Loss-sensitive Training of Probabilistic Conditional Random Fields

Maksims N. Volkovs
Department of Computer Science
University of Toronto
Toronto, Canada

Hugo Larochelle
Département d'informatique
Université de Sherbrooke
Sherbrooke, Canada

Richard S. Zemel
Department of Computer Science
University of Toronto
Toronto, Canada

Abstract

We consider the problem of training probabilistic conditional random fields (CRFs) in the context of a task where performance is measured using a specific loss function. While maximum likelihood is the most common approach to training CRFs, it ignores the inherent structure of the task's loss function. We describe alternatives to maximum likelihood which take that loss into account. These include a novel adaptation of a loss upper bound from the structured SVMs literature to the CRF context, as well as a new loss-inspired KL divergence objective which relies on the probabilistic nature of CRFs. These loss-sensitive objectives are compared to maximum likelihood using ranking as a benchmark task. This comparison confirms the importance of incorporating loss information in the probabilistic training of CRFs, with the loss-inspired KL outperforming all other objectives.

1 Introduction

Conditional random fields (CRFs) [1] form a flexible family of models for capturing the interaction between an input $\mathbf{x}$ and a target $\mathbf{y}$. CRFs have been designed for a vast variety of problems, including natural language processing [2, 3, 4], speech processing [5], computer vision [6, 7, 8] and bioinformatics [9, 10] tasks. One reason for their popularity is that they provide a flexible framework for modeling the conditional distributions of targets constrained by some specific structure, such as chains [1], trees [11], 2D grids [7, 12], permutations [13] and many more.

While there has been a lot of work on developing appropriate CRF potentials and energy functions, as well as on deriving efficient (approximate) inference procedures for some given target structure, much less attention has been paid to the loss function under which the CRF's performance is ultimately evaluated. Indeed, CRFs are usually trained by maximum likelihood (ML) or the maximum a posteriori criterion (MAP or regularized ML), which ignores the task's loss function. Yet, several tasks are associated with loss functions[1] that are also structured and do not correspond to a simple 0/1 loss: labelwise error (Hamming loss) for item labeling, BLEU score for machine translation, normalized discounted cumulative gain (NDCG) for ranking, etc. Ignoring this structure can prove as detrimental to performance as ignoring the target's structure.

The inclusion of loss information into learning is an idea that has been more widely explored in the context of structured support vector machines (SSVMs) [14, 15]. SSVMs and CRFs are closely related models, both trying to shape an energy or score function over the joint input and target space to fit the available training data. However, while an SSVM attempts to satisfy margin constraints without invoking a probabilistic interpretation of the model, a CRF follows a probabilistic approach and instead aims at calibrating its probability estimates to the data. Similarly, while an SSVMs relies

---

[1]Without loss of generality, for tasks where a performance measure is instead provided (i.e. where higher values is better), we assume it can be converted into a loss, e.g. by setting the loss to the negative of the performance measure.

on maximization procedures to identify the most violated margin constraints, a CRF relies on (approximate) inference or sampling procedures to estimate probabilities under its distribution and compare it to the empirical distribution.

While there are no obvious reasons to prefer one approach to the other, a currently unanswered question is whether the known methods that adapt SSVM training to some given loss (i.e., upper bounds based on margin and slack scaling [15]) can also be applied to the probabilistic training of CRFs. Another question is how such methods would compare to other loss-sensitive training objectives which rely on the probabilistic nature of CRFs and which may have no analog in the SSVM framework.

We investigate these questions in this paper. First, we describe upper bounds similar to the margin and slack scaling upper bounds of SSVMs, but that correspond to maximum likelihood training of CRFs with loss-augmented and loss-scaled energy functions. Second, we describe two other loss-inspired training objectives for CRFs which rely on the probabilistic nature of CRFs: the standard average expected loss objective and a novel loss-inspired KL-divergence objective. Finally, we compare these loss-sensitive objectives on ranking benchmarks based on the NDCG performance measure. To our knowledge, this is the first systematic evaluation of loss-sensitive training objectives for probabilistic CRFs.

2 Conditional Random Fields

This work is concerned with the general problem of supervised learning, where the relationship between an input $\mathbf{x}$ and a target $\mathbf{y}$ must be learned from a training set of instantiated pairs $\mathcal{D} = \{\mathbf{x}_t, \mathbf{y}_t\}$. More specifically, we are interested in learning a predictive mapping from $\mathbf{x}$ to $\mathbf{y}$.

Conditional random fields (CRFs) tackle this problem by defining directly the conditional distribution $p(\mathbf{y}|\mathbf{x})$ through some energy function $E(\mathbf{y}, \mathbf{x}; \theta)$ as follows:

$$p(\mathbf{y}|\mathbf{x}) = \exp(-E(\mathbf{y}, \mathbf{x}; \theta))/Z(\mathbf{x}), \ \ Z(\mathbf{x}) = \sum_{\mathbf{y} \in \mathcal{Y}(\mathbf{x})} \exp(-E(\mathbf{y}, \mathbf{x}; \theta)) \tag{1}$$

where $\mathcal{Y}(\mathbf{x})$ is the set of all possible configurations for $\mathbf{y}$ given the input $\mathbf{x}$ and θ is the model's parameter vector. The parametric form of the energy function $E(\mathbf{y}, \mathbf{x}; \theta)$ will depend on the nature of $\mathbf{x}$ and $\mathbf{y}$. A popular choice is that of a linear function of a set of features on $\mathbf{x}$ and $\mathbf{y}$, i.e., $E(\mathbf{y}, \mathbf{x}; \theta) = -\sum_i \theta_i f_i(\mathbf{x}, \mathbf{y})$.

2.1 Maximum Likelihood Objective

The most popular approach to training CRFs is conditional maximum likelihood. It corresponds to the minimization with respect to θ of the objective $\mathcal{L}_{\mathrm{ML}}(\mathcal{D}; \theta)$:

$$-\frac{1}{|\mathcal{D}|} \sum_{(\mathbf{x}_t, \mathbf{y}_t) \in \mathcal{D}} \log p(\mathbf{y}_t|\mathbf{x}_t) = \frac{1}{|\mathcal{D}|} \sum_{(\mathbf{x}_t, \mathbf{y}_t) \in \mathcal{D}} E(\mathbf{y}_t, \mathbf{x}_t; \theta) + \log \left(\sum_{\mathbf{y} \in \mathcal{Y}(\mathbf{x})} \exp(-E(\mathbf{y}, \mathbf{x}_t; \theta)) \right). \tag{2}$$

To this end, one can use any gradient-based optimization procedure, which can convergence to a local optimum, or even a global optimum if the problem is convex (e.g., by choosing an energy function $E(\mathbf{y}, \mathbf{x}; \theta)$ linear in θ). The gradients have an elegant form:

$$\frac{\partial \mathcal{L}_{\mathrm{ML}}(\mathcal{D}; \theta)}{\partial \theta} = \frac{1}{|\mathcal{D}|} \sum_{(\mathbf{x}_t, \mathbf{y}_t) \in \mathcal{D}} \frac{\partial E(\mathbf{y}_t, \mathbf{x}_t; \theta)}{\partial \theta} - \mathbf{E}_{\mathbf{y}|\mathbf{x}_t} \left[\frac{\partial E(\mathbf{y}, \mathbf{x}_t; \theta)}{\partial \theta} \right]. \tag{3}$$

Hence exact gradient evaluations are possible when the conditional expectation in the second term is tractable. This is the case for CRFs with a chain or tree structure, for which belief propagation can be used. When gradients are intractable, two

approximate alternatives can be considered. The first is to approximate the intractable expectation using either Markov chain Monte Carlo sampling or variational inference algorithms such as mean-field or loopy belief propagation, the latter being the most popular. The second approach is to use alternative objectives such as pseudolikelihood [16] or piece-wise training [17][2].

3 Loss-sensitive Training Objectives

Unfortunately, maximum likelihood and its associated approximations all suffer from the problem that the loss function under which the performance of the CRF is evaluated is ignored. In the well-specified case and for large datasets, this would probably not be a problem because of the asymptotic consistency and efficiency properties of maximum likelihood. However, almost all practical problems do not fall in the well-specified setting, which justifies the exploration of alternative training objectives.

Let $\hat{\mathbf{y}}(\mathbf{x}_t)$ denote the prediction made by a CRF for some given input $\mathbf{x}_t$. Most commonly[3], this prediction will be $\hat{\mathbf{y}}(\mathbf{x}_t) = \arg\max_{\mathbf{y} \in \mathcal{Y}(\mathbf{x}_t)} p(\mathbf{y}|\mathbf{x}_t) = \arg\min_{\mathbf{y} \in \mathcal{Y}(\mathbf{x}_t)} E(\mathbf{y}, \mathbf{x}_t)$. We assume that we are given some loss $l_t(\hat{\mathbf{y}}(\mathbf{x}_t))$ under which the performance of the CRF on some dataset $\mathcal{D}$ will be measured. We will also assume that $l_t(\mathbf{y}_t) = 0$. The goal is then to achieve a low average $\frac{1}{|\mathcal{D}|} \sum_{(\mathbf{x}_t, \mathbf{y}_t) \in \mathcal{D}} l_t(\hat{\mathbf{y}}(\mathbf{x}_t))$ under that loss.

Directly minimizing this average loss is hard, because $l_t(\hat{\mathbf{y}}(\mathbf{x}_t))$ is not a smooth function of the CRF parameters θ. In fact, the loss itself $l_t(\hat{\mathbf{y}}(\mathbf{x}_t))$ is normally not a smooth function of the prediction $\hat{\mathbf{y}}(\mathbf{x}_t)$, and $\hat{\mathbf{y}}(\mathbf{x})$ is also not a smooth function of the model parameters θ. This non-smoothness makes it impossible to apply gradient-based optimization.

However, one could attempt to indirectly optimize the average loss by deriving smooth objectives that also depend on the loss. In the next sections, we describe three separate formulations of this approach.

3.1 Loss Upper Bounds

The loss function provides important information as to how good a potential prediction $\mathbf{y}$ is with respect to the ground truth $\mathbf{y}_t$. In particular, it specifies an ordering from the best prediction ($\mathbf{y} = \mathbf{y}_t$) to increasingly bad predictions with increasing value of their associated loss $l_t(\mathbf{y})$. It might then be desirable to ensure that the CRF assigns particularly low probability (i.e., high energy) to the worst possible predictions, as measured by the loss.

A first way of achieving this is to augment the energy function at a given training example $(\mathbf{x}_t, \mathbf{y}_t)$ by including the loss function for that example, producing a *Loss-Augmented* energy:

$$E_t^{\text{LA}}(\mathbf{y}, \mathbf{x}_t; \theta) = E(\mathbf{y}, \mathbf{x}_t; \theta) - l_t(\mathbf{y}) . \tag{4}$$

By artificially reducing the energy of bad values of $\mathbf{y}$ as a function of their loss, this will force the CRF to increase even more the value of $E(\mathbf{y}, \mathbf{x}; \theta)$ for those values of $\mathbf{y}$ with high loss. This idea is similar to the concept of margin re-scaling in structured support vector machines (SSVMs) [14, 15], a similarity that has been highlighted previously by Hazan and Urtasun [18]. Moreover, as in SSVMs, it can be shown that by replacing the regular energy function with this loss-augmented energy function in the maximum likelihood objective of Equation 2, we obtain a new Loss-Augmented

---

[2]Variational inference-based training can also be interpreted as training based on a different objective.

[3]For loss functions that decompose into loss terms over subsets of target variables, it may be more appropriate to use the mode of the marginals over each subset as the prediction.

objective that upper bounds the average loss:

$$
\begin{aligned}
\mathcal{L}_{\mathrm{LA}}(\mathcal{D};\theta) &= \frac{1}{|\mathcal{D}|}\sum_{(\mathbf{x}_t,\mathbf{y}_t)\in\mathcal{D}} E_t^{\mathrm{LA}}(\mathbf{y}_t,\mathbf{x}_t) + \log\left(\sum_{\mathbf{y}\in\mathcal{Y}(\mathbf{x})}\exp(-E_t^{\mathrm{LA}}(\mathbf{y},\mathbf{x}_t))\right) \\
&\geq \frac{1}{|\mathcal{D}|}\sum_{(\mathbf{x}_t,\mathbf{y}_t)\in\mathcal{D}} E_t^{\mathrm{LA}}(\mathbf{y}_t,\mathbf{x}_t) + \log\left(\exp(-E_t^{\mathrm{LA}}(\hat{\mathbf{y}}(\mathbf{x}_t),\mathbf{x}_t))\right) \\
&= \frac{1}{|\mathcal{D}|}\sum_{(\mathbf{x}_t,\mathbf{y}_t)\in\mathcal{D}} E(\mathbf{y}_t,\mathbf{x}_t) - E(\hat{\mathbf{y}}(\mathbf{x}_t),\mathbf{x}_t) + l_t(\hat{\mathbf{y}}(\mathbf{x}_t)) - l_t(\mathbf{y}_t) \\
&\geq \frac{1}{|\mathcal{D}|}\sum_{(\mathbf{x}_t,\mathbf{y}_t)\in\mathcal{D}} l_t(\hat{\mathbf{y}}(\mathbf{x}_t)) \, .
\end{aligned}
$$

We see that the higher $l_t(\mathbf{y})$ is for some given $\mathbf{y}$, the more important the energy term associated with it will be in the global objective. Hence, introducing the loss this way will indeed force the optimization to focus more on increasing the energy for configurations of $\mathbf{y}$ associated with high loss.

As an alternative to subtracting the loss, we could further increase the weight of terms associated with high loss by also multiplying the original energy function, as follows:

$$
E_t^{\mathrm{LS}}(\mathbf{y},\mathbf{x}_t;\theta) = l_t(\mathbf{y})(E(\mathbf{y},\mathbf{x}_t;\theta) - E(\mathbf{y}_t,\mathbf{x}_t;\theta)) - l_t(\mathbf{y}) \, . \tag{5}
$$

The advantage of this *Loss-Scaled* energy is that when a configuration of $\mathbf{y}$ with high loss already has higher energy than the target (i.e., $E(\mathbf{y},\mathbf{x}_t;\theta) - E(\mathbf{y}_t,\mathbf{x}_t;\theta) > 0$), then the energy is going to be further increased, reducing its weight in the optimization. In other words, focus in the optimization is put on bad configurations of $\mathbf{y}$ only when they have lower energy than the target. Finally, we can also show that the Loss-Scaled objective obtained from this loss-scaled energy leads to an upper bound on the average loss:

$$
\begin{aligned}
\mathcal{L}_{\mathrm{LS}}(\mathcal{D};\theta) &= \frac{1}{|\mathcal{D}|}\sum_{(\mathbf{x}_t,\mathbf{y}_t)\in\mathcal{D}} E_t^{\mathrm{LS}}(\mathbf{y}_t,\mathbf{x}_t) + \log\left(\sum_{\mathbf{y}\in\mathcal{Y}(\mathbf{x})}\exp(-E_t^{\mathrm{LS}}(\mathbf{y},\mathbf{x}_t))\right) \\
&= \frac{1}{|\mathcal{D}|}\sum_{(\mathbf{x}_t,\mathbf{y}_t)\in\mathcal{D}} \log\left(\sum_{\mathbf{y}\in\mathcal{Y}(\mathbf{x})}\exp(-l_t(\mathbf{y})(E(\mathbf{y},\mathbf{x}_t;\theta) - E(\mathbf{y}_t,\mathbf{x}_t;\theta)) + l_t(\mathbf{y}))\right) \\
&\geq \frac{1}{|\mathcal{D}|}\sum_{(\mathbf{x}_t,\mathbf{y}_t)\in\mathcal{D}} l_t(\hat{\mathbf{y}}(\mathbf{x}_t))(1 + E(\mathbf{y}_t,\mathbf{x}_t;\theta) - E(\hat{\mathbf{y}}(\mathbf{x}_t),\mathbf{x}_t;\theta)) \\
&\geq \frac{1}{|\mathcal{D}|}\sum_{(\mathbf{x}_t,\mathbf{y}_t)\in\mathcal{D}} l_t(\hat{\mathbf{y}}(\mathbf{x}_t)) \, .
\end{aligned}
$$

There is a connection with SSVM training objectives here as well. This loss-scaled CRF is the probabilistic equivalent of SSVM training with slack re-scaling [15].

Since both the loss-augmented and loss-scaled CRF objectives follow the general form of the maximum likelihood objective but with different energy functions, the form of the gradient is also that of Equation 3. The two key differences are that the energy function is now different, and the conditional expectation on $\mathbf{y}$ given $\mathbf{x}_t$ is according to the CRF distribution with the associated loss-sensitive energy. In general (particularly for the loss-scaled CRF), it will not be

possible to run belief propagation to compute the expectation[4], but adapted forms of loopy belief propagation or MCMC (e.g., Gibbs sampling) could be used.

3.2 Expected Loss

A second approach to deriving a smooth version of the average loss is to optimize the average *Expected Loss*, where the expectation is based on the CRF's distribution:

$$\mathcal{L}_{\mathrm{EL}}(\mathcal{D};\theta) = \frac{1}{|\mathcal{D}|} \sum_{(\mathbf{x}_t,\mathbf{y}_t)\in\mathcal{D}} \mathbb{E}_{\mathbf{y}|\mathbf{x}_t}\left[l_t(\mathbf{y})\right] = \frac{1}{|\mathcal{D}|} \sum_{(\mathbf{x}_t,\mathbf{y}_t)\in\mathcal{D}} \sum_{\mathbf{y}\in\mathcal{Y}(\mathbf{x}_t)} l_t(\mathbf{y})p(\mathbf{y}|\mathbf{x}_t) \,. \tag{6}$$

While this objective is not an upper bound, it becomes increasingly closer to the average loss as the entropy of $p(\mathbf{y}|\mathbf{x}_t)$ becomes smaller and puts all its mass on $\hat{\mathbf{y}}(\mathbf{x}_t)$.

The parameter gradient has the following form:

$$\frac{\partial \mathcal{L}_{\mathrm{EL}}(\mathcal{D};\theta)}{\partial \theta} = \frac{1}{|\mathcal{D}|} \sum_{(\mathbf{x}_t,\mathbf{y}_t)\in\mathcal{D}} \mathbb{E}_{\mathbf{y}|\mathbf{x}_t}\left[l_t(\mathbf{y})\right] \mathbb{E}_{\mathbf{y}|\mathbf{x}_t}\left[\frac{\partial E(\mathbf{y},\mathbf{x}_t;\theta)}{\partial \theta}\right] - \mathbb{E}_{\mathbf{y}|\mathbf{x}_t}\left[l_t(\mathbf{y})\frac{\partial E(\mathbf{y}_t,\mathbf{x}_t;\theta)}{\partial \theta}\right] \,. \tag{7}$$

If the required expectations cannot be computed tractably, MCMC sampling can be used to approximate them. Another alternative is to use a fixed set of representative samples [13].

3.3 Loss-inspired Kullback-Leibler

Both the average expected loss and the loss upper bound objectives have in common that their objectives are perfectly minimized when the posteriors $p(\mathbf{y}|\mathbf{x}_t)$ put all their mass on the targets $\mathbf{y}_t$. In practice, this is bound not to happen, since this is likely to correspond to an overfitted solution which will be avoided using additional regularization.

Instead of relying on a generic regularizer such as the ℓ^2-norm of the parameter vector, perhaps the loss function itself might provide cues as to how best to regularize the CRF. Indeed, we can think of the loss as a ranking of all potential predictions, from perfect to adequate to worse. Hence, if we are not to put all probability mass on $p(\mathbf{y}_t|\mathbf{x}_t)$, we could make use of the information provided by the loss in order to determine how to distribute the excess mass $1 - p(\mathbf{y}_t|\mathbf{x}_t)$ on other configurations of $\mathbf{y}$. In particular, it would be sensible to distribute it on other values of $\mathbf{y}$ in proportion to the loss $l_t(\mathbf{y})$.

To achieve this, we propose to first convert the loss into a distribution over the target $q(\mathbf{y}|t)$ and then minimize the *Kullback-Leibler* (KL) divergence between this target distribution and the CRF posterior:

$$\mathcal{L}_{\mathrm{KL}}(\mathcal{D};\theta) = \frac{1}{|\mathcal{D}|} \sum_{(\mathbf{x}_t,\mathbf{y}_t)\in\mathcal{D}} D_{\mathrm{KL}}(q(\cdot|t)\|p(\cdot|\mathbf{x}_t)) = -\frac{1}{|\mathcal{D}|} \sum_{(\mathbf{x}_t,\mathbf{y}_t)\in\mathcal{D}} \sum_{\mathbf{y}\in\mathcal{Y}(\mathbf{x}_t)} q(\mathbf{y}|t) \log p(\mathbf{y}|\mathbf{x}_t) - C \tag{8}$$

where constant $C = H(q(\cdot|t))$ is the entropy of the target distribution, which does not depend on parameter vector θ.

There are several ways of defining the target distribution $q(\mathbf{y}|t)$. In this work, we define it as follows:

$$q(\mathbf{y}|t) = \exp(-l_t(\mathbf{y})/T)/Z_t, \ \ Z_t = \sum_{\mathbf{y}\in\mathcal{Y}(x_t)} \exp(-l_t(\mathbf{y})/T) \tag{9}$$

where the temperature parameter T controls how peaked this distribution is around $\mathbf{y}_t$. The maximum likelihood objective is recovered as T approaches 0.

---

[4]In the loss-augmented case, one exception is if the loss decomposes into individual losses over each target variable y_i and the CRF follows a tree structure in its output. In this case, the loss terms can be integrated into the CRF unary features and belief propagation will perform exact inference.
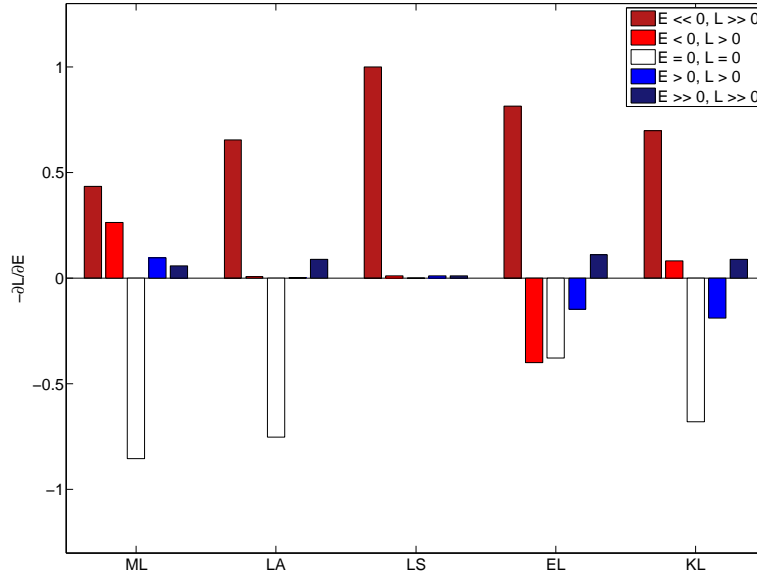
Figure 1: Negative derivatives of the objective with respect to energy for each of the five training objectives presented. The five training objectives ar: maximum likelihood (ML), loss-augmented ML (LA), loss-scaled ML (LS), expected loss (EL) and loss-inspired Kullback-Leibler divergence (KL). For each objective we consider five different configurations: from left to right, their energies are $[-1, -0.5, 0, 0.5, 1]$ and the losses are $[5, 1, 0, 1, 5]$. The middle one therefore corresponds to a ground-truth configuration; those to its left are currently more likely under the model, and loss increases with distance from this middle one. The derivatives for each objective are normalized by the ℓ^2 norm.

The gradient with respect to θ is simply the expectation of the gradient for maximum likelihood $\mathcal{L}_{\mathrm{ML}}$ according to the target distribution $q(\mathbf{y}|t)$. Here too, if the expectation is not tractable, one can using sampling to approximate it. In particular, since we have total control over the form of $q(\mathbf{y}|t)$, it is easy to define it such that it can be sampled from exactly.

3.4 Analyzing the Behavior of the Training Objectives

Figure 1 shows how the gradient with respect to the energy changes for each objective as we consider configurations $\mathbf{y}$ with varying energy and loss values. From this figure we see significant differences in the behaviors of the introduced objective functions. Only the expected-loss and loss-inspired Kullback-Leibler objectives will attempt to lower the energies of configurations that have non-zero loss. The maximum likelihood objective aims to raise the energies of the non-zero loss configurations, in proportion to how probable they are. On the other hand the loss-augmented and loss-scaled objectives concentrate on the most probable configurations that have the highest loss (worst violators), with the loss-scaled objective having the most extreme behavior and putting all the gradient on the worst violator. This behavior is expected as the energies get amplified by the addition (multiplication) of the loss which artificially raises the probability of the already probable violators.

The behavior of the expected-loss objective is counter-intuitive as it tries to lower the energy of all configurations that

have low loss, including those that are already more probable than the zero-loss one. In this example, it even pushes down more the energy of a non-zero loss configuration more than that of the zero-loss (target) configuration. The loss-inspired KL objective adjusts this and only lowers the energy of the zero-loss (ground-truth) and the low-loss configuration that has low probability.

4   Learning with Multiple Ground Truths

In certain applications, for some given input $\mathbf{x}_t$, there is not only a single target $\mathbf{y}_t$ that is correct (see Section 6 for the case of ranking). This information can easily be encoded within the loss function, by setting $l_t(\mathbf{y}) = 0$ for all such valid predictions.

In this context, maximum likelihood training corresponds to the objective:

$$\mathcal{L}_{\mathrm{ML}}(\mathcal{D}; \theta) = -\frac{1}{|\mathcal{D}|} \sum_{(\mathbf{x}_t, \mathbf{r}_t) \in \mathcal{D}} \sum_{\mathbf{y}_t \in \mathcal{Y}_0(\mathbf{x}_t)} \log p(\mathbf{y}_t | \mathbf{x}_t) \tag{10}$$

where $\mathcal{Y}_0(\mathbf{x}_t) = \{\mathbf{y} | \mathbf{y} \in \mathcal{Y}(\mathbf{x}_t),\ l_t(\mathbf{y}) = 0\}$. This is equivalent to maximizing the likelihood of all predictions $\mathbf{y}$ that are consistent with the loss, i.e. that have zero loss. The loss-augmented variant is similarly adjusted. As for loss-scaling, we replace the energy at the ground truth with the average energy of all valid ground truths in the loss-scaled energy:

$$E_t^{\mathrm{LS}}(\mathbf{y}, \mathbf{x}_t; \theta) = l_t(\mathbf{y}) \left(E(\mathbf{y}, \mathbf{x}_t; \theta) - \frac{1}{|\mathcal{Y}_t(\mathbf{x}_t)|} \sum_{\mathbf{y}_t \in \mathcal{Y}_0(\mathbf{x}_t)} E(\mathbf{y}_t, \mathbf{x}_t; \theta) \right) - l_t(\mathbf{y}) . \tag{11}$$

No changes to the average expected loss and loss-inspired KL objectives are necessary as they consider all valid $\mathbf{y}$.

In the setting of multiple ground truths, a clear distinction can be made between the average expected loss and the other objectives, in terms of the solutions they encourage. Indeed, the expected loss will be minimized as long as $\sum_{\mathbf{y}_t \in \mathcal{Y}_t(\mathbf{x}_t)} p(\mathbf{y}_t | \mathbf{x}_t) = 1$, i.e. probability mass is only put on configurations of $\mathbf{y}$ that have zero loss. On the other hand, the maximum likelihood and loss upper bound objectives add the requirement that the mass be equally distributed amongst those configurations. As for the loss-inspired KL, it requires that the sum of the probability mass sum to a constant smaller than 1, specifically $1 - \sum_{\mathbf{y}_t \in \mathcal{Y}_t(\mathbf{x}_t)} q(\mathbf{y}_t | t)$.

5   Related Work

While maximum likelihood is the dominant approach to training CRFs in the literature, others have proposed ways of adapting the CRF training objective for specific tasks. For sequence labeling problems, Kakade et al. [19] proposed to maximize the label-wise marginal likelihood instead of the joint label sequence likelihood, to reflect the fact that the task's loss function is the sum of label-wise classification errors. Suzuki et al. [20], Gross et al. [21] went a step further by proposing to directly optimize a smoothed version of the label-wise classification error (Suzuki et al. [20] also described how to apply it to optimize an F-score). Their approach is similar to the average expected loss described in Section 3.2, however they do not discuss how to generalize it to arbitrary loss functions. The average expected loss objective for CRFs was formulated by Taylor et al. [22] and Volkovs and Zemel [13], in the context of ranking.

Work in other frameworks than CRFs for structured output prediction have looked at how to incorporate loss information into learning. Tsochantaridis et al. [15] describe how to upper bound the average loss with margin and slack scaling. McAllester et al. [23] propose a perceptron-like algorithm based on an update which in expectation is close to the gradient on the true expected loss (i.e., the expectation is with respect to the true generative process). Both SSVMs and perceptron

algorithms require procedures for computing the so-called loss-adjusted MAP assignment of the output $\mathbf{y}$ which, for richly structured losses, can be intractable. One advantage of CRFs is that they can instead leverage the vast MCMC literature to sample from CRFs with loss-adjusted energies. Moreover, they open the door to alternative (i.e. not necessarily upper-bounding) objectives.

Finally, while Hazan and Urtasun [18] described how margin scaling can be applied to CRFs, we give for the first time the equivalent of slack scaling for CRFs in Section 3.1.

6 Experiments

We evaluate the usefulness of the different loss-sensitive training objectives on ranking tasks. In this setting, the input $\mathbf{x} = (\mathbf{q}, \mathbf{D})$ corresponds to a pair made of a query vector $\mathbf{q}$ and a set of documents $\mathbf{D} = \{\mathbf{d}^{(i)}\}$, and $\mathbf{y}$ is a vector corresponding to a ranking[5] of each document $\mathbf{d}^{(i)}$ among the whole set of documents $\mathbf{D}$.

Ranking is particularly interesting as a benchmark task for loss-sensitive training of CRFs for two reasons. The first is the complexity of the output space $\mathcal{Y}(\mathbf{q}, \mathbf{D})$, which corresponds to all possible permutations of documents $\mathbf{D}$, making the application of CRFs to this setting more challenging than sequential labeling problems with chain structure.

The second is that learning to rank is an example of a task with multiple ground truths (see Section 4), which is a more challenging setting than the single ground truth case. Indeed, for each training input $\mathbf{x}_t = (\mathbf{q}_t, \mathbf{D}_t)$, we are not given a single target rank $\mathbf{y}_t$, but a vector $\mathbf{r}_t$ of relevance level values for each document. The higher the level, the more relevant the document is and the better its rank should be. Moreover, two documents $\mathbf{d}_t^{(i)}$ and $\mathbf{d}_t^{(j)}$ with the same relevance level (i.e., $r_{ti} = r_{tj}$) are indistinguishable in their ranking, meaning that they can be swapped within some ranking without affecting the quality of that ranking.

The quality of a ranking is measured by the Normalized Discounted Cumulative Gain:

$$NDCG(\mathbf{y}, \mathbf{r}_t) = N_t \sum_{i=1}^{m_t} \frac{r_{ti} \log(2)}{\log(1 + y_i)} \tag{12}$$

where $N_t = 1/NDCG(\arg\operatorname{sort}(-\mathbf{r}_t), \mathbf{r}_t)$ is a normalization constant that insures the maximum value of NDCG is 1, which is achieved when documents are ordered in decreasing order of their relevance levels. Note that this is not a standard definition of NDCG, we use it here because this form was adopted for evaluation of the baselines on the Microsoft's LETOR4.0 datset collection [24]. To convert NDCG into a loss, we simply define $l_t(\mathbf{y}) = 1 - NDCG(\mathbf{y}, \mathbf{r}_t)$.

A common approach to ranking is to learn a scoring function $f(\mathbf{q}, \mathbf{d}^{(i)})$ which outputs for all documents $\mathbf{d}^{(i)} \in \mathbf{D}$ a score corresponding to how relevant document $\mathbf{d}^{(i)}$ is for query $\mathbf{q}$. Here, we follow the same approach by incorporating this scoring function into the energy function of the CRF. We use an energy function linear in the scores:

$$E(\mathbf{y}, \mathbf{q}, \mathbf{D}) = \sum_{i=1}^{|\mathbf{D}|} \alpha_{y_i} f(\mathbf{q}, \mathbf{d}^{(i)}) \tag{13}$$

where α is a weight vector of decreasing values (i.e., $\alpha_i > \alpha_j$ for $i < j$). In our experiments, we use a weighting inspired by the NDCG measure: $\alpha_i = \log(2)/\log(i + 1)$. Using this energy function, we can show that the prediction $\hat{\mathbf{y}}(\mathbf{q}, \mathbf{D})$ is obtained by sorting the documents in decreasing order of their scores:

$$\hat{\mathbf{y}}(\mathbf{q}, \mathbf{D}) = \underset{\mathbf{y} \in \mathcal{Y}(\mathbf{q}, \mathbf{D})}{\arg\min} E(\mathbf{y}, \mathbf{q}, \mathbf{D}) = \arg\operatorname{sort}([-f(\mathbf{q}, \mathbf{d}^{(1)}), \dots, -f(\mathbf{q}, \mathbf{d}^{(|\mathbf{D}|)})]) . \tag{14}$$

---

[5]For example, if $y_i = 3$, then document $\mathbf{d}^{(i)}$ is ranked third amongst all documents $\mathbf{D}$ for the query $\mathbf{q}$.
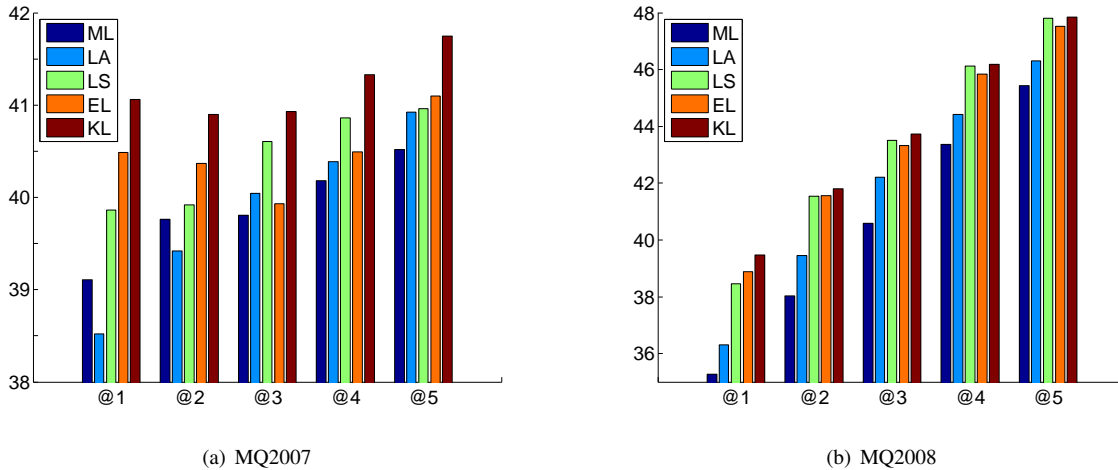
Figure 2: NDCG@1-5 results on MQ2008 and MQ2007 datasets for different learning objectives.

As for the scoring function, we use a simple linear function $f(\mathbf{q}, \mathbf{d}^{(i)}) = \theta^T \phi(\mathbf{q}, \mathbf{d}^{(i)})$ on a joint query-document feature representation $\phi(\mathbf{q}, \mathbf{d}^{(i)})$. A standard feature representation is provided in each ranking datasets we considered.

We trained CRFs according to maximum likelihood as well as the different loss-sensitive objectives described in Section 3. In all cases, stochastic gradient descent was used by iterating over queries and performing a gradient step update for each query. Moreover, because the size of $\mathcal{Y}(\mathbf{q}, \mathbf{D})$ is factorial in the number of documents, explicit summation over that set is only tractable for a small number of documents. To avoid this problem we use an approach similar to the one suggested by Petterson et al. [25]. Every time a query q_t is visited and its associated set of documents $\mathbf{D}_t$ is greater than 6, we randomly select a subset of 6 documents $\widetilde{\mathbf{D}} \subset \mathbf{D}_t$, ensuring that it contains at least one document of every relevance level found for that query. The exact parameter gradients can then be computed for this reduced set by enumerating all possible permutations, and the CRF can be updated.

6.1 Datasets

In our experiments we use the LETOR [24] benchmark datasets. These data sets were chosen because they are publicly available, include several baseline results, and provide evaluation tools to ensure accurate comparison between methods. In LETOR4.0 there are two learning to rank data sets MQ2007 and MQ2008. MQ2007 contains 1692 queries with 69623 documents and MQ2008 contains 784 queries and a total of 15211 documents. Each query document pair is assigned one of three relevance judgments: 2 = highly relevant, 1 = relevant and 0 = irrelevant. Both datasets come with five precomputed folds with 60/20/20 slits for training validation and testing. The results show for each model the averages of the test set results for the five folds.

6.2 Results

We experimented with five objective functions, namely: maximum likelihood (ML), loss-augmented ML (LA), loss-scaled ML (LS), expected loss (EL) and loss-inspired Kullback-Leibler divergence (KL). For the loss-augmented objective we introduced an additional weight $\alpha > 0$ modifying the energy to: $E_t(\mathbf{y}, \mathbf{x}_t; \theta) = E(\mathbf{y}, \mathbf{x}_t; \theta) - \alpha l_t(\mathbf{y})$. In this form α controls the contribution of the loss to the overall energy. For all objectives we did a sweep over learning rates in

Table 1: NDCG@1-5 results on MQ2008 and MQ2007 datasets.

	MQ2007: NDCG					MQ2008: NDCG				
	@1	@2	@3	@4	@5	@1	@2	@3	@4	@5
Regression	38.94	39.60	39.86	40.53	41.11	36.67	40.62	42.89	45.60	47.73
SVM-Struct	40.96	40.73	40.62	40.84	41.42	36.26	39.84	42.85	45.08	46.95
ListNet	40.02	40.63	40.91	**41.44**	41.70	37.54	41.12	43.24	45.68	47.47
AdaRank	38.76	39.67	40.44	40.67	41.02	38.26	**42.11**	**44.20**	**46.53**	**48.21**
KL	**41.06**	**40.90**	**40.93**	41.33	**41.75**	**39.47**	41.80	43.74	46.18	47.84

[0.5, 0.01, 0.01, 0.001]. Moreover we experimented with α in [1, 10, 20, 50] for the loss-augmented objective and T in [1, 10, 20, 50] for the KL objective. For each fold the setting that gave the best validation NDCG was chosen and the corresponding model was then tested on the test set.

The results for the five objective functions are shown in Figures 2(a) and 2(b). First, we see that in almost all cases loss-augmentation produces better results than the base maximum likelihood approach. Second, loss-scaling further improves on the loss-augmentation results and has similar performance to the expected objective. Finally, among all objectives, KL consistently produces the best results on both datasets. Taken together, these results strongly support our claim that incorporating the loss into the learning procedure of CRFs is important.

Comparisons of the CRFs trained on the KL objective with other models is also shown in Table 1, where the performance of linear regression and other linear baselines listed on LETOR's website is provided. We see that KL outperforms the baselines on the MQ2007 dataset on all truncations except 4. Moreover, on MQ2008 the performance KL is comparable to the best baseline AdaRank, with KL beating AdaRank on NDCG@1. We note also that KL consistently outperforms LETOR's SVM-Struct baseline.

7 Conclusion

In this work, we explored different approaches to incorporating loss function information into the training objective of a probabilistic CRF. We discussed how to adapt ideas from the SSVM literature to the probabilistic context of CRFs, introducing for the first time the equivalent of slack scaling to CRFs. We also described objectives that rely on the probabilistic nature of CRFs, including a novel loss-inspired KL objective. In an empirical comparison on ranking benchmarks, this new KL objective was shown to consistently outperform all other loss-sensitive objectives.

To our knowledge, this is the broadest comparison of loss-sensitive training objectives for probabilistic CRFs yet to be made. It strongly suggests that the most popular approach to CRF training, maximum likelihood, is likely to be suboptimal. While ranking was considered as the benchmark task here, in future work, we would like to extend our empirical analysis to other tasks such as labeling tasks.

References

[1] John Lafferty, Andrew McCallum, and Fernando Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *ICML*, pages 282–289. Morgan Kaufmann, San Francisco, CA, 2001.

[2] Fei Sha and Fernando Pereira. Shallow parsing with conditional random fields. In *HLT/NAACL*, 2003.

[3] Sunita Sarawagi and William W. Cohen. Semi-markov conditional random fields for information extraction. In Lawrence K. Saul, Yair Weiss, and Léon Bottou, editors, *NIPS*, pages 1185–1192. MIT Press, Cambridge, MA, 2005.

[4] Dan Roth and Wen-tau Yih. Integer linear programming inference for conditional random fields. In *Proceedings of the 22nd international conference on Machine learning*, ICML, pages 736–743, New York, NY, USA, 2005. ACM.

[5] Asela Gunawardana, Milind Mahajan, Alex Acero, and John C. Platt. Hidden conditional random fields for phone classification. In *Interspeech*, pages 1117–1120, 2005.

[6] Xuming He, Richard S. Zemel, and Miguel Á. Carreira-Perpiñán. Multiscale conditional random fields for image labeling. In *CVPR*, pages 695–702, 2004.

[7] Sanjiv Kumar and Martial Hebert. Discriminative fields for modeling spatial dependencies in natural images. In *NIPS*, 2003.

[8] Ariadna Quattoni, Michael Collins, , and Trevor Darrell. Conditional random fields for object recognition. In *NIPS*, 2005.

[9] Kengo Sato and Yasubumi Sakakibara. Rna secondary structural alignment with conditional random fields. *Bioinformatics*, 2005.

[10] Yan Liu, Jaime Carbonell, Peter Weigele, and Vanathi Gopalakrishnan. Protein fold recognition using segmentation conditional random fields (scrfs). *Journal of Computational Biology*, 2006.

[11] Trevor Cohn and Philip Blunsom. Semantic role labelling with tree conditional random fields. In *CoNLL*, pages 169–172, 2005.

[12] Zaiqing Nie, Ji rong Wen, Bo Zhang, and Wei ying Ma. 2d conditional random fields for web information extraction. In *ICML*, pages 1044–1051. ACM Press, 2005.

[13] Maksims Volkovs and Richard Zemel. BoltzRank: Learning to maximize expected ranking gain. In Léon Bottou and Michael Littman, editors, *ICML*, pages 1089–1096, Montreal, June 2009. Omnipress.

[14] Ben Taskar, Carlos Guestrin, and Daphne Koller. Max-margin markov networks. In Sebastian Thrun, Lawrence Saul, and Bernhard Schölkopf, editors, *NIPS*. MIT Press, Cambridge, MA, 2004.

[15] Ioannis Tsochantaridis, Thorsten Joachims, Thomas Hofmann, and Yasemin Altun. Large margin methods for structured and interdependent output variables. *JMLR*, 6:1453–1484, 2005. ISSN 1532-4435.

[16] Julian Besag. Statistical analysis of non-lattice data. *The Statistician*, 24(3):179–195, 1975.

[17] C. Sutton and A. Mccallum. Piecewise Training for Undirected Models. In *UAI*, 2005.

[18] Tamir Hazan and Raquel Urtasun. A primal-dual message-passing algorithm for approximated large scale structured prediction. In J. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R.S. Zemel, and A. Culotta, editors, *NIPS*, pages 838–846. MIT Press, 2010.

[19] Sham Kakade, Yee Whye Teh, and Sam T. Roweis. An alternate objective function for markovian fields. In *ICML*, pages 275–282, 2002.

[20] Jun Suzuki, Erik McDermott, and Hideki Isozaki. Training conditional random fields with multivariate evaluation measures. In *ICCL-ACL*, pages 217–224, Stroudsburg, PA, USA, 2006. Association for Computational Linguistics.

[21] Samuel S. Gross, Olga Russakovsky, Chuong B. Do, and Serafim Batzoglou. Training conditional random fields for maximum labelwise accuracy. In B. Schölkopf, J. Platt, and T. Hoffman, editors, *NIPS*, pages 529–536. MIT Press, Cambridge, MA, 2007.

[22] Michael J. Taylor, John Guiver, Stephen Robertson, and Tom Minka. SoftRank: optimizing non-smooth rank metrics. In *WSDM*, pages 77–86, 2008.

[23] David McAllester, Tamir Hazan, and Joseph Keshet. Direct loss minimization for structured prediction. In J. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R.S. Zemel, and A. Culotta, editors, *NIPS*, pages 1594–1602. MIT Press, 2010.

[24] T. Liu, J. Xu, W. Xiong, and H. Li. LETOR: Benchmark dataset for search on learning to rank for information retrieval. In *ACM SIGIR Workshop on Learning to Rank for Information Retrieval*, 2007.

[25] J. Petterson, T. S. Caetano, J. J. McAuley, and J. Yu. Exponential family graph matching and ranking. In *Neural Information Processing Systems*, pages 1455–1463, 2009.