

# On Extending Neural Networks with Loss Ensembles for Text Classification

**Hamideh Hajiabadi**

Ferdowsi University of Mashhad (FUM)  
Mashhad, Iran

hamideh.hajiabadi@mail.um.ac.ir

**Diego Molla-Aliod**

Macquarie University  
Sydney, New South Wales, Australia

diego.molla-aliod@mq.edu.au

**Reza Monsefi**

Ferdowsi University of Mashhad (FUM)  
Mashhad, Iran

monsefi@um.ac.ir

## Abstract

Ensemble techniques are powerful approaches that combine several weak learners to build a stronger one. As a meta learning framework, ensemble techniques can easily be applied to many machine learning techniques. In this paper we propose a neural network extended with an ensemble loss function for text classification. The weight of each weak loss function is tuned within the training phase through the gradient propagation optimization method of the neural network. The approach is evaluated on several text classification datasets. We also evaluate its performance in various environments with several degrees of label noise. Experimental results indicate an improvement of the results and strong resilience against label noise in comparison with other methods.

## 1 Introduction

In statistics and machine learning, ensemble methods use multiple learning algorithms to obtain better predictive performance (Mannor and Meir, 2001). It has been proved that ensemble methods can boost weak learners whose accuracies are slightly better than random guessing into arbitrarily accurate strong learners (Bai et al., 2014; Zhang et al., 2016). When it could not be possible to directly design a strong complicated learning system, ensemble methods would be a possible solution. In this paper, we are inspired by ensemble techniques to combine several weak loss functions in order to design a stronger ensemble loss function for text classification.

In this paper we will focus on multi-class classification where the class to predict is encoded as

a vector  $y$  with the one-hot encoding of the target label, and the output of a classifier  $\hat{y} = f(x; \theta)$  is a vector of probability estimates of each label given input sample  $x$  and training parameters  $\theta$ . Then, a loss function  $L(y, \hat{y})$  is a positive function that measures the error of estimation (Steinwart and Christmann, 2008). Different loss functions have different properties, and some well-known loss functions are shown in Table 1. Different loss functions lead to different Optimum Bayes Estimators having their own unique characteristics. So, in each environment, picking a specific loss function will affect performance significantly (Xiao et al., 2017; Zhao et al., 2010).

In this paper, we propose an approach for combining loss functions which performs substantially better especially when facing annotation noise. The framework is designed as an extension to regular neural networks, where the loss function is replaced with an ensemble of loss functions, and the ensemble weights are learned as part of the gradient propagation process. We implement and evaluate our proposed algorithm on several text classification datasets.

The paper is structured as follows. An overview of several loss functions for classification is briefly introduced in Section 2. The proposed framework and the proposed algorithm are explained in Section 3. Section 4 contains experimental results on classifying several text datasets. The paper is concluded in Section 5.

## 2 Background

A typical machine learning problem can be reduced to an expected loss function minimization problem (Bartlett et al., 2006; Painsky and Rosset, 2016). Rosasco et al. (2004) studied the impact of choosing different loss functions from the viewpoint of statistical learning theory. In this section,

| Name of loss function   | $L(y, \hat{y})$   |
|---|---|
| Zero-One (Xiao et al., 2017)  | $L_{0-1} = \begin{cases} 0 & z \geq 0 \\ 1 & z < 0 \end{cases}$   |
| Hinge Loss (Masnadi-Shirazi and Vasconcelos, 2009; Steinwart, 2002) | $L_H = \begin{cases} 0 & z \geq 1 \\ \max(0, 1 - z) & z < 1 \end{cases}$  |
| Smoothed Hinge (Zhao et al., 2010)                                  | $L_{S-H} = \begin{cases} 0 & z \geq 1 \\ \frac{1-z^2}{2} & 0 \leq z < 1 \\ \max(0, 1 - z) & z \leq 0 \end{cases}$ |
| Square Loss   | $L_S = \ y - \hat{y}\ _2^2$   |
| Correntropy Loss (Liu et al., 2007, 2006)                           | $L_C = \exp \frac{\ y - \hat{y}\ _2^2}{\sigma^2}$   |
| Cross-Entropy Loss (Masnadi-Shirazi et al., 2010)                   | $L_{C-E} = \log(1 + \exp(-z))$  |
| Absolute Loss   | $L_A = \ y - \hat{y}\ _1$   |

Table 1: Several well-known loss functions, where  $z = y \cdot \hat{y} \in \mathcal{R}$ .

several well-known loss functions are briefly introduced, followed by a review of ensemble methods.

In the literature, loss functions are divided into margin-based and distance-based categories. Margin-based loss functions are often used for classification purposes (Steinwart and Christmann, 2008; Khan et al., 2013; Chen et al., 2017). Since we evaluate our work on classification of text datasets, in this paper we focus on margin-based loss functions.

A margin-based loss function is defined as a penalty function  $L(y, \hat{y})$  based in a margin  $z = y \cdot \hat{y}$ . In any given application, some margin-based loss functions might have several disadvantages and advantages and we could not certainly tell which loss function is preferable in general. For example, consider the Zero-One loss function which penalizes all the misclassified samples with the constant value of 1 and the correctly classified samples with no loss. This loss function would result in a robust classifier when facing outliers but it would have a terrible performance in an application with margin focus (Zhao et al., 2010).

A loss function is margin enforcing if minimization of the expected loss function leads to a classifier enhancing the margin (Masnadi-Shirazi and Vasconcelos, 2009). Learning a classifier with an acceptable margin would increase generalization. Enhancing the margin would be possible if the loss function returns a small amount of loss for the correct samples close to the classification hyperplane. For example, Zero-One does not penalize correct samples at all and therefore it does not enhance the

margin, while Hinge Loss is a margin enhancing loss function.

The general idea of ensemble techniques is to combine different expert ideas aiming at boosting the accuracy based on enhanced decision making. Predominantly, the underlying idea is that the decision made by a committee of experts is more reliable than the decision of one expert alone (Bai et al., 2014; Mannor and Meir, 2001). Ensemble techniques as a framework have been applied to a variety of real problems and better results have been achieved in comparison to using a single expert.

Having considered the importance of the loss function in learning algorithms, in order to reach a better learning system, we are inspired by ensemble techniques to design an ensemble loss function. The weight applied to each weak loss function is tuned through the gradient propagation optimization of a neural network working on a text classification dataset.

Other works (Shi et al., 2015; BenTaieb et al., 2016) have combined two loss functions where the weights are specified as a hyperparameter set prior to the learning process (e.g. during a fine-tuning process with crossvalidation). In this paper, we combine more than two functions and the hyperparameter is not set a-priori but it is learned during the training process.

### 3 Proposed Approach

Let  $(x, y)$  be a sample where  $x \in \mathcal{R}^N$  is the input and  $y \in \{0, 1\}^C$  is the one-hot encoding of the label ( $C$  is the number of classes). Let

$\theta$  be the parameters of a neural network classifier with a top softmax layer so that the probability estimates are  $\hat{y} = \text{softmax}(f(x; \theta))$ . Let  $\{L_i(y, \hat{y})\}_{i=1}^M$  denote  $M$  weak loss functions. In addition to finding the optimal  $\theta$ , the goal is to find the best weights,  $\{\lambda_1, \lambda_2, \dots, \lambda_M\}$ , to combine  $M$  weak loss functions in order to generate a better application-tailored loss function. We need to add a further constraint to avoid yielding near zero values for all  $\lambda_i$  weights. The proposed ensemble loss function is defined as below.

$$L = \sum_{j=1}^M \lambda_j L_j(y, \hat{y}), \sum_{j=1}^M \lambda_j = 1 \quad (1)$$

The optimization problem could be defined as follows, given  $T$  training samples.

$$\begin{aligned} & \underset{\theta, \lambda}{\text{minimize}} \sum_{i=1}^T \sum_{j=1}^M \lambda_j L_j(y_i, \hat{y}_i) \\ & \text{s.t.} \sum_{j=1}^M \lambda_j = 1, \lambda_i \geq 0 \end{aligned} \quad (2)$$

To make the optimization algorithm simpler, we use  $\lambda_i^2$  instead of  $\lambda_i$ , so the second constraint  $\lambda_i \geq 0$  can be omitted. We then incorporate the constraint as a regularization term based on the concept of Augmented Lagrangian. The modified objective function using Augmented Lagrangian is presented as follows.

$$\begin{aligned} & \underset{\theta, \lambda}{\text{minimize}} \sum_{i=1}^T \sum_{j=1}^M \lambda_j^2 L_j(y_i, \hat{y}_i) + \\ & \eta_1 \left( \sum_{j=1}^M \lambda_j^2 - 1 \right) + \eta_2 \left( \sum_{j=1}^M \lambda_j^2 - 1 \right)^2 \end{aligned} \quad (3)$$

Note that the amount of  $\eta_2$  must be significantly greater than  $\eta_1$  (Nocedal and Wright, 2006). The first and the second terms of the objective function cause  $\lambda_j^2$  values to approach zero but the third term satisfies  $\sum_{j=1}^M \lambda_j^2 = 1$ .

Figure 1 illustrates the framework of the proposed approach with the dashed box representing the contribution of this paper. In the training phase, the weight of each weak loss function is trained through the gradient propagation optimization method. The accuracy of the model is calculated in a test phase not shown in the figure.

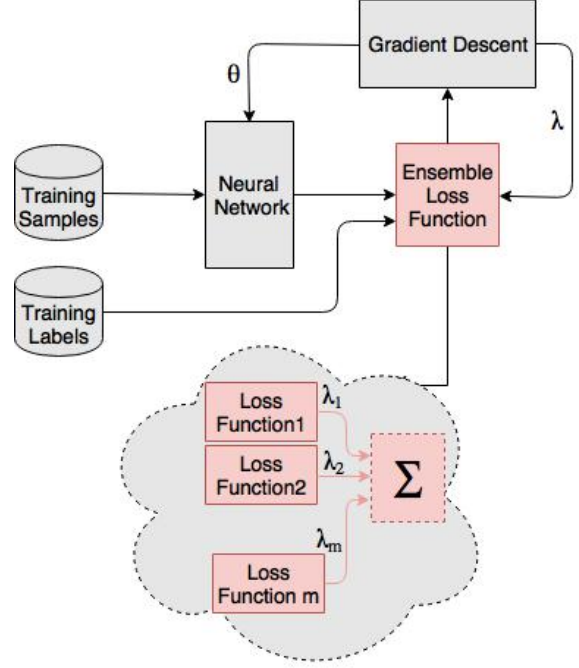


Figure 1: The proposed learning diagram

## 4 Experimental Results

We have applied the proposed ensemble loss function to several text datasets. Table 2 provides a brief description of the datasets. To reach a better ensemble loss function we choose three loss functions with different approaches in facing with outliers, as weak loss functions: Correntropy Loss which does not assign a high weight to samples with big errors, Hinge Loss which penalizes linearly and Cross-entropy Loss function which highly penalizes the samples whose predictions are far from the targets. We compared results with 3 loss functions which are widely used in neural networks: Cross-entropy, Square Loss, and Hinge Loss.

We picked  $\eta_1$  near zero and  $\eta_2 = 200$  in (3).

Since this work is a proof of concept, the neural networks of each application are simply a softmax of the linear combination of input features plus bias:

$$\hat{y} = \text{softmax}(x \cdot W + b)$$

where the input features  $x$  are the word frequencies in the input text. Thus,  $\theta$  in our notation is composed of  $W$  and  $b$ . We use Python and its TensorFlow package for implementing the proposed approach. The results are shown in Table 3. The table compares the results of using individual loss functions and the ensemble loss.

| Name of Datasets            | Description  |
|-----------------------------|--|
| 20-newsgroup                | This data set is a collection of 20,000 messages, collected from 20 different net-news newsgroups.   |
| Movie-reviews in corpus     | The NLTK corpus movie-reviews data set has the reviews, and they are labeled already as positive or negative.  |
| Email-Classification (TREC) | It is a collection of sample emails (i.e. a text corpus). In this corpus, each email has already been labeled as Spam or Ham.                                |
| Reuters-21578               | The data was originally collected and labeled by Carnegie Group, Inc. and Reuters, Ltd. in the course of developing the CON-STRUE text categorization system |

Table 2: Description of dataset

| Dataset                     | Cross-entropy | Hinge | Square      | Ensemble    |
|-----------------------------|---------------|-------|-------------|-------------|
| 20-newsgroups               | 0.80          | 0.69  | 0.82        | <b>0.85</b> |
| Movie-review                | 0.83          | 0.81  | <b>0.85</b> | 0.83        |
| Email-Classification (TREC) | 0.88          | 0.78  | 0.96        | <b>0.97</b> |
| Reuters                     | 0.79          | 0.79  | <b>0.81</b> | <b>0.81</b> |

Table 3: Accuracy

We have also compared the robustness of the proposed loss function with the use of individual loss functions. In particular, we add label noise by randomly modifying the target label in the training samples, and keep the evaluation set intact. We conducted experiments with 10% and 30% of noise, where e.g. 30% of noise means randomly changing 30% of the labels in the training data. Tables 4 and 5 show the results, with the best results shown in boldface. We can observe that, in virtually all of the experiments, the ensemble loss is at least as good as the individual losses, and in only two cases the loss is (slightly) worse. And, in general, the ensemble loss performed comparatively better as we increased the label noise.

| Dataset                     | Cross-entropy | Hinge | Square | Ensemble    |
|-----------------------------|---------------|-------|--------|-------------|
| 20-newsgroups               | 0.79          | 0.67  | 0.69   | <b>0.83</b> |
| Movie-reviews               | 0.75          | 0.74  | 0.73   | <b>0.78</b> |
| Email-Classification (TREC) | 0.86          | 0.57  | 0.82   | <b>0.96</b> |
| Reuters                     | <b>0.76</b>   | 0.69  | 0.71   | 0.73        |

Table 4: Accuracy in data with 10% label noise

## 5 Conclusion

This paper proposed a new loss function based on ensemble methods. This work focused on text classification tasks and can be considered as an initial attempt to explore the use of ensemble loss functions. The proposed loss function shows an improvement when compared with the use of well-known individual loss functions. Furthermore, the approach is more robust against the presence of label noise. Moreover, according to our experiments, the gradient descent method quickly converged.

| Dataset                     | Cross-entropy | Hinge | Square | Ensemble    |
|-----------------------------|---------------|-------|--------|-------------|
| 20-newsgroups               | 0.57          | 0.64  | 0.55   | <b>0.82</b> |
| movie-review                | 0.55          | 0.54  | 0.55   | <b>0.6</b>  |
| Email-Classification (TREC) | 0.80          | 0.46  | 0.81   | <b>0.93</b> |
| Reuters                     | 0.64          | 0.54  | 0.53   | <b>0.68</b> |

Table 5: Accuracy in data with 30% label noise

We have used a very simple neural architecture in this work but in principle this method could be used for systems that use any neural networks. In future work we will explore the integration of more complex neural networks such as those using convolutions and recurrent networks. We also plan to study the application of this method to other tasks such as sequence labeling (e.g. for NER and PoS tagging). Another possible extension could focus on handling sparseness by adding a regularization term.

## References

- Qinxun Bai, Henry Lam, and Stan Sclaroff. 2014. A Bayesian framework for online classifier ensemble. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 1584–1592.
- Peter L Bartlett, Michael I Jordan, and Jon D McAuliffe. 2006. Convexity, classification, and risk bounds. *Journal of the American Statistical Association* 101(473):138–156.
- Aïcha BenTaieb, Jeremy Kawahara, and Ghassan Hamarneh. 2016. Multi-loss convolutional networks for gland analysis in microscopy. In *Biomedical Imaging (ISBI 2016)*, pages 642–645.
- Badong Chen, Lei Xing, Bin Xu, Haiquan Zhao, Nanning Zheng, and Jose C Principe. 2017. Kernel risk-sensitive loss: Definition, properties and application to robust adaptive filtering. *IEEE Transactions on Signal Processing* 65(11):2888–2901.
- Inayatullah Khan, Peter M Roth, Abdul Bais, and Horst Bischof. 2013. Semi-supervised image classification with huberized Laplacian support vector machines. In *Emerging Technologies (ICET), 2013 IEEE 9th International Conference on*. IEEE, pages 1–6.
- W. Liu, P. P. Pokharel, and J. C. Principe. 2007. **Correntropy: Properties and applications in non-Gaussian signal processing**. *IEEE Transactions on Signal Processing* 55(11):5286–5298. <https://doi.org/10.1109/TSP.2007.896065>.
- Weifeng Liu, P. P. Pokharel, and J. C. Principe. 2006. **Correntropy: A localized similarity measure**. In *The 2006 IEEE International Joint Conference on Neural Network Proceedings*, pages 4919–4924. <https://doi.org/10.1109/IJCNN.2006.247192>.
- Shie Mannor and Ron Meir. 2001. Weak learners and improved rates of convergence in boosting. In *Advances in Neural Information Processing Systems*, pages 280–286.
- Hamed Masnadi-Shirazi, Vijay Mahadevan, and Nuno Vasconcelos. 2010. On the design of robust classifiers for computer vision. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*. IEEE, pages 779–786.
- Hamed Masnadi-Shirazi and Nuno Vasconcelos. 2009. On the design of loss functions for classification: theory, robustness to outliers, and SavageBoost. In *Advances in neural information processing systems*, pages 1049–1056.
- Jorge Nocedal and Stephen J Wright. 2006. Penalty and augmented Lagrangian methods. *Numerical Optimization* pages 497–528.
- Amichai Painsky and Saharon Rosset. 2016. Isotonic modeling with non-differentiable loss functions with application to Lasso regularization. *IEEE transactions on pattern analysis and machine intelligence* 38(2):308–321.
- Lorenzo Rosasco, Ernesto De Vito, Andrea Caponnetto, Michele Piana, and Alessandro Verri. 2004. Are loss functions all the same? *Neural Computation* 16(5):1063–1076.
- Qinfeng Shi, Mark Reid, Tiberio Caetano, Anton Van Den Hengel, and Zhenhua Wang. 2015. **A hybrid loss for multiclass and structured prediction**. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 37(1):2–12. <https://doi.org/10.1109/TPAMI.2014.2306414>.
- Ingo Steinwart. 2002. Support vector machines are universally consistent. *Journal of Complexity* 18(3):768–791.
- Ingo Steinwart and Andreas Christmann. 2008. *Support Vector Machines*. Springer Science & Business Media.
- Yingchao Xiao, Huangang Wang, and Wenli Xu. 2017. Ramp loss based robust one-class SVM. *Pattern Recognition Letters* 85:15–20.
- Peng Zhang, Tao Zhuo, Yanning Zhang, Hanqiao Huang, and Kangli Chen. 2016. Bayesian tracking fusion framework with online classifier ensemble for immersive visual applications. *Multimedia Tools and Applications* 75(9):5075–5092.
- Lei Zhao, Musa Mammadov, and John Yearwood. 2010. From convex to nonconvex: A loss function analysis for binary classification. In *Data Mining Workshops (ICDMW), 2010 IEEE International Conference on*. IEEE, pages 1281–1288.