] and both methods presented

Algorithm 1 Bisecting search

```

1: procedure BISECTING( $\lambda_0$ )
Input: Initial  $\lambda$  for the search  $\lambda_0 \in \mathbb{R}_+$ 
Output:  $\hat{y} \in \mathcal{Y}$ .
Initialize:  $H = G = L = \mathbb{R}_+, \lambda = \lambda_0, \hat{\Phi} = 0$ .
2:   while  $H \neq \emptyset$  and  $G \neq \emptyset$  do
3:      $y' \leftarrow \mathcal{O}(\lambda)$ 
4:      $u \leftarrow [h(y') \lambda g(y')], v \leftarrow [g(y') \frac{1}{\lambda} h(y')]$ 
5:      $H \leftarrow H \cap \{h' | \min u \leq h' \leq \max u\}$  ▷
       Update
6:      $G \leftarrow G \cap \{g' | \min v \leq g' \leq \max v\}$ 
7:     if  $v_1 \leq v_2$  then ▷ Increase  $\lambda$ 
8:        $L \leftarrow L \cap \{\lambda' \in \mathbb{R} | \lambda' \geq \lambda\}$ 
9:     else ▷ Decrease  $\lambda$ 
10:       $L \leftarrow L \cap \{\lambda' \in \mathbb{R} | \lambda' \leq \lambda\}$ 
11:       $\lambda \leftarrow \frac{1}{2}(\min L + \max L)$ 
12:      if  $h(y')g(y') \geq \hat{\Phi}$  then
13:         $\hat{y} \leftarrow y', \hat{\Phi} \leftarrow h(y')g(y')$ 

```

above, cannot be guaranteed to find a label optimizing $\Phi(y)$, even approximately, and even with unlimited accesses to the oracle. This problem calls for a more powerful oracle.

4.5 Angular search with the constrained- λ -oracle

The *constrained λ -oracle* defined in (8) has two inequality constraints to restrict the search space. Using this modified oracle, we can present an algorithm that is guaranteed to find the most violating constraint, as captured by the following theorem, proved in Appendix F:

Theorem 2. Angular search described in Algorithm 2 finds the optimum $y^* = \operatorname{argmax}_{y \in \mathcal{Y}} \Phi(y)$ using at most $t = 2M + 1$ iteration where $M = |\mathcal{Y}|$ is the number of the labels.

This is already an improvement over the previous methods, as at least we are guaranteed to return the actual most violating label. However, it is still disappointing since the number of iterations, and thus number of oracle accesses might actually be larger than the number of labels. This defies the whole point, since we might as well just enumerate over all M possible labels. Unfortunately, even with a constrained oracle, this is almost the best we can hope for. In fact, even if we allow additional linear constraints, we might still need M oracle accesses, as indicated by the following theorem, proved in Appendix E.

Theorem 3. Any search algorithm accessing labels only through a λ -oracle with any number of linear constraints cannot find y^* using less than M iterations in the worst case, where $M = |\mathcal{Y}|$ is the number of labels.

Fortunately, even though we cannot guarantee optimizing $\Phi(y)$ exactly using a small number of oracle accesses, we

	Angular	Bisecting	Sarawagi
Yeast (N=160)			
Success	22.4%	16.5%	16.4%
Queries per search	3.8	10.3	43.2
Average time (ms)	4.7	3.6	18.5
RCV1 (N=160)			
Success	25.6%	18.2%	18%
Queries per search	4.8	12.7	49
Average time (ms)	4.4	5.2	20.9

Table 1: Comparison of the search algorithm.

convergence. Thus, it is suggested to use a priority queue, and dequeue the angle with the highest upper bound on Φ . A similar strategy is to have a *label cache*, the subset of previous most violated labels, denoted as \mathcal{C} . With the label cache, we can discard a large part of the search space $\{\vec{y} | \Phi(\vec{y}) \leq \max_{y' \in \mathcal{C}} \Phi(\vec{y}')\}$ immediately. Algorithm 2 also uses the constrained λ -oracle to avoid returning previously found labels. Finally, for λ_0 , we suggest to use $\lambda_0 = \frac{\hat{H}}{\hat{G}}$, with \hat{H} calculated from the current weights w .

See Appendix G for an illustration of the angular search.

5 Experiments

In this section we compare the behavior of various search algorithms on standard benchmark datasets, and its effect on the learning algorithm. Specifically, we show that Angular search with SGD is not only much faster than the other alternatives, but also results in an accurate predictor which outperforms both margin rescaling and other slack rescaling methods.

Unlike the simple structure used in [2], we apply our approach to complicated structures. Specifically, we experiment with multi-label dataset modeled by a Markov Random Field with pairwise potentials as in [5]. Since the inference of margin rescaling is NP-hard in this case, we rely on linear programming relaxation to compute the λ -oracle. Note that this complicates the problem, since the number of labels becomes even larger with additional fractional solutions. Also notice that all of our results above apply with mild modifications to this harder setting. Two standard benchmark multi-label datasets, Yeast [4] (14 labels) and RCV1 [8], are tested. For RCV1 we reduce the data to the 50 most frequent labels. For angular search, we stop the search whenever $\Phi(\hat{y}) > 0.999 \cdot \Phi(y^*)$ holds, to avoid numerical issues.

5.1 Comparison of the search algorithms

We first focus on the search method and run cutting-plane optimization, where at each iteration we call all three algorithms to find the most violating label: Angu-

Yeast				
	Acc	Label loss	MiF1	MaF1
Slack	.54	.205	.661	.651
Margin	.545	.204	.666	.654
RCV1				
	Acc	Label loss	MiF1	MaF1
Slack	.676	.023	.755	.747
Margin	.662	.023	.753	.742

Table 2: Results on Multi-label Dataset with Markov Random Field.

lar search, Bisecting search, and Sarawagi and Gupta’s search (Sarawagi) [11]. The cutting-plane algorithm calls the search procedures to find the most violating label \hat{y} , and adds it to the active set if the violation is larger than some margin ϵ , i.e., $\Delta(\hat{y}, y_i)(1 + f(\hat{y}) - f(y_i)) > \xi_i + \epsilon$. We use the result of Angular search for the actual update and repeat. Table 1 compares the performance of the search in terms of the time spend, the number of oracle queries, and the success rate of finding the most violating label. Success is measured by the percentage in which the search algorithm finds such a violating label. As expected by Theorem 1, Bisecting and Sarawagi’s search miss the violating label in cases where Angular search successfully finds one. For RCV1 dataset, not only is Angular search more accurate, but it also uses about 2.6 times less oracle queries than Bisecting and 10.1 times less queries than Sarawagi’s search. As for runtime, Angular search is 1.18 times faster than Bisecting search, and 4.7 times faster than Sarawagi’s algorithm.

In order to understand the effect of the search procedure on the overall performance, we next compare combinations of learning and search algorithms. In figure 4 we compare convergence rate and accuracy for the different combinations. These show that Angular search with SGD converges much faster than other schemes. Additional plots showing convergence w.r.t. the number of queries and iterations can be found in Appendix I.

Table 2 shows a comparison of predictive performance for the multi-label datasets. In this case margin and slack rescaling perform similarly, with a slight advantage for slack rescaling on RCV1.

5.2 Hierarchical Multi-label Classification

We further run experiments on a hierarchical multi-label classification problem [3]. In hierarchical multi-label classification, each label y is a leaf node in a given graph, and it shares ancestor nodes. It can be described as a graphical model where a potential of a multi-label $\mathcal{Y} = \{y_1, \dots, y_k\}$ is the sum of all potentials of its ancestors, i.e., $\Phi(Y) = \sum_{n' \in \bigcup_{n \in \mathcal{Y}} \text{Anc}(n)} \Phi(n)$. We extracted 1500 instances with dimensionality 17944 and a graph structure of 156 nodes

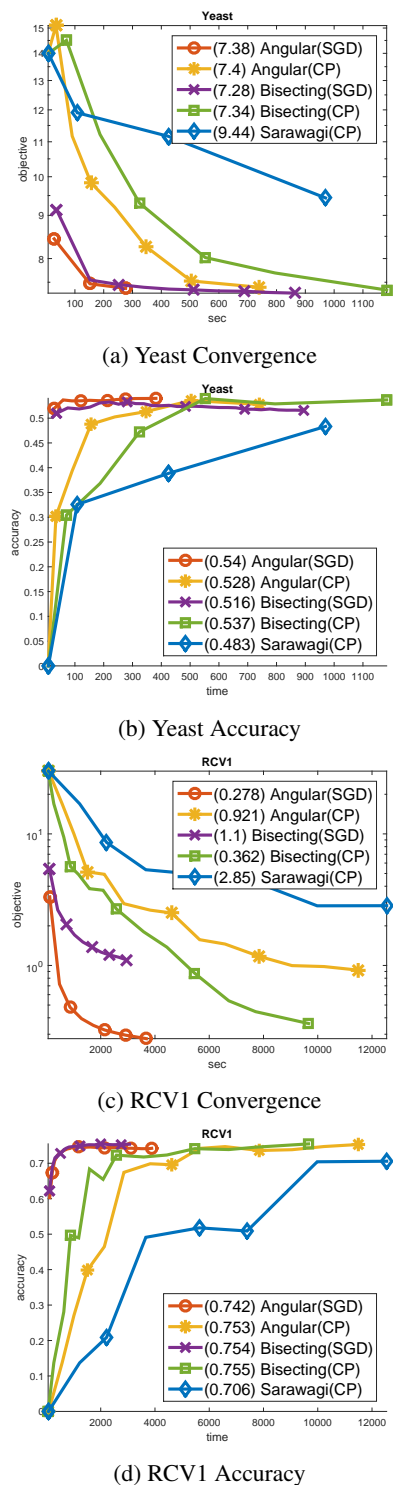


Figure 4: Convergence and accuracy of various learning and search algorithms.

with 123 labels from SWIKI-2011. SWIKI-2011 is a multi-label dataset of Wikipedia pages with a DAG structure from LSHTC competition.¹ We used 750 instances as training set, 250 instances as holdout set, and 500 instances as test set. The hamming distance of label vectors including inner nodes is used as label loss. We show that slack rescaling in such large label structure is tractable and outperforms margin rescaling. Unfortunately, the cutting-plane approach failed to scale up to the large problem size, however, using our approach we show that slack rescaling in such large label structure is tractable and outperforms margin rescaling. See Table 3.

	Acc	Label loss	MiF1	MaF1
Slack	.3798	.0105	.3917	.3880
Margin	.3327	.0110	.3394	.3378

Table 3: Result on hierarchical multi-label dataset

6 Summary

As we saw in our experiments, and has also been previously noted, slack rescaling is often beneficial compared to margin rescaling in terms of predictive performance. However, the margin-rescaled argmax (5) is often much easier computationally due to its additive form. Here we show how an oracle for solving an argmax of the form (5), or perhaps a slightly modified form (the constrained- λ oracle), is sufficient for also obtaining exact solutions to the slack-rescaling argmax (6). This allows us to train slack-rescaled SVMs using SGD, obtaining better predictive performance than using margin rescaling. Prior work in this direction provided only approximations [11] or handled only specific models and losses [2], and more significantly, was specific to cutting-plane optimization. In this work we provide a generic method relying on a simple explicitly specified oracle that is guaranteed to be exact and efficient even when the number of labels is infinite. Our approach allows using SGD and is thus more suitable for large scale problems.

References

- [1] Bakir, G. H., Hofmann, T., Schölkopf, B., Smola, A. J., Taskar, B., and Vishwanathan, S. V. N. (2007). *Predicting Structured Data*. The MIT Press.
- [2] Bauer, A., Gornitz, N., Biegler, F., Muller, K.-R., and Kloft, M. (2014). Efficient algorithms for exact inference in sequence labeling svms. *Neural Networks and Learning Systems, IEEE Transactions on*, 25(5):870–881.
- [3] Cai, L. and Hofmann, T. (2004). Hierarchical document categorization with support vector machines. In

¹<http://lshtc.iit.demokritos.gr/>

- Proceedings of the thirteenth ACM international conference on Information and knowledge management*, pages 78–87. ACM.
- [4] Elisseeff, A. and Weston, J. (2001). A kernel method for multi-labelled classification. In *Advances in neural information processing systems*, pages 681–687.
- [5] Finley, T. and Joachims, T. (2008). Training structural svms when exact inference is intractable. In *Proceedings of the 25th international conference on Machine learning*, pages 304–311. ACM.
- [6] Joachims, T., Finley, T., and Yu, C.-N. J. (2009). Cutting-plane training of structural svms. *Machine Learning*, 77(1):27–59.
- [7] Lacoste-Julien, S., Jaggi, M., Schmidt, M., and Pletscher, P. (2013). Block-coordinate frank-wolfe optimization for structural svms. In *ICML 2013 International Conference on Machine Learning*, pages 53–61.
- [8] Lewis, D. D., Yang, Y., Rose, T. G., and Li, F. (2004). Rcv1: A new benchmark collection for text categorization research. *The Journal of Machine Learning Research*, 5:361–397.
- [9] Nghia, N. D., Chinh, D. D., and Duong, P. C. (1995). Minimizing the product of two discrete convex functions. *ACTA MATHEMATICA VIETNAMICA*, 20(2):265–267.
- [10] Ratliff, N., Bagnell, J. A. D., and Zinkevich, M. (2007). (Online) subgradient methods for structured prediction. In *AISTATS*.
- [11] Sarawagi, S. and Gupta, R. (2008). Accurate max-margin training for structured output spaces. In *Proceedings of the 25th international conference on Machine learning*, pages 888–895. ACM.
- [12] Shalev-Shwartz, S., Singer, Y., Srebro, N., and Cotter, A. (2011). Pegasos: Primal estimated sub-gradient solver for svm. *Mathematical programming*, 127(1):3–30.
- [13] Shalev-Shwartz, S. and Zhang, T. (2013). Accelerated proximal stochastic dual coordinate ascent for regularized loss minimization. *Mathematical Programming*, pages 1–41.
- [14] Taskar, B., Guestrin, C., and Koller, D. (2003). Max-margin Markov networks. In *Advances in Neural Information Processing Systems*. MIT Press.
- [15] Tsochantaridis, I., Hofmann, T., Joachims, T., and Al-tun, Y. (2004). Support vector machine learning for interdependent and structured output spaces. In *Proceedings of the twenty-first international conference on Machine learning*, page 104. ACM.

Supplementary Material

Fast and Scalable Structural SVM with Slack Rescaling

Appendix A Details of binary search

Lemma 1. Let $\bar{F}(\lambda) = \frac{1}{4} \max_{y \in \mathcal{Y}^+} \left(\frac{1}{\lambda} h(y) + \lambda g(y) \right)^2$, then

$$\max_{y \in \mathcal{Y}} \Phi(y) \leq \min_{\lambda > 0} \bar{F}(\lambda)$$

and $\bar{F}(\lambda)$ is a convex function in λ .

Proof. First, let $\mathcal{Y}^+ = \{y | y \in \mathcal{Y}, h(y) > 0\}$, then $\max_{y \in \mathcal{Y}} \Phi(y) = \max_{y \in \mathcal{Y}^+} \Phi(y)$, since any solution y such that $h(y) < 0$ is dominated by y_i , which has zero loss. Second, we prove the bound w.r.t. $y \in \mathcal{Y}^+$. In the following proof we use a quadratic bound (for a similar bound see [9]).

$$\begin{aligned} \max_{y \in \mathcal{Y}^+} \Phi(y) &= \max_{y \in \mathcal{Y}^+} h(y)g(y) = \max_{y \in \mathcal{Y}^+} \frac{1}{4} \left(2\sqrt{h(y)g(y)} \right)^2 \\ &= \frac{1}{4} \left(\max_{y \in \mathcal{Y}^+} \min_{\lambda > 0} \left\{ \frac{1}{\lambda} h(y) + \lambda g(y) \right\} \right)^2 \\ &\leq \frac{1}{4} \left(\min_{\lambda > 0} \max_{y \in \mathcal{Y}^+} \left\{ \frac{1}{\lambda} h(y) + \lambda g(y) \right\} \right)^2 \end{aligned} \tag{13}$$

To see the convexity of $\bar{F}(\lambda)$, we differentiate twice to obtain:

$$\frac{\partial^2 \bar{F}(\lambda)}{\partial \lambda^2} = \frac{1}{4} \max_{y \in \mathcal{Y}^+} 6 \frac{1}{\lambda^4} h(y)^2 + 2g(y)^2 > 0$$

□

Similar to [11], we obtain a convex upper bound on our objective. Evaluation of the upper bound $\bar{F}(\lambda)$ requires using only the λ -oracle. Importantly, this alternative bound $\bar{F}(\lambda)$ does not depend on the slack variable ξ_i , so it can be used with algorithms that optimize the unconstrained formulation (4), such as SGD, SDCA and FW. As in [11], we minimize $\bar{F}(\lambda)$ using *binary search* over λ . The algorithm keeps track of y_{λ_t} , the label returned by the λ -oracle for intermediate values λ_t encountered during the binary search, and returns the maximum label $\max_t \Phi(y_{\lambda_t})$. This algorithm focuses on the upper bound $\min_{\lambda > 0} \bar{F}(\lambda)$, and interacts with the target function Φ only through evaluations $\Phi(y_{\lambda_t})$ (similar to [11]).

Appendix B An example of label mapping




Figure 5: A snapshot of labels during optimization with Yeast dataset. Each $2^{14} - 1$ labels is shown as a pot in the figure 5. X-axis is the $\Delta(y, y_i)$ and Y-axis is $1 + f_W(y) - f_W(y_i)$.

Appendix C Monotonicity of h and g in λ

Proof. Let $g_1 = g(y_{\lambda_1})$, $h_1 = h(y_{\lambda_1})$, $g_2 = g(y_{\lambda_2})$, and $h_2 = h(y_{\lambda_2})$.

$$\begin{aligned} h_1 + \lambda_1 g_1 &\geq h_2 + \lambda_1 g_2, & h_2 + \lambda_2 g_2 &\geq h_1 + \lambda_2 g_1 \\ \Leftrightarrow h_1 - h_2 + \lambda_1(g_1 - g_2) &\geq 0, & -h_1 + h_2 + \lambda_2(g_2 - g_1) &\geq 0 \\ \Leftrightarrow (g_2 - g_1)(\lambda_2 - \lambda_1) &\geq 0 \end{aligned}$$

For h , change the role of g and h . □

Appendix D Improvements for the binary search

Appendix D.1 Early stopping

If $L = [\lambda_m, \lambda_M]$, and both endpoints have the same label, i.e., $y_{\lambda_m} = y_{\lambda_M}$, then we can terminate the binary search safely because from lemma 4, it follows that the solution y_λ will not change in this segment.

Appendix D.2 Suboptimality bound

Let $K(\lambda)$ be the value of the λ -oracle. i.e.,

$$K(\lambda) = \max_{y \in \mathcal{Y}} h(y) + \lambda g(y). \quad (14)$$

Lemma 5. Φ^* is upper bounded by

$$\Phi(y^*) \leq \frac{K(\lambda)^2}{4\lambda} \quad (15)$$

Proof.

$$\begin{aligned}
 h(y) + \lambda g(y) &\leq K(\lambda) \\
 \iff g(y)(h(y) + \lambda g(y)) &\leq g(y)K(\lambda) \\
 \iff \Phi(y) &\leq g(y)K(\lambda) - \lambda g(y)^2 \\
 &= -\lambda \left(g(y) - \frac{K(\lambda)}{2\lambda} \right)^2 + \frac{K(\lambda)^2}{4\lambda} \leq \frac{K(\lambda)^2}{4\lambda}
 \end{aligned}$$

□

Appendix E Proof of the limitation of the λ -oracle search

Theorem 1. Let $\hat{H} = \max_y h(y)$ and $\hat{G} = \max_y g(y)$. For any $\epsilon > 0$, there exists a problem with 3 labels such that for any $\lambda \geq 0$, $y_\lambda = \operatorname{argmin}_{y \in \mathcal{Y}} \Phi(y) < \epsilon$, while $\Phi(y^*) = \frac{1}{4} \hat{H} \hat{G}$. Let $\hat{H} = \max_y h(y)$ and $\hat{G} = \max_y g(y)$. For any $\epsilon > 0$ and $\lambda > 0$, there exists a problem of 3 labels that $y_\lambda = \operatorname{argmin}_{y \in \mathcal{Y}} \Phi(y) < \epsilon$, and $\Phi(y^*) - \Phi(y_\lambda) = \frac{1}{4} \hat{H} \hat{G}$.

Proof. We will first prove following lemma which will be used in the proof.

Lemma 6. Let $A = [A_1 \ A_2] \in \mathbb{R}^2$, $B = [B_1 \ B_2] \in \mathbb{R}^2$, and $C = [C_1 \ C_2] \in \mathbb{R}^2$, and $A_1 < B_1 < C_1$. If B is under the line \overline{AC} , i.e., $\exists t, 0 \leq t \leq 1, D = tA + (1-t)C$, $D_1 = B_1$, $D_2 > B_2$. Then, $\forall \lambda \geq 0$, $v = [1 \ \lambda] \in \mathbb{R}^2$, such that

$$v \cdot B > v \cdot A \text{ and } v \cdot B > v \cdot C \quad (16)$$

Proof. Translate vectors A, B , and C into coordinates of $[0, A_2], [a, b], [C_1, 0]$ by adding a vector $[-A_1, -C_2]$ to each vectors A, B , and C , since it does not change $B - A$ or $B - C$. Let $X = C_1$ and $Y = A_2$.

If $0 \leq \lambda \leq \frac{X}{Y}$, then $v \cdot A = \lambda Y \leq X = v \cdot C$. $v \cdot (B - C) > 0 \iff (a - X) + \lambda b > 0$ corresponds to all the points above line \overline{AC} . Similarly, if $\lambda \geq \frac{X}{Y}$, (16) corresponds to $a + \lambda(b - Y) > 0$ is also all the points above \overline{AC} . □

From lemma 6, if $y_1, y_2 \in \mathcal{Y}$, then all the labels which lies under line y_1 and y_2 will not be found by λ -oracle. In the adversarial case, this holds when label lies on the line also. Therefore, Theorem 1 holds when there exists three labels, for arbitrary small $\epsilon > 0$, $A = [\epsilon, \hat{G}]$, $B = [\hat{H}, \epsilon]$, and $C = [\frac{1}{2}\hat{H}, \frac{1}{2}\hat{G}]$, $\mathcal{Y} = \{A, B, C\}$. In this case $\hat{\Phi} \approx 0$. □

Appendix F Angular search

We first introduce needed notations. $\partial^\perp(a)$ be the perpendicular slope of a , i.e., $\partial^\perp(a) = -\frac{1}{\partial(a)} = -\frac{a_1}{a_2}$. For $A \subseteq \mathbb{R}^2$, let label set restricted to A as $\vec{\mathcal{Y}}_A = \vec{\mathcal{Y}} \cap A$, and $y_{\lambda, A} = \mathcal{O}(\lambda, A) = \operatorname{argmax}_{y \in \mathcal{Y}, \vec{y} \in A} h(y) + \lambda g(y) = \operatorname{argmax}_{\vec{y} \in \vec{\mathcal{Y}}_A} [\vec{y}]_1 + \lambda [\vec{y}]_2$. Note that if $A = \mathbb{R}^2$, $y_{\lambda, \mathbb{R}^2} = y_\lambda$. For $P, Q \in \mathbb{R}^2$, define $\Lambda(P, Q)$ to be the area below the line \overline{PQ} , i.e., $\Lambda(P, Q) = \{\vec{y} \in \mathbb{R}^2 \mid [\vec{y}]_2 - [P]_2 \leq \partial^\perp(Q - P)([\vec{y}]_2 - [P]_2)\}$. $\Upsilon_\lambda = \{\vec{y} \in \mathbb{R}^2 \mid \vec{\Phi}(\vec{y}) = [\vec{y}]_1 \cdot [\vec{y}]_2 \geq \vec{\Phi}(\vec{y}_{\lambda, A})\}$ be the area above C_λ , and $\underline{\Upsilon}_\lambda = \{\vec{y} \in \mathbb{R}^2 \mid \vec{\Phi}(\vec{y}) = [\vec{y}]_1 \cdot [\vec{y}]_2 \leq \vec{\Phi}(\vec{y}_{\lambda, A})\}$ be the area below C_λ .

Recall the *constrained λ -oracle* defined in (8):

$$y_{\lambda, \alpha, \beta} = \mathcal{O}_c(\lambda, \alpha, \beta) = \max_{y \in \mathcal{Y}, \alpha h(y) \geq g(y), \beta h(y) < g(y)} \mathcal{L}_\lambda(y)$$

where $\alpha, \beta \in \mathbb{R}_+$ and $\alpha \geq \beta > 0$. Let $A(\alpha, \beta) \subseteq \mathbb{R}^2$ be the restricted search space, i.e., $A(\alpha, \beta) = \{a \in \mathbb{R}^2 \mid \beta < \partial(a) \leq \alpha\}$. Constrained λ -oracle reveals maximal \mathcal{L}_λ label within restricted area defined by α and β . The area is bounded by two lines whose slope is α and β . Define a pair (α, β) , $\alpha, \beta \in \mathbb{R}_+$, $\alpha \geq \beta > 0$ as an *angle*. The angular search recursively divides an angle into two different angles, which we call the procedure as a *split*. For $\alpha \geq \beta \geq 0$, let $\lambda = \frac{1}{\sqrt{\alpha\beta}}$, $z = \vec{y}_{\lambda, \alpha, \beta}$ and $z' = [\lambda[z]_2, \frac{1}{\lambda}[z]_1]$. Let P be the point among z and z' which has the greater slope (any if two equal), and Q be the other

point, i.e., if $\partial(z) \geq \partial(z')$, $P = z$ and $Q = z'$, otherwise $P = z'$ and $Q = z$. Let $R = \left[\sqrt{\lambda[z]_1 \cdot [z]_2} \sqrt{\frac{1}{\lambda}[z]_1 \cdot [z]_2} \right]$. Define $\text{split}(\alpha, \beta)$ as a procedure divides (α, β) into two angles $(\alpha^+, \gamma^+) = (\partial(P), \partial(R))$ and $(\gamma^+, \beta^+) = (\partial(R), \partial(Q))$. First, show that $\partial(P)$ and $\partial(Q)$ are in between α and β , and $\partial(R)$ is between $\partial(P)$ and $\partial(Q)$.

Lemma 7. For each $\text{split}(\alpha, \beta)$,

$$\beta \leq \partial(Q) \leq \partial(R) \leq \partial(P) \leq \alpha$$

Proof. $\beta \leq \partial(z) \leq \alpha$ follows from the definition of constrained λ -oracle in (8).

$$\partial(z') = \frac{1}{\lambda^2 \partial(z)} = \frac{\alpha\beta}{\partial(z)} \implies \beta \leq \partial(z') \leq \alpha \implies \beta \leq \partial(Q) \leq \partial(P) \leq \alpha.$$

$$\partial(Q) \leq \partial(R) \leq \partial(P) \iff \min \left\{ \partial(z), \frac{1}{\lambda^2 \partial(z)} \right\} \leq \frac{1}{\lambda} \leq \max \left\{ \partial(z), \frac{1}{\lambda^2 \partial(z)} \right\} \text{ from } \forall a, b \in \mathbb{R}_+, b \leq a \implies b \leq \sqrt{ab} \leq a. \quad \square$$

After each split, the union of the divided angle (α^+, γ) and (γ, β^+) can be smaller than angle (α, β) . However, following lemma shows it is safe to use (α^+, γ) and (γ, β^+) when our objective is to find y^* .

Lemma 8.

$$\forall a \in \vec{\mathcal{Y}}_{A(\alpha, \beta)}, \Phi(a) > \Phi(y_{\lambda, \alpha, \beta}) \implies \beta^+ < \partial(a) < \alpha^+$$

Proof. From lemma 2, $\vec{\mathcal{Y}}_{A(\alpha, \beta)} \subseteq \Lambda(P, Q)$. Let $U = \{a \in \mathbb{R}^2 | \partial(a) \geq \alpha_+ = \partial(P)\}$, $B = \{a \in \mathbb{R}^2 | \partial(a) \leq \beta_+ = \partial(Q)\}$, and two contours of function $C = \{a \in \mathbb{R}^2 | \vec{\Phi}(a) = \Phi(y_{\lambda, \alpha, \beta})\}$, $S = \{a \in \mathbb{R}^2 | \mathcal{L}_\lambda(a) = \mathcal{L}_\lambda(\vec{y}_{\lambda, \alpha, \beta})\}$. S is the upper bound of $\Lambda(P, Q)$, and C is the upper bound of $\underline{C} = \{a \in \mathbb{R}^2 | \vec{\Phi}(a) \leq \Phi(y_{\lambda, \alpha, \beta})\}$. P and Q are the intersections of C and S . For area of U and B , S is under C , therefore, $\Lambda(P, Q) \cap U \subseteq \underline{C}$, and $\Lambda(P, Q) \cap B \subseteq \underline{C}$. It implies that $\forall a \in (\Lambda(P, Q) \cap U) \cup (\Lambda(P, Q) \cap B) \implies \vec{\Phi}(a) \leq \Phi(y_{\lambda, \alpha, \beta})$. And the lemma follows from $A(\alpha, \beta) = U \cup B \cup \{a \in \mathbb{R}^2 | \beta^+ < \partial(a) < \alpha^+\}$. \square

We associate a quantity we call a *capacity* of an angle, which is used to prove the suboptimality of the algorithm. For an angle (α, β) , the capacity of an angle $v(\alpha, \beta)$ is

$$v(\alpha, \beta) := \sqrt{\frac{\alpha}{\beta}}$$

Note that from the definition of an angle, $v(\alpha, \beta) \geq 1$. First show that the capacity of angle decreases exponentially for each split.

Lemma 9. For any angle (α, β) and its split (α^+, γ^+) and (γ^+, β^+) ,

$$v(\alpha, \beta) \geq v(\alpha^+, \beta^+) = v(\alpha^+, \gamma^+)^2 = v(\gamma^+, \beta^+)^2$$

Proof. Assume $\partial(P) \geq \partial(Q)$ (the other case is follows the same proof with changing the role of P and Q), then $\alpha^+ = \partial(P)$ and $\beta^+ = \partial(Q)$. $\partial(Q) = \frac{1}{\lambda^2 \partial(P)} = \frac{\alpha\beta}{\partial(P)}$, $v(\alpha^+, \beta^+) = v(\partial(P), \partial(Q)) = \lambda \partial(P) = \frac{\partial(P)}{\sqrt{\alpha\beta}}$. Since α is the upper bound and β is the lower bound of $\partial(P)$, $\sqrt{\frac{\beta}{\alpha}} \leq v(\partial(P), \partial(Q)) \leq \sqrt{\frac{\alpha}{\beta}}$. Last two equalities in the lemma are from $v(\partial(P), \partial(R)) = v(\partial(R), \partial(Q)) = \sqrt{\frac{\partial(P)}{\sqrt{\alpha\beta}}}$ by plugging in the coordinate of R . \square

Lemma 10. Let $\mathcal{B}(a) = \frac{1}{4} \left(a + \frac{1}{a} \right)^2$. The suboptimality bound of an angle (α, β) with $\lambda = \frac{1}{\sqrt{\alpha\beta}}$ is

$$\frac{\max_{\vec{y} \in \vec{\mathcal{Y}}_{A(\alpha, \beta)}} \vec{\Phi}(\vec{y})}{\Phi(y_{\lambda, \alpha, \beta})} \leq \mathcal{B}(v(\alpha, \beta)).$$

Proof. From lemma 2, $\vec{\mathcal{Y}}_{A(\alpha,\beta)} \subseteq \Lambda(P,Q) = \Lambda(z,z')$. Let $\partial(z) = \gamma$. From 7, $\beta \leq \gamma \leq \alpha$. Let $m = \operatorname{argmax}_{a \in \Lambda(z,z')} \vec{\Phi}(a)$. m is on line $\overline{zz'}$ otherwise we can move m increasing direction of each axis till it meets the boundary $\overline{zz'}$ and Φ only increases, thus $m = tz + (1-t)z'$. $\vec{\Phi}(m) = \max_t \vec{\Phi}(tz + (1-t)z')$. $\frac{\partial \vec{\Phi}(tz + (1-t)z')}{\partial t} = 0 \implies t = \frac{1}{2}$. $m = \frac{1}{2}[z_1 + \lambda z_2 \quad z_2 + \frac{z_1}{\lambda}]$.

$$\begin{aligned} \frac{\max_{\vec{y} \in \vec{\mathcal{Y}}_{A(\alpha,\beta)}} \vec{\Phi}(\vec{y})}{\Phi(y_{\lambda,\alpha,\beta})} &= \frac{1}{4} \left(\sqrt{\frac{z_1}{\lambda z_2}} + \sqrt{\frac{\lambda z_2}{z_1}} \right)^2 \\ &= \frac{1}{4} \left(\sqrt{\frac{\sqrt{\alpha\beta}}{\gamma}} + \sqrt{\frac{\gamma}{\sqrt{\alpha\beta}}} \right)^2 \end{aligned}$$

Since $v(a) = v\left(\frac{1}{a}\right)$ and $v(a)$ increases monotonically for $a \geq 1$,

$$\mathcal{B}(a) \leq \mathcal{B}(b) \iff \max\left\{a, \frac{1}{a}\right\} \leq \max\left\{b, \frac{1}{b}\right\}$$

If $\frac{\sqrt{\alpha\beta}}{\gamma} \geq \frac{\gamma}{\sqrt{\alpha\beta}}$, then $\frac{\sqrt{\alpha\beta}}{\gamma} \leq \sqrt{\frac{\alpha}{\beta}}$ since $\gamma \geq \beta$. If $\frac{\gamma}{\sqrt{\alpha\beta}} \geq \frac{\sqrt{\alpha\beta}}{\gamma}$, then $\frac{\gamma}{\sqrt{\alpha\beta}} \leq \sqrt{\frac{\alpha}{\beta}}$ since $\gamma \leq \alpha$. Therefore, $\frac{\max_{\vec{y} \in \vec{\mathcal{Y}}_{A(\alpha,\beta)}} \vec{\Phi}(\vec{y})}{\Phi(y_{\lambda,\alpha,\beta})} = \mathcal{B}\left(\frac{\sqrt{\alpha\beta}}{\gamma}\right) \leq \mathcal{B}(v(\alpha, \beta))$. \square

Now we can prove the theorems.

Theorem 2. Angular search described in algorithm 2 finds optimum $y^* = \operatorname{argmax}_{y \in \mathcal{Y}} \Phi(y)$ at most $t = 2M + 1$ iteration where M is the number of the labels.

Proof. Denote $y_t, \alpha_t, \beta_t, z_t, z'_t, K_t^1$, and K_t^2 for $y, \alpha, \beta, z, z', K^1$, and K^2 at iteration t respectively. $\mathcal{A}(\alpha_t, \beta_t)$ is the search space at each iteration t . At the first iteration $t = 1$, the search space contains all the labels with positive Φ , i.e., $\{y | \Phi(y) \geq 0\} \subseteq \mathcal{A}(\infty, 0)$. At iteration $t > 1$, firstly, when $y_t = \emptyset$, the search area $\mathcal{A}(\alpha_t, \beta_t)$ is removed from the search since $y_t = \emptyset$ implies there is no label inside $\mathcal{A}(\alpha_t, \beta_t)$. Secondly, when $y_t \neq \emptyset$, $\mathcal{A}(\alpha_t, \beta_t)$ is dequeued, and K_t^1 and K_t^2 is enqueued. From lemma 8, at every step, we are ensured that do not loose y^* . By using strict inequalities in the constrained oracle with valuable s , we can ensure y_t which oracle returns is an unseen label. Note that split only happens if a label is found, i.e., $y_t \neq \emptyset$. Therefore, there can be only M splits, and each split can be viewed as a branch in the binary tree, and the number of queries are the number of nodes. Maximum number of the nodes with M branches are $2M + 1$. \square

Theorem 4. In angular search, described in Algorithm 2, at iteration t ,

$$\Phi(\hat{y}^t) \geq \Phi(y^*) (v_1^{-\frac{4}{t+1}})$$

where $\hat{y}^t = \operatorname{argmax}_t y^t$ is the optimum up to t , $v_1 = \max\left\{\frac{\lambda_0}{\partial(y_1)}, \frac{\partial(y_1)}{\lambda_0}\right\}$, λ_0 is the initial λ used, and y_1 is the first label returned by the constrained λ -oracle.

Proof. After $t \geq 2^r - 1$ iteration as in algorithm 2 where r is an integer, for all the angle (α, β) in the queue Q , $v(\alpha, \beta) \leq (v_1)^{2^{1-r}}$. This follows from the fact that since the algorithm uses the depth first search, after $2^r - 1$ iterations all the nodes at the search is at least r . At each iteration, for a angle, the capacity is square rooted from the lemma 9, and the depth is increased by one. And the theorem follows from the fact that after $t \geq 2^r - 1$ iterations, all splits are at depth $r' \geq r$, and at least one of the split contains the optimum with suboptimality bound with lemma 10. Thus,

$$\frac{\Phi(y^*)}{\Phi(\hat{y})} \leq \mathcal{B}\left((v_1)^{2^{1-r}}\right) < (v_1)^{2^{2-r}} \leq (v_1)^{\frac{4}{t+1}}$$

\square

Theorem 5. Assuming $\Phi(y^*) > \phi$, angular search described in algorithm 2 with $\lambda_0 = \frac{\hat{G}}{\hat{H}}$, $\alpha_0 = \frac{\hat{G}^2}{\phi}$, $\beta_0 = \frac{\phi}{\hat{H}^2}$, finds ϵ -optimal solution, $\Phi(y) \geq (1 - \epsilon)\Phi(y^*)$, in T queries and $O(T)$ operations where $T = 4 \log \left(\frac{\hat{G}\hat{H}}{\phi} \right) \cdot \frac{1}{\epsilon}$, and δ -optimal solution, $\Phi(y) \geq \Phi(y^*) - \delta$, in T' queries and $O(T')$ operations where $T' = 4 \log \left(\frac{\hat{G}\hat{H}}{\phi} \right) \cdot \frac{\Phi(y^*)}{\delta}$.

Proof. $\Phi(y^*) > \phi \Leftrightarrow \frac{\phi}{\hat{H}^2} < \frac{g(y^*)}{h(y^*)} = \partial(y^*) < \frac{\hat{G}^2}{\phi}$. $v_1 = \max \left\{ \frac{\lambda_0}{\partial(y_1)}, \frac{\partial(y_1)}{\lambda_0} \right\}$ from Theorem 4. Algorithm finds y^* if $\beta \leq \partial(y^*) \leq \alpha$, thus set $\alpha = \frac{\hat{G}^2}{\phi}$ and $\beta = \frac{\phi}{\hat{H}^2}$. Also from the definition of constrained λ -oracle, $\beta = \frac{\phi}{\hat{H}^2} \leq \partial(y_1) \leq \alpha = \frac{\hat{G}^2}{\phi}$. Therefore, $v_1 \leq \max \left\{ \frac{\lambda_0}{\partial(y_1)}, \frac{\partial(y_1)}{\lambda_0} \right\}$. And the upper bound of two terms equal when $\lambda_0 = \frac{\hat{G}}{\hat{H}}$, then $v_1 \leq \frac{\hat{G}\hat{H}}{\phi}$. δ bound follows plugging in the upper bound of v_1 , and $\epsilon = \frac{\delta}{\Phi(y^*)}$. \square

Appendix G Illustration of the angular search

Following figure 6 illustrates Angular search. Block dots are the labels from figure 5. Blue X denotes the new label returned by the oracle. Red X is the maximum point. Two straight lines are the upper bound and the lower bound used by the constrained oracle. Constrained oracle returns a blue dot between the upper and lower bounds. We can draw a line that passes blue X that no label can be above the line. Then, split the angle into half. This process continues until the y^* is found.




Figure 6: Illustration of the Angular search.

Appendix H Limitation of the constraint λ -oracle search

Theorem 3. Any search algorithm accessing labels only through λ -oracle with any number of the linear constraints cannot find y^* in less than M iterations in the worst case where M is the number of labels.

Proof. We show this in the perspective of a game between a searcher and an oracle. At each iteration, the searcher query the oracle with λ and the search space denoted as \mathcal{A} , and the oracle reveals a label according to the query. And the claim is that with any choice of $M - 1$ queries, for each query the oracle can either give an consistent label or indicate that there is no label in \mathcal{A} such that after $M - 1$ queries the oracle provides an unseen label y^* which has bigger Φ than all previous revealed labels.

Denote each query at iteration t with $\lambda_t > 0$ and a query closed and convex set $\mathcal{A}_t \subseteq \mathbb{R}^2$, and denote the revealed label at iteration t as y_t . We will use $y_t = \emptyset$ to denote that there is no label inside query space \mathcal{A}_t . Let $\mathcal{Y}_t = \{y_{t'} | t' < t\}$.

Algorithm 3 describes the pseudo code for generating such y_t . The core of the algorithm is maintaining a rectangular area \mathcal{R}_t for each iteration t with following properties. Last two properties are for y_t .

Algorithm 3 Construct a consistent label set \mathcal{Y} .

Input: $\{\lambda_t, \mathcal{A}_t\}_{t=1}^{M-1}, \lambda_t > 0, \mathcal{A}_t \subseteq \mathbb{R}^2, \mathcal{A}_t$ is closed and convex region.

Output: $\{y_t \in \mathbb{R}^2\}_{t=1}^{M-1}, y^* \in \mathbb{R}^2$
Initialize: $\mathcal{R}_0 = \{(a, b) | 0 < a, 0 < b\}, \mathcal{Y}_0 = \emptyset$.

```

1: for  $t = 1, 2, \dots, M - 1$  do
2:   if  $\mathcal{Y}_{t-1} \cap \mathcal{A}_t = \emptyset$  then
3:      $\tilde{y} = \operatorname{argmax}_{y \in \mathcal{Y}_t} h(y) + \lambda_t g(y)$ .
4:      $\tilde{\mathcal{R}} = \mathcal{R}_{t-1} \cap \{y | h(y) + \lambda_t g(y) < h(\tilde{y}) + \lambda_t g(\tilde{y}) \text{ or } y \notin \mathcal{A}_t\}$ .
5:   else
6:      $\tilde{y} = \emptyset, \tilde{\mathcal{R}} = \mathcal{R}_{t-1} - \mathcal{A}_t$ .
7:   if  $\tilde{\mathcal{R}} \neq \emptyset$  then
8:      $y_t = \emptyset, \mathcal{R}_t = \operatorname{FindRect}(\tilde{\mathcal{R}})$ 
9:   else
10:     $y_t = \operatorname{FindPoint}(\mathcal{R}_{t-1}, \lambda_t)$ .
11:     $\mathcal{R}_t = \operatorname{FindRect}(\operatorname{Shrink}(\mathcal{R}_{t-1}, y_t, \lambda_t))$ .
12:   if  $y_t \neq \emptyset$  then
13:      $\mathcal{Y} = \mathcal{Y} \cup \{y_t\}$ .
14: Pick any  $y^* \in \mathcal{R}_{M-1}$ 
    
```

1. $\forall t' < t, \forall y \in \mathcal{R}_t, \Phi(y) > \Phi(y_{t'})$.
2. $\forall t' < t, \forall y \in \mathcal{R}_t \cap \mathcal{A}_{t'}, h(y_{t'}) + \lambda_{t'} g(y_{t'}) > h(y) + \lambda_{t'} g(y)$.
3. $\mathcal{R}_t \subseteq \mathcal{R}_{t-1}$.
4. \mathcal{R}_t is a non-empty open set.
5. $y_t \in \mathcal{R}_t \cap \mathcal{A}_t$
6. $y_t = \operatorname{argmax}_{y \in \mathcal{Y}_t \cap \mathcal{A}_t} h(y) + \lambda_t g(y)$.

Note that if these properties holds till iteration M , we can simply set y^* as any label in \mathcal{R}_M which proves the claim.

First, we show that property 4 is true. \mathcal{R}_0 is a non-empty open set. Consider iteration t , and assume \mathcal{R}_{t-1} is a non-empty open set. Then $\tilde{\mathcal{R}}$ is an open set since \mathcal{R}_{t-1} is an open set. There are two unknown functions, *Shrink* and *FindRect*. For open set $A \subseteq \mathbb{R}^2, y \in \mathbb{R}^2$, let $\operatorname{Shrink}(A, y, \lambda) = A - \{y' | \Phi(y') \leq \Phi(y) \text{ or } h(y') + \lambda g(y') \geq h(y) + \lambda g(y)\}$. Note that $\operatorname{Shrink}(A, y, \lambda) \subseteq A$, and $\operatorname{Shrink}(A, y, \lambda)$ is an open set. Assume now that there exists a y such that $\operatorname{Shrink}(\mathcal{R}_{t-1}, y, \lambda_t) \neq \emptyset$ and $\operatorname{FindPoint}(\mathcal{R}_{t-1}, \lambda_t)$ returns such y . Function *FindPoint* will be given later. $\operatorname{FindRect}(A)$ returns an open non-empty rectangle inside A . Note that $\operatorname{Rect}(A) \subseteq A$, and since input to *Rect* is always non empty open set, such rectangle exists. Since \mathcal{R}_0 is non-empty open set, $\forall t, \mathcal{R}_t$ is a non-empty open set.

Property 3 and 5 are easy to check. Property 1 and 2 follows from the fact that $\forall t \in \{t | y_t \neq \emptyset\}, \forall t' > t, \mathcal{R}_{t'} \subseteq \operatorname{Shrink}(\mathcal{R}_{t-1}, y_t, \lambda_{t-1})$.

Property 6 follows from the facts that if $\mathcal{Y}_{t-1} \cap \mathcal{A}_t \neq \emptyset, \tilde{\mathcal{R}} = \emptyset \implies \mathcal{R}_{t-1} \subseteq \{y | h(y) + \lambda_t g(y) > h(\tilde{y}) + \lambda_t g(\tilde{y}) \text{ and } y \in \mathcal{A}_t\}$, otherwise $\mathcal{Y}_{t-1} \cap \mathcal{A}_t = \emptyset$, and $\mathcal{R}_{t-1} \subseteq \mathcal{A}_t$.

$\operatorname{FindPoint}(A, \lambda)$ returns any $y \in A - \{y \in \mathbb{R}^2 | \lambda y_2 = y_1\}$. Given input A is always an non-empty open set, such y exists. $\operatorname{Shrink}(\mathcal{R}_{t-1}, y, \lambda_t) \neq \emptyset$ is ensured from the fact that two boundaries, $c = \{y' | \Phi(y') = \Phi(y)\}$ and $d = \{h(y') + \lambda g(y') = h(y) + \lambda g(y)\}$ meets at y . Since c is a convex curve, c is under d on one side. Therefore the intersection of set above c and below d is non-empty and also open. \square

Appendix I Additional Plots from the Experiments

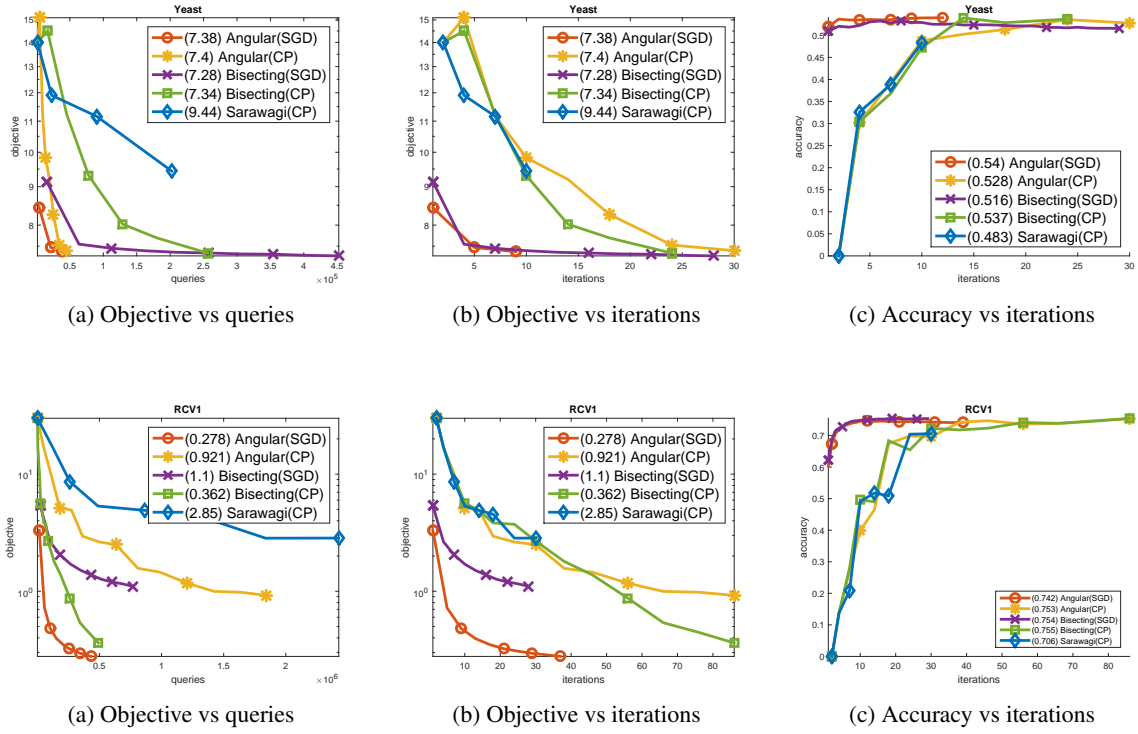


Figure 8: Additional experiment plot (RCV)