
Training Deep Neural Networks via Direct Loss Minimization

Yang Song

Dept. of Physics, Tsinghua University, Beijing 100084, China

SONGYANG12@MAILS.TSINGHUA.EDU.CN

Alexander G. Schwing

ASCHWING@CS.TORONTO.EDU

Richard S. Zemel

ZEMEL@CS.TORONTO.EDU

Raquel Urtasun

URTASUN@CS.TORONTO.EDU

Dept. of Computer Science, University of Toronto, Toronto, Ontario M5S 2E4, Canada

Abstract

Supervised training of deep neural nets typically relies on minimizing cross-entropy. However, in many domains, we are interested in performing well on metrics specific to the application. In this paper we propose a direct loss minimization approach to train deep neural networks, which provably minimizes the application-specific loss function. This is often non-trivial, since these functions are neither smooth nor decomposable and thus are not amenable to optimization with standard gradient-based methods. We demonstrate the effectiveness of our approach in the context of maximizing average precision for ranking problems. Towards this goal, we develop a novel dynamic programming algorithm that can efficiently compute the weight updates. Our approach proves superior to a variety of baselines in the context of action classification and object detection, especially in the presence of label noise.

1. Introduction

Standard supervised neural network training involves computing the gradient of the loss function with respect to the parameters of the model, and therefore requires the loss function to be differentiable. Many interesting loss functions are, however, *non-differentiable* with respect to the output of the network. Notable examples are functions based on discrete outputs, as is common in labeling and ranking problems. In many cases these losses are also *non-decomposable*, in that they cannot be expressed as simple sums over the output units of the network.

In the context of structured prediction problems, in which

Proceedings of the 33rd International Conference on Machine Learning, New York, NY, USA, 2016. JMLR: W&CP volume 48. Copyright 2016 by the author(s).

the output is multi-dimensional, researchers have developed max-margin training methods that are capable of minimizing an upper bound on non-decomposable loss functions. Standard learning in this paradigm involves changing the parameters such that the model assigns a higher score to the groundtruth output than to any other output. This is typically encoded by a constraint, enforcing that the groundtruth score should be higher than that of a selected, contrastive output. The latter is defined as the result of inference performed using a modified score function which combines the model score and the task loss, representing the metric that we care about for the application. This modified scoring function encodes the fact that we should penalize higher scoring configurations that are inferior in terms of the task loss. Various efficient methods have been proposed for incorporating complex discrete loss functions into this max-margin approach (Yue et al., 2007; Volkovs & Zemel, 2009; Tarlow & Zemel, 2012; Mohapatra et al., 2014). Importantly, however, this form of learning does not directly optimize the task loss, but rather an upper bound.

An alternative approach, frequently used in deep neural networks, is to train with a surrogate loss that can be easily optimized, *e.g.*, cross-entropy (LeCun & Huang, 2005; Bengio et al., 2015). The problem of this procedure is that for many application domains the task loss differs significantly from the surrogate loss.

The seminal work of McAllester et al. (2010) showed how to compute the gradient of complex non-differentiable loss functions when dealing with linear models. In this paper we extend their theorem to the non-linear case. This is important in practice as it provides us with a new learning algorithm to train deep neural networks end-to-end to minimize the application specific loss function. As shown in our experiments on action classification and object detection, this is very beneficial, particularly when dealing with noisy labels.

2. Direct Loss Minimization for Neural Networks

In this section we present a novel formulation for learning neural networks by minimizing the task loss. Towards this goal, our first main result is a theorem extending the direct loss minimization framework of (McAllester et al., 2010) to non-linear models.

A neural network can be viewed as defining a composite scoring function $F(x, y, w)$, which depends on the input data $x \in \mathcal{X}$, some parameters $w \in \mathbb{R}^A$, and the output $y \in \mathcal{Y}$. Inference is then performed by picking the output with maximal score, *i.e.*:

$$y_w = \arg \max_{\hat{y} \in \mathcal{Y}} F(x, \hat{y}, w).$$

Given a dataset of input-output pairs $\mathcal{D} = \{(x, y)\}$, a standard machine learning approach is to optimize the parameters w of the scoring function F by optimizing cross-entropy. This is equivalent to maximizing the likelihood of the data, where the probability over each output configuration is given by the output of a softmax function, attached to the last layer of the network.

However, in many practical applications, we want prediction to succeed in an application-specific metric. This metric is typically referred to as the *task loss*, $L(y, y_w) \geq 0$, which measures the compatibility between the annotated configuration y and the prediction y_w . For learning, in this paper, we are thus interested in minimizing the task loss

$$w^* = \arg \min_w \mathbb{E} [L(y, y_w)], \quad (1)$$

where $\mathbb{E}[\cdot]$ denotes an expectation taken over the underlying distribution behind the given dataset.

Solving this program is non-trivial, as many loss functions of interest are non-decomposable and non-smooth, and thus are not amenable to gradient-based methods. Examples of such metrics include average precision (AP) from information retrieval, intersection-over-union which is used in image labeling, and normalized discounted cumulative gain (NDCG) which is popular in ranking. In general many metrics are discrete, are not simple sums over the network outputs, and are not readily differentiable.

During training of a classifier it is hence common to employ a surrogate error metric, such as cross-entropy or hinge-loss, where one can directly compute the gradients with respect to the parameters. In the context of structured prediction models, several approaches have been developed to effectively optimize the structured hinge loss with non-decomposable task losses. While these methods include the task loss in the objective, they are not directly minimizing

it, and hence these surrogate losses are at best highly correlated with the desired metric. Finding efficient techniques to directly minimize the metric of choice is therefore desirable.

McAllester et al. (2010) showed that it is possible to asymptotically optimize the task-loss when the function F is linear in the parameters *i.e.*, $F(x, y, w) = w^\top \phi(x, y)$. This work has produced encouraging results. For example, McAllester et al. (2010) used this technique for phoneme-to-speech alignment on the TIMIT dataset optimizing the τ -alignment loss and the τ -insensitive loss, while Keshet et al. (2011) illustrated applicability of the method to hidden Markov models for speech. Direct loss minimization was also shown to work well for inverse optimal control by Doerr et al. (2015).

The first contribution of our work is to generalize this theorem to arbitrary scoring functions, *i.e.*, non-linear and non-convex functions. This allows us to derive a new training algorithm for deep neural networks which directly minimizes the task loss.

Theorem 1 (General Loss Gradient Theorem). *When given a finite set \mathcal{Y} , a scoring function $F(x, y, w)$, a data distribution, as well as a task-loss $L(y, \hat{y})$, then, under some mild regularity conditions (see the supplementary material for details), the direct loss gradient has the following form:*

$$\begin{aligned} & \nabla_w \mathbb{E} [L(y, y_w)] \\ &= \pm \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} \mathbb{E} [\nabla_w F(x, y_{\text{direct}}, w) - \nabla_w F(x, y_w, w)], \end{aligned} \quad (2)$$

with

$$\begin{aligned} y_w &= \arg \max_{\hat{y} \in \mathcal{Y}} F(x, \hat{y}, w), \\ y_{\text{direct}} &= \arg \max_{\hat{y} \in \mathcal{Y}} F(x, \hat{y}, w) \pm \epsilon L(y, \hat{y}). \end{aligned} \quad (3)$$

Proof. We refer the reader to the supplementary material for a formal proof of the theorem. \square

According to Theorem 1, to obtain the gradient we need to find the solution of two inference problems. The first computes y_w , which is a standard inference task, solved by employing the forward propagation algorithm. The second inference problem is prediction using a scoring function which is perturbed by the task loss $L(y, \hat{y})$. This is typically non-trivial to solve, particularly when the task loss is not decomposable. We borrow terminology from the structured prediction literature, where this perturbed inference problem is commonly referred to as *loss-augmented inference* (Tsochantaridis et al., 2005; Chen et al., 2015). In the

Algorithm: Direct Loss Minimization for Deep Networks

Repeat until stopping criteria

1. Forward pass to compute $F(x, \hat{y}; w)$
2. Obtain y_w and y_{direct} via inference and loss-augmented inference
3. Single backward pass via chain rule to obtain gradient $\nabla_w \mathbb{E}[L(y, y_w)]$
4. Update parameters using stepsize η : $w \leftarrow w - \eta \nabla_w \mathbb{E}[L(y, y_w)]$

Figure 1. Our algorithm for direct loss minimization.

following section we derive an efficient dynamic programming algorithm to perform loss-augmented inference when the task loss is average precision.

Note that loss-augmented inference, as specified in Eq. (3), can take the task loss into account in a positive or a negative way. Depending on the sign, the gradient direction changes. McAllester et al. (2010) provides a nice intuition for the two different directions. The positive update performs a step away from a worse configuration, while the negative direct loss gradient encourages moves towards better outputs. This can be seen when considering that maximization of the loss returns the worst output configuration, while maximization of its negation returns the label with the lowest loss. It is however an empirical question, whether the positive update or the negative version performs better. McAllester et al. (2010) reports better results with the negative update, while we find in our experiments that the positive update is better when applied to deep neural networks. We will provide intuition for this phenomenon in Sec. 4.

Note the relation between direct loss minimization and optimization of the structured hinge-loss. While we compute the gradient via the difference between the loss-augmented inference result and the prediction, structured hinge-loss requires computation of the difference between the loss-augmented inference solution and the ground truth. In addition, for any finite ϵ , direct loss minimization is sub-gradient descent of some corresponding ramp loss (Keshet & McAllester, 2011).

In Fig. 1 we summarize the resulting learning algorithm, which consists of the following four steps. First we use a standard forward pass to evaluate F . We then perform inference and loss-augmented inference as specified in Eq. (3) to obtain the prediction y_w and y_{direct} . We combine the predictions to obtain the gradient $\nabla_w \mathbb{E}[L(y, y_w)]$ via a single backward pass which is then used to update the parameters. For notational simplicity we omit details like momentum-based gradient updates, the use of mini-batches, etc., which are easily included.

3. Direct Loss Minimization for Average Precision

In order to directly optimize the task-loss we are required to compute the gradient defined in Eq. (2). As mentioned above, we need to solve both the standard inference task as well as the loss-augmented inference problem given in Eq. (3). While the former is typically assumed to be solvable, the latter depends on L and might be very complex to solve, e.g., when the loss is not decomposable.

In this paper we consider ranking problems, where the desired task loss L is average precision (AP), a concrete example of a non-decomposable and non-smooth target loss function. For the linear setting, efficient algorithms for positive loss-augmented inference with AP loss were proposed by Yue et al. (2007) and Mohapatra et al. (2014). Their results can be extended to the non-linear setting only in the positive case, where $y_{\text{direct}} = \arg \max_{\hat{y} \in \mathcal{Y}} F(x, \hat{y}, w) + \epsilon L(y, \hat{y})$. For the negative setting inequalities required in their proof do not hold and thus their method is not applicable. In this section we propose a more general algorithm that can handle both cases with the same time complexity as (Yue et al., 2007), while being more intuitive to prove and understand.

Alternatives to optimizing average precision are methods such as RankNet (Burges et al., 2005), LambdaRank (Burges et al., 2007) and LambdaMART (Burges, 2010). For an overview, we refer the reader to Burges (2010) and references therein. Our goal here is simply to show direct loss minimization of AP as an example of our general framework.

To compute the AP loss we are given a set of positive, i.e., relevant, and negative samples. The sample x_i belongs to the positive class if $i \in \mathcal{P} = \{1, \dots, |\mathcal{P}|\}$, and x_i is part of the negative class if $i \in \mathcal{N} = \{|\mathcal{P}| + 1, \dots, |\mathcal{P}| + |\mathcal{N}|\}$.

We define the output to be composed of pairwise comparisons, with $y_{i,j} = 1$ if sample i is ranked higher than sample j , $y_{i,i} = 0$, and $y_{i,j} = -1$ otherwise. We subsumed all these pairwise comparisons in $y = (\dots, y_{i,j}, \dots)$. Similarly $x = (x_1, \dots, x_N)$ contains all inputs, and $N =$

Algorithm: AP loss-augmented inference

1. Set $h(1, 0) = \mp \epsilon \frac{1}{|\mathcal{P}|}$ and $h(0, 1) = 0$
2. For $i = 1, \dots, |\mathcal{P}|, j = 1, \dots, |\mathcal{N}|$, recursively fill the matrix

$$h(i, j) = \max \begin{cases} h(i-1, j) \mp \epsilon \frac{1}{|\mathcal{P}|} \frac{i}{i+j} + B(i, j), & i \geq 1, j \geq 0 \\ h(i, j-1) + G(i, j), & i \geq 0, j \geq 1 \end{cases}$$

3. Backtrack to obtain configuration \hat{y}^*

Figure 2. Our algorithm for AP loss-augmented maximization or minimization.

$|\mathcal{P}| + |\mathcal{N}|$ refers to the total number of data points in the training set. In addition, we assume the ranking across all samples y to be complete, *i.e.*, consistent. During inference we obtain a ranking by predicting scores $\phi(x_i, w)$ for all data samples x_i which are easily sorted afterwards.

For learning we generalize the feature function defined by Yue et al. (2007) and Mohapatra et al. (2014) to a non-linear scoring function, using

$$F(x, y, w) = \frac{1}{|\mathcal{P}||\mathcal{N}|} \sum_{i \in \mathcal{P}, j \in \mathcal{N}} y_{i,j} (\phi(x_i, w) - \phi(x_j, w)).$$

where $\phi(x_i, w)$ is the output of the deep neural network when using the i -th example as input.

AP is unfortunately a non-decomposable loss function, *i.e.*, it does not decompose into functions dependent only on the individual $y_{i,j}$. To define the non-decomposable AP loss formally, we construct a vector $\hat{p} = \text{rank}(\hat{y}) \in \{0, 1\}^{|\mathcal{P}|+|\mathcal{N}|}$ by sorting the data points according to the ranking defined by the configuration \hat{y} . This vector contains a 1 for each positive sample and a value 0 for each negative element. In applications such as object detection, an example is said to be positive if the intersection over union of its bounding box and the ground truth box is bigger than a certain threshold (typically 50%). Using the rank operator we obtain the AP loss by comparing two vectors $p = \text{rank}(y)$ and $\hat{p} = \text{rank}(\hat{y})$ via

$$L_{AP}(p, \hat{p}) = 1 - \frac{1}{|\mathcal{P}|} \sum_{j: \hat{p}_j=1} \text{Prec}@j, \quad (4)$$

where $\text{Prec}@j$ is the percentage of relevant samples in the prediction \hat{p} that are ranked above position j .

To solve the loss-augmented inference task we have to solve the following program

$$\arg \max_{\hat{y}} F(x, \hat{y}, w) \pm \epsilon L_{AP}(\text{rank}(y), \text{rank}(\hat{y})). \quad (5)$$

In the following we derive a dynamic programming algorithm that can handle both the positive and negative case

and has the same complexity as (Yue et al., 2007). Towards this goal, we first note that Observation 1 of Yue et al. (2007) holds for both the positive and the negative case. For completeness, we repeat their observation here and adapt it to our notation.

Observation 1 (Yue et al. (2007)). *Consider rankings which are constrained by fixing the relevance at each position in the ranking (e.g., the 3rd sample in the ranking must be relevant). Every ranking satisfying the same set of constraints will have the same L_{AP} . If the positive samples are sorted by their scores in descending order, and the irrelevant samples are likewise sorted by their scores, then the interleaving of the two sorted lists satisfying the constraints will maximize Eq. (5) for that constrained set of rankings.*

Observation 1 means that we only need to consider the interleaving of two sorted lists of \mathcal{P} and \mathcal{N} to solve the program given in Eq. (5). From now on we therefore assume that the elements of \mathcal{P} and \mathcal{N} are sorted in descending order of their predicted score.

We next assert the optimal substructure property of our problem. Let the restriction to subsets of i positive and j negative examples be given by $\mathcal{P}_i = \{1, \dots, i\}$ and $\mathcal{N}_j = \{|\mathcal{P}| + 1, \dots, |\mathcal{P}| + j\}$. The cost function value obtained when restricting loss-augmented inference to the subsets can be computed as:

$$\begin{aligned} h(i, j) &= \max_{\hat{y}} \frac{1}{|\mathcal{P}||\mathcal{N}|} \sum_{m \in \mathcal{P}_i} \sum_{n \in \mathcal{N}_j} \hat{y}_{m,n} (\phi(x_m, w) - \phi(x_n, w)) \\ &\quad \pm \epsilon L_{AP}^{i,j}(\text{rank}(y), \text{rank}(\hat{y})), \end{aligned} \quad (6)$$

where $L_{AP}^{i,j}$ refers to the AP loss restricted to subsets of i positive and j negative elements.

Lemma 1. *Suppose that $\text{rank}(\hat{y}^*)$ is the optimal ranking for Eq. (6) when restricted to i positive and j negative samples. Any of its sub-sequences starting at position 1 is then*

also an optimal ranking for the corresponding restricted sub-problem.

We provide the proof of this lemma in the supplementary material. Based on Lemma 1 we can construct Bellman equations to recursively fill in a matrix of size $\mathcal{P} \times \mathcal{N}$, resulting in an overall time complexity of $O(|\mathcal{P}||\mathcal{N}|)$. We can then obtain the optimal loss-augmented predicted ranking via back-tracking.

The Bellman recursion computes the optimal cost function value of the sub-sequence containing data from \mathcal{P}_i and \mathcal{N}_j , based on previously computed values as follows:

$$h(i, j) = \max \begin{cases} h(i-1, j) \mp \epsilon \frac{1}{|\mathcal{P}|} \frac{i}{i+j} + B(i, j) & i \geq 1, j \geq 0 \\ h(i, j-1) + G(i, j) & i \geq 0, j \geq 1 \end{cases}$$

with initial conditions $h(1, 0) = \mp \epsilon \frac{1}{|\mathcal{P}|}$ and $h(0, 1) = 0$. Note that we used the pre-computed matrices of scores

$$B(i, j) = -\frac{1}{|\mathcal{P}||\mathcal{N}|} \sum_{k \in \mathcal{N}_j} (\phi(x_i, w) - \phi(x_k, w))$$

$$G(i, j) = \frac{1}{|\mathcal{P}||\mathcal{N}|} \sum_{k \in \mathcal{P}_i} (\phi(x_k, w) - \phi(x_j, w)).$$

Intuitively, the Bellman recursion considers two cases: (i) how does the maximum score $h(i, j)$ for the solution restricted to sequences \mathcal{P}_i and \mathcal{N}_j change if we add a new positive sample; (ii) how does the maximum score $h(i, j)$ change when adding a new negative sample. In both cases we need to add the respective scores $B(i, j)$ or $G(i, j)$. When adding a positive sample we additionally need to consider a contribution from the precision which contains i positive elements from a total of $i + j$ elements.

The matrices $B(i, j)$ and $G(i, j)$ store the additional contribution to Eq. (6) obtained when adding a positive or a negative sample respectively. They can be efficiently computed ahead of time using the recursion

$$B(i, j) = B(i, j-1) - \frac{1}{|\mathcal{P}||\mathcal{N}|} (\phi(x_i, w) - \phi(x_j, w))$$

$$G(i, j) = G(i-1, j) + \frac{1}{|\mathcal{P}||\mathcal{N}|} (\phi(x_i, w) - \phi(x_j, w)),$$

with initial conditions $B(i, 0) = 0, G(0, j) = 0$.

We summarize our dynamic programming algorithm for AP loss-augmented inference in Fig. 2. Note that after having completed the matrix $h(i, j)$ we can backtrack to obtain the best loss-augmented ranking. We hence presented a general solution to solve positive and negative loss-augmented inference defined in Eq. (3) for the AP loss given in Eq. (4).

4. Experimental Evaluation

To evaluate the performance of our approach we perform experiments on both synthetic and real datasets. We compare the positive and negative version of our direct loss minimization approach to a diverse set of baselines.

4.1. Synthetic data

Dataset: We generate synthetic data via a neural network which assigns scalar scores to input vectors. The neural network consists of four layers of rectified linear units with parameters randomly drawn from independent Gaussians with zero mean and unit variance. The input for a training example is drawn from a 10 dimensional Gaussian, and the output produced by the network is its score. We generate 20,000 examples, and sort them in descending order based on their scores. The top 20% of the samples are assigned to the positive set \mathcal{P} and the remaining examples to the negative set. We then randomly divide the generated data into a training set containing 10,000 elements and a test set containing the rest. We compare various loss functions in terms of their ability to train the network to effectively reproduce the original scoring function. To ensure that we do not suffer from model mis-specification we employ the same network structure when training the parameters from random initializations. This synthetic experiment provides a good test environment to compare these training methods, as the mapping from input to scores is fairly complex yet deterministic.

Algorithms: We evaluate the positive and negative versions of our direct loss minimization when using two different task losses: AP and 0-1 loss. For the latter we predict for each sample x_i whether it is a member of the set \mathcal{P} or whether it is part of the set \mathcal{N} . We named these algorithms, “pos-AP,” “neg-AP,” “pos-01,” and “neg-01.” We also evaluate training the network using hinge loss, when employing AP and 0-1 loss as the task losses. We called these baselines “hinge-AP” and “hinge-01.” Note that “hinge-AP” is equivalent to the approach of Yue et al. (2007). We use the perceptron updates as additional baselines, which we call “per-AP” and “per-01.” Finally, the last baseline uses maximum-likelihood (*i.e.*, cross entropy) to train the network. We call this approach “x-ent.” The parameters of all algorithms are individually determined via grid search to produce the best AP on the training set.

Results: Fig. 3 shows AP results on the test set. We observe that the perceptron update does not perform well, since it does not take the task loss into account and has convergence issues. Contrary to the claim in (McAllester et al., 2010) for linear models, the negative update for direct loss is not competitive in our setting (*i.e.*, non-linear models). This might be due to the fact that it tends to

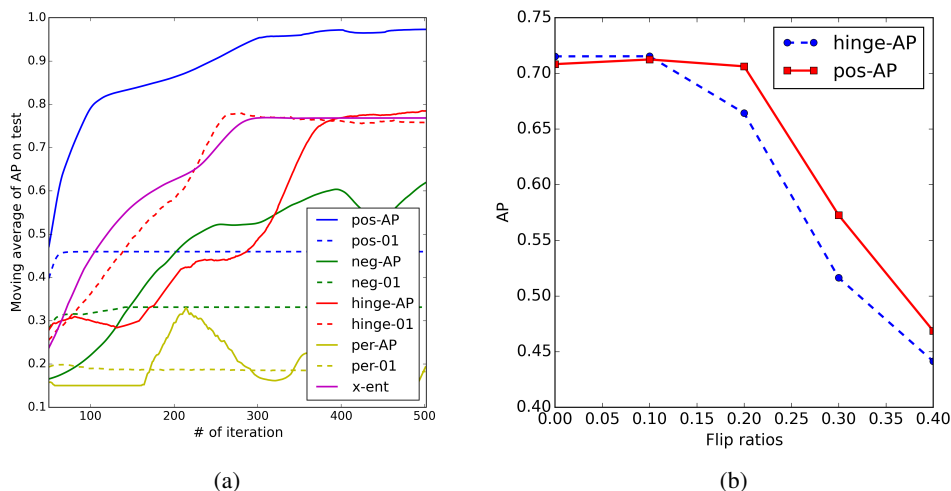


Figure 3. Experiments on synthetic data: (a) Average Precision (AP) on the test set as a function of the number of iterations (best view in color). (b) The robustness of pos-AP compared to hinge-AP.

overly correct the classifier in noisy situations. Note the resemblance of the perceptron method and the negative direct loss minimization update: they are both trying to move “towards better.” Therefore we expect the negative update to behave similarly to the perceptron, *i.e.*, more vulnerable to noisy data. This resemblance also explains their strong performance on the TIMIT dataset (Cheng et al., 2009; McAllester et al., 2010), which is relatively easy and can be handled well with linear models. For the same reason, we expect the negative update to lead to better performance in less noisy situations, *e.g.*, if the data is nearly linearly separable. To further examine this hypothesis we provide additional experimental results in the supplementary material.

The direct loss minimization of 0-1 loss does not work well. It is likely that the sharp changes of the 0-1 loss result in a ragged energy landscape that is hard to optimize, while smoothing as performed by AP loss or surrogate costs helps in this setting.

The hinge-AP algorithm of Yue et al. (2007) performs slightly better than other algorithms based on 0-1 loss. This shows that taking the AP loss into consideration helps improve the performance measured by AP. It is also important to note that when employing positive updates our direct loss minimization outperforms all baselines by a large margin. Note the resemblance between the structural SVM update and positive direct loss minimization: they are both moving “away from worse.” We believe that hinge loss does not deal with label noise well as the update depends on the ground truth more than the direct loss update. Another possible reason is that the hinge loss upper bound is

looser in this case. Taking the 0-1 loss as an example, the hinge loss of each outlier is much larger than the cost measured by the 0-1 loss (which is at most 1) in noisy case. This illustrates why the gap between hinge loss and target loss is large in noisy situations in general.

To further examine the performance of pos-AP and hinge-AP, we performed another small scale synthetic experiment. We hypothesize that when the data contains outliers, and more generally the labels are noisy, hinge loss will become a worse approximation. We control the noise level in the data as a way of testing this hypothesis. We randomly generate 1000 10-dimensional data from $\mathcal{N}(0, 10)$. Datum sample x is assigned to be positive when $\|x\|_2^2 > 1200$ and negative when $\|x\|_2^2 < 1000$. The noise is incorporated by randomly flipping a fixed percentage of labels. We use the same neural network structure for this task, tune the parameters on the training set, and report their results on the independent test set. As shown in Fig. 3(b) our method, pos-AP, is more robust to noise.

Based on the results obtained from the synthetic experiments, during the experimental evaluation on real datasets we focus on comparing the positive non-linear direct loss minimization (pos-AP) to the strong baseline of hinge-AP Yue et al. (2007), as well as the standard approach of training based on cross-entropy.

4.2. Action classification task

Dataset: In the next experiment we use the PASCAL VOC2012 action classification dataset provided by Everingham et al. (2014). The dataset contains 4588 images and 6278 “trainval” person bounding boxes. For each of the

noise level	0			10%			20%			30%			40%		
	method	x-ent	hinge-AP	pos-AP	x-ent	hinge-AP	pos-AP	x-ent	hinge-AP	pos-AP	x-ent	hinge-AP	pos-AP	x-ent	hinge-AP
jumping	76.6	77.0	77.3	68.7	64.0	74.0	51.6	44.9	65.1	42.3	42.4	36.3	22.8	27.2	52.0
phoning	45.4	46.2	44.8	36.3	31.2	39.2	25.7	22.1	35.4	11.8	10.8	15.4	9.5	10.2	16.8
playing instrument	72.6	74.0	74.0	67.9	67.6	71.4	60.6	57.3	69.9	38.3	40.0	62.8	25.1	15.9	60.6
reading	50.3	50.6	50.6	40.1	36.8	43.4	27.0	22.1	40.1	17.9	15.9	17.3	13.8	15.8	9.6
riding bike	92.3	92.9	93.1	84.9	88.2	90.1	72.9	73.3	88.2	54.1	40.6	79.3	32.9	18.3	22.3
riding horse	89.1	92.0	92.4	78.4	82.8	84.9	70.2	77.8	79.4	45.7	48.0	69.7	25.1	28.4	53.2
running	82.3	84.0	84.8	77.9	76.9	76.2	64.6	57.3	75.8	40.5	44.0	71.3	17.7	9.3	29.8
taking photo	41.2	45.8	43.2	33.0	40.7	33.0	19.0	21.7	22.4	15.9	12.1	19.5	11.0	11.3	8.4
using computer	68.7	68.3	66.5	57.6	59.3	60.2	42.3	45.1	56.0	21.7	21.5	41.4	13.2	15.1	11.1
walking	61.3	68.0	68.2	51.9	53.3	52.2	39.0	35.1	46.1	29.1	24.6	46.3	11.2	16.2	32.6
mean	68.0	69.9	69.5	59.7	60.1	62.5	47.3	45.7	57.8	31.7	30.0	45.9	18.2	16.8	29.6

Table 1. Comparison of our direct AP loss minimization approach to two strong baselines, which utilize surrogate loss functions, on the action classification task. Each method is evaluated for various amounts of label noise in the test dataset.

10 target classes, we divide the training dataset into equal-sized training, validation and test sets. We tuned the learning rate, regularization weight, and ϵ for all the algorithms based on their performance on the validation dataset, and report the results on the test set. For all algorithms we used the entire available training set in a single batch and performed 300 iterations.

Algorithms: We train our non-linear direct loss minimization as well as all the baselines individually for each class. As baselines we again use a deep network trained with cross entropy and also consider the structured SVM method proposed by Yue et al. (2007). The deep network used in these experiments follows the architecture of Krizhevsky et al. (2012), with the top dimension adjusted to a single output. We initialize the parameters using the weights trained on ILSVRC2012 (Russakovsky et al., 2015). Inspired by the RCNN (Girshick et al., 2014), we cropped the regions of each image with a padding of 16 pixels and interpolated them to a size of $227 \times 227 \times 3$ to fit the input data dimension of the network. All the algorithms we compare to as well as our approach use raw pixels as input.

Results: Intuitively we expect direct loss minimization to outperform surrogate loss functions whenever there is a significant number of outliers in the data. To evaluate this hypothesis, we conduct experiments by randomly flipping a fixed number of labels. Our experiments shown in Tab. 1 and Fig. 4 confirm our intuitions, and direct loss works much better than hinge loss in the presence of label noise. When there is no noise, both algorithms perform similarly.

4.3. Object detection task

Dataset: For object detection we use the PASCAL VOC2012 object detection dataset collected by Everingham et al. (2014). The dataset contains 5717 images for




Figure 4. This figure summarizes the effect of label noise in the action classification task, by showing the mean AP on the test set for different proportions of flipped labels. We compare the performance of the strongest baseline method, the hinge-loss trained network which uses AP as the task loss, versus our direct loss method, pos-AP.

training, 5823 images for validation and 10991 images for test. For each image, we use the fast mode of selective search by Uijlings et al. (2013) to produce around 2000 bounding boxes. We train algorithms on the training set and report results on the validation set.

Algorithms: On this dataset we follow the RCNN paradigm (Girshick et al., 2014). We adjust the dimension of the top layer of the network (Krizhevsky et al., 2012) to be one and fine-tune using weights pre-trained on ILSVRC2012 (Russakovsky et al., 2015). We train direct loss minimization for all 20 classes separately. In contrast to the action classification task, we cannot calculate the overall AP in each iteration, due to the large number of bounding boxes. Instead, we use the AP on each mini-batch to approximate the overall AP. We find that using a batch

	aeroplane	bicycle	bird	boat	bottle	bus	car	cat	chair	cow	diningtable	dog	horse	motorbike	person	pottedplant	sheep	sofa	train	tvmonitor	mean
x-ent (0 label noise)	63.8	61.0	42.6	30.7	23.5	63.2	51.7	58.5	20.1	37.0	32.0	52.8	50.8	62.5	50.1	23.5	48.3	33.1	48.5	57.4	45.6
hinge-AP (0 label noise)	67.5	60.6	43.6	30.8	25.3	64.5	54.9	64.4	21.9	34.5	34.2	57.0	48.8	63.9	56.3	25.1	49.6	37.4	54.3	57.3	47.6
pos-AP (0 label noise)	65.1	59.8	43.7	31.4	27.7	64.6	53.1	63.7	25.6	40.2	36.2	58.1	52.8	63.6	56.2	28.1	50.0	38.9	50.0	61.3	48.5
hinge-AP (20% label noise)	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
pos-AP (20% label noise)	52.8	54.0	33.6	20.9	20.0	50.6	45.9	55.7	23.1	26.4	35.2	47.2	39.7	54.2	53.3	22.5	42.6	32.5	40.5	55.0	40.3

Table 2. Comparison of our direct AP loss minimization approach to surrogate loss function optimization on the object detection task.

size of 512 balances computational complexity and performance, though using a larger batch size (such as 2048) generally results in better performance. For our final results, we use a learning rate of 0.1, a regularization parameter of $1 \cdot 10^{-7}$, and $\epsilon = 0.1$ for all classes.

As baselines, we evaluate a network which uses cross-entropy and is trained separately for each class. Again, the network structure was chosen to be identical and we use the parameters provided by Krizhevsky et al. (2012) for initialization. In addition, we consider the structured SVM algorithm, which optimizes a surrogate of the AP loss. This structured SVM was trained using the same batch size as our direct loss minimization. We use a learning rate of 1, and a regularization parameter of $1 \cdot 10^{-7}$ for all classes. As usual, we compare hinge-AP and pos-AP in the presence of 20% label noise.

Results: Tab. 2 shows competitive results of stochastic direct loss minimization, outperforming the strongest baseline by 0.9. Our direct loss minimization performs better than hinge loss in this case. This is because the data for training detectors are slightly noisier compared to the action classification task, which is likely due to the common method of data augmentation based on intersection-over-union thresholds. To our astonishment, it becomes so hard for hinge-AP to learn well with noise in this detection task that it barely learns anything, while pos-AP only suffers from a reasonable decrease.

5. Conclusion

In this paper we have proposed a direct loss minimization approach to train deep neural networks. We have demonstrated the effectiveness of our approach in the context of maximizing average precision for ranking problems. This involves minimizing a non-smooth and non-decomposable loss. Towards this goal we have proposed a dynamic programming algorithm that can efficiently compute the weight updates. Our experiments showed that this is beneficial when compared to a large variety of baselines in the context of action classification and object detection, particularly in the presence of noisy labels. In the future, we plan to investigate direct loss minimization in the context of

other non-decomposable losses, such as intersection over union for semantic segmentation and shortest-path predictions in graphs.

Acknowledgments

YS would like to thank the Department of Physics, Tsinghua University for providing financial support for his stay in Toronto and travel to ICML. We thank David A. McAllester for discussions and all the reviewers for helpful suggestions. This work was partially supported by ONR Grant N00014-14-1-0232, and a Google researcher award.

References

- Bengio, Y., Goodfellow, I. J., and Courville, A. Deep learning. Book in preparation for MIT Press, 2015. URL <http://www.iro.umontreal.ca/~bengioy/dlbook>.
- Burges, C. J. C. From RankNet to LambdaRank to LambdaMART: An Overview. Technical report, Microsoft Research, 2010.
- Burges, C. J. C., Shaked, T., Renshaw, E., Lazier, A., Deeds, M., Hamilton, N., and Hullender, G. Learning to rank using gradient descent. In *Proc. ICML*, 2005.
- Burges, C. J. C., Ragno, R., and Le, Q. V. Learning to rank with nonsmooth cost functions. *Proc. NIPS*, 2007.
- Chen, L.-C., Schwing, A. G., Yuille, A. L., and Urtasun, R. Learning Deep Structured Models. In *Proc. ICML*, 2015.
- Cheng, C.-C., Sha, F., and Saul, L. K. Matrix updates for perceptron training of continuous density hidden markov models. In *Proc. ICML*, 2009.
- Doerr, A., Ratliff, N., Bohg, J., Toussaint, M., and Schaal, S. Direct loss minimization inverse optimal control. *Proc. of robotics: science and systems (R: SS)*, 2015.
- Everingham, M., Eslami, A. S. M., van Gool, L., Williams, C. K. I., Winn, J., and Zisserman, A. The pascal visual object classes challenge: A retrospective. *IJCV*, 2014.

- Girshick, R., Donahue, J., Darrell, T., and Malik, J. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proc. CVPR*, 2014.
- Keshet, J. and McAllester, D. A. Generalization bounds and consistency for latent structural probit and ramp loss. In *Proc. NIPS*, 2011.
- Keshet, J., Cheng, C.-C., Stoehr, M., and McAllester, D. Direct Error Rate Minimization of Hidden Markov Models. In *Proc. Interspeech*, 2011.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Proc. NIPS*, 2012.
- LeCun, Y. and Huang, F. J. Loss Functions for Discriminative Training of Energy-Based Models. In *Proc. AIS-TATS*, 2005.
- McAllester, D. A., Keshet, J., and Hazan, T. Direct loss minimization for structured prediction. In *Proc. NIPS*, 2010.
- Mohapatra, P., Jawahar, C. V., and Kumar, M. P. Efficient Optimization for Average Precision SVM. In *Proc. NIPS*, 2014.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., and Fei-Fei, L. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 2015.
- Tarlow, D. and Zemel, R. S. Structured Output Learning with High Order Loss Functions. In *Proc. AISTATS*, 2012.
- Tsochantaridis, I., Joachims, T., Hofmann, T., and Altun, Y. Large Margin Methods for Structured and Interdependent Output Variables. *JMLR*, 2005.
- Uijlings, J., van de Sande, K., Gevers, T., and Smeulders, A. Selective search for object recognition. *IJCV*, 2013.
- Volkovs, M. N. and Zemel, R. S. BoltzRank: Learning to Maximize Expected Ranking Gain. In *Proc. ICML*, 2009.
- Yue, Y., Finley, T., Radlinski, F., and Joachims, T. A support vector method for optimizing average precision. In *Proc. SIGIR*, 2007.