

UNIVERSUM PRESCRIPTION: REGULARIZATION USING UNLABELED DATA

Xiang Zhang & Yann LeCun

{xiang, yann}@cs.nyu.edu

Courant Institute of Mathematical Sciences, New York University

719 Broadway, 12th Floor, New York, NY 10003, USA

ABSTRACT

This paper shows that simply prescribing “none of the above” labels to unlabeled data has a beneficial regularization effect to supervised learning. We call it universum prescription by the fact that the prescribed labels cannot be one of the supervised labels. In spite of its simplicity, universum prescription obtained competitive results in training deep convolutional networks for CIFAR-10, CIFAR-100 and STL-10 datasets. A qualitative justification of these approaches using Rademacher complexity is presented. The effect of a regularization parameter – probability of sampling from unlabeled data – is also studied empirically.

1 INTRODUCTION

The idea of exploiting the wide abundance of unlabeled data to improve the accuracy of supervised learning tasks is a very natural one. In this paper, we study what is perhaps the simplest way to exploit unlabeled data in the context of deep learning. We assume that the unlabeled samples do not belong to any of the categories of the supervised task, and we force the classifier to produce a “none of the above” output for these samples. This is by no means a new idea, but we show empirically and theoretically that doing so has a beneficial regularization effect on supervised task and reduces the generalization gap, the expected difference between the test error and the training error. We study three different ways to prescribe “none of the above” outputs, dubbed uniform prescription, dustbin class, and background class and show that they improve the test error of convolutional networks trained on CIFAR-10, CIFAR-100 (Krizhevsky (2009)), and STL-10 (Coates et al. (2011)). The method is justified theoretically using Rademacher complexity (Bartlett & Mendelson (2003)).

Our work is a direct extension to learning in the presence of universum (Weston et al. (2006)), originated from Vapnik (1998) and Vapnik (2006). The definition of universum is a set of unlabeled data that are known not to belong to any of the classes but in the same domain of the training data. We extended the idea of using universum from support vector machines to deep learning.

Using unlabeled data to facilitate supervised learning is sometimes called semi-supervised learning as surveyed by Chapelle et al. (2006b) and Zhu & Goldberg (2009). The most related ones are information regularization (Corduneanu & Jaakkola (2006)) and transduction learning (Chapelle et al. (2006a)) (Gammerman et al. (1998)). In these approaches, prescribing supervised labels to unlabeled data is part of the overall algorithm. They are the opposite case of universum prescription.

Representation or feature learning (reviewed by Bengio et al. (2013) and Bengio & LeCun (2007)) and transfer learning (Thrun & Pratt (1998)) are also related to our work. They include the idea of pretraining (Erhan et al. (2010) Hinton et al. (2006) Ranzato et al. (2006)), which transfers the features learnt from unlabeled data to some supervised task. Universum prescription incorporates unlabeled data as part of the supervised training process, imposing neither sparsity nor reconstruction.

The methods in this article could be thought of as a simple form of multi-task learning (Baxter (2000)) (Caruana (1993)), where an auxiliary task is to control overfitting under the universum assumption (see section 2). It can also be thought of as using hints (Abu-Mostafa (1990)) (Sudderth & Holden (1991)) for training where the hint is functional regularity from unlabeled data.

Universum prescription is also related to the idea of dark knowledge (Bucilu et al. (2006)) (Hinton et al. (2015)). The idea is to prescribe “soft” targets from an ensemble of models to a single model,

and improvement on classification accuracy is observed. Our methods prescribe “soft” targets to unlabeled data as well, except that the targets are agnostic to the classification problem.

Regularization – techniques for the control of overfitting or generalization gap – has been studied extensively. Most of the practical approaches implement a secondary optimization objective, such as L_1 or L_2 norm. Some other methods such as dropout (Srivastava et al. (2014)) and dropconnect (Wan et al. (2013)) cheaply simulate model averaging to control the model variance.

As part of the general statistical learning theory (Vapnik (1995), Vapnik (1998)), the justification for regularization is well-developed. There are many formulations, such as probably approximately correct (PAC) learning (Valiant (1984)), the trade-off between bias and variance (Geman et al. (1992)), and the prescription of Bayesian a priori (Mozer & Smolensky (1989)). We qualitatively justify the methods using Radamacher complexity (Bartlett & Mendelson (2003)), similar to Wan et al. (2013).

2 UNIVERSUM PRESCRIPTION

In this section we attempt to formalize the the trick of prescribing “none of the above” labels. We call it universum prescription because these labels could not belong to any supervised class. Consider the problem of exclusive k -way classification. In inference we can find the most probable class $y \in \{1, 2, \dots, k\}$ given input x . In learning we hope to find a hypothesis function $h \in \mathcal{H}$ mapping to \mathbb{R}^k so that the label is determined by $y = \operatorname{argmin}_i h_i(x)$. The following assumptions are made.

1. (Loss assumption) The loss used as the optimization objective is negative log-likelihood:

$$L(h, x, y) = h_y(x) + \log \left[\sum_{i=1}^k \exp(-h_i(x)) \right]. \quad (1)$$

2. (Universum assumption) The proportion of samples belonging to one of the k classes in the unlabeled data is negligible.

The loss assumption assumes that the probability of class y given an input x can be thought of as

$$\Pr[Y = y|x, h] = \frac{\exp(-h_y(x))}{\sum_{i=1}^k \exp(-h_i(x))}, \quad (2)$$

where $(X, Y) \sim \mathbf{D}$ and \mathbf{D} is the distribution where labeled data are sampled. We use lowercase letters for values, uppercase letters for random variables and bold uppercase letters for distribution. The loss assumption is simply a necessary detail rather than a limitation.

The universum assumption implicates that labeled classes are a negligible subset. In many practical cases we only care about a small number of classes, either by problem design or due to high cost in the labeling process. At the same time, a very large amount of unlabeled data is easily obtained. Put in mathematics, assuming we draw unlabeled data from distribution \mathbf{U} , the assumption states that

$$\Pr_{(X, Y) \sim \mathbf{U}}[X, Y \in \{1, 2, \dots, k\}] \approx 0. \quad (3)$$

This is opposite to the assumptions of information regularization (Corduneanu & Jaakkola (2006)) and transduction learning (Chapelle et al. (2006a)) (Gammerman et al. (1998)). All the methods discussed below prescribe agnostic targets to the unlabeled data. During learning, we randomly present an unlabeled sample to the optimization procedure with probability p .

2.1 UNIFORM PRESCRIPTION

It is known that negative log-likelihood is simply a reduced form of cross-entropy

$$L(h, x, y) = - \sum_{i=1}^k Q[Y = i|x] \log \Pr[Y = i|x, h] \quad (4)$$

in which the target probability $Q[Y = y|x] = 1$ and $Q[Y = i|x] = 0$ for $i \neq y$. Under the universum assumption, if we are presented with an unlabeled sample x , we would hope to prescribe some Q so

that every class has some equally minimal probability. Q also has to satisfy $\sum_{i=1}^k Q[Y = i|x] = 1$ by the probability axioms. The only possible choice for Q is then $Q[Y|x] = 1/k$. The learning algorithm then uses the cross-entropy loss instead of negative log-likelihood.

It is worth noting that uniform output has the maximum entropy among all possible choices. In the case that the hypothesis h is parameterized as a deep neural network, uniform output is achieved when these parameters are constantly 0. Therefore, uniform prescription may have the effect of reducing the magnitude of parameters, similar to norm-based regularization.

2.2 DUSTBIN CLASS

Another way of prescribing agnostic target is to append a “dustbin” class to the supervised task. This requires some changes to the hypothesis function h such that it outputs $k + 1$ targets. For deep learning models one can simply extend the last parameterized layer. All unlabeled data are prescribed to this extra “dustbin” class. The learning algorithm remains unchanged.

The effect of dustbin class is clearly seen in the loss function of an unlabeled sample $(x, k + 1)$

$$L(h, x, k + 1) = h_{k+1}(x) + \log \left[\sum_{i=1}^{k+1} \exp(-h_i(x)) \right]. \quad (5)$$

The second term is a “soft” maximum for all dimensions of $-h$. When an unlabeled sample is present, the algorithm attempts to introduce smoothness by minimizing probability spikes.

2.3 BACKGROUND CLASS

We could further simplify dustbin class by removing parameters for class $k + 1$. For some given threshold constant τ , we could change the probability of a labeled sample to

$$\Pr[Y = y|x, h] = \frac{\exp(-h_y(x))}{\exp(-\tau) + \sum_{i=1}^k \exp(-h_i(x))}, \quad (6)$$

and an unlabeled sample

$$\Pr[Y = k + 1|x, h] = \frac{\exp(-\tau)}{\exp(-\tau) + \sum_{i=1}^k \exp(-h_i(x))}. \quad (7)$$

This will result in changes to the loss function of a labeled sample (x, y) as

$$L(h, x, y) = h_y(x) + \log \left[\exp(-\tau) + \sum_{i=1}^k \exp(-h_i(x)) \right], \quad (8)$$

and an unlabeled sample

$$L(h, x, k + 1) = \tau + \log \left[\exp(-\tau) + \sum_{i=1}^k \exp(-h_i(x)) \right]. \quad (9)$$

We call this method background class and τ background constant. Similar to dustbin class, the algorithm attempts to minimize the spikes of outputs, but limited to a certain extent by the inclusion of $\exp(-\tau)$ in the partition function. In our experiments τ is always set to 0.

3 THEORETICAL JUSTIFICATION

In this part, we derive a qualitative justification for universum prescription using probably approximately correct (PAC) learning (Valiant (1984)). By a “qualitative” theory, we are comparing with numerical or combinatorial bounds such as growth function (Massart (2000), Vapnik (1998)), Vapnik-Chervonenkis dimension (Vapnik & Chervonenkis (1971)), covering numbers (Dudley (1967)) and others. Our theory is based on Rademacher complexity (Bartlett & Mendelson (2003)), similar to the work by Wan et al. (2013) in which both dropout (Srivastava et al. (2014)) and dropconnect (Wan et al. (2013)) are justified. Rademacher complexity is usually a lower-bound of other numerical or combinatorial complexity measurement, therefore our qualitative intuition is more accurate using it.

Definition 1 (Empirical Rademacher complexity). *Let \mathcal{F} be a family of functions mapping from \mathcal{U} to \mathbb{R} , and $S = (x_1, x_2, \dots, x_m)$ a fixed sample of size m with elements in \mathcal{X} . Then, the empirical Rademacher complexity of \mathcal{F} with respect to the sample S is defined as:*

$$\hat{\mathfrak{R}}_S(\mathcal{F}) = \mathbb{E}_{\boldsymbol{\eta}} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^m \eta_i f(x_i) \right] \quad (10)$$

where $\boldsymbol{\eta} = (\eta_1, \dots, \eta_m)^T$, with η_i 's independent random variables taking values from a discrete uniform distribution on $\{-1, 1\}$.

Definition 2 (Rademacher complexity). *Let \mathbf{D} denote the distribution from which the samples were drawn. For any integer $m \geq 1$, the Rademacher complexity of \mathcal{F} is the expectation of the empirical Rademacher complexity over all samples of size m drawn according to \mathbf{D} :*

$$\mathfrak{R}_m(\mathcal{F}, \mathbf{D}) = \mathbb{E}_{S \sim \mathbf{D}^m} [\hat{\mathfrak{R}}_S(\mathcal{F})] \quad (11)$$

It could be argued that the distribution for $\boldsymbol{\eta}$ is arbitrary. There are other possibilities such as Gaussian complexity (Bartlett & Mendelson (2003)), but they can all be generalized to stochastic complexity as in Zhang (2013) and result in the same conclusions. In the case that f has multiple outputs, one can simply add the complexity measurements for each outputs together and the theory still holds.

Two qualitative properties of Rademacher complexity is worth noting here. First of all, Rademacher complexity is always non-negative by the convexity of supremum

$$\hat{\mathfrak{R}}_S(\mathcal{F}) = \mathbb{E}_{\boldsymbol{\eta}} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^m \eta_i f(x_i) \right] \geq \sup_{f \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^m \mathbb{E}[\eta_i] f(x_i) = 0. \quad (12)$$

Secondly, if for a fixed input all functions in \mathcal{F} output the same value, then it's Rademacher complexity is 0. Assume for any $f \in \mathcal{F}$ we have $f(x) = f_0(x)$, then

$$\hat{\mathfrak{R}}_S(\mathcal{F}) = \mathbb{E}_{\boldsymbol{\eta}} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^m \eta_i f(x_i) \right] = \mathbb{E}_{\boldsymbol{\eta}} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^m \eta_i f_0(x) \right] = \frac{1}{m} \sum_{i=1}^m \mathbb{E}[\eta_i] f_0(x) = 0. \quad (13)$$

Therefore, one way to qualitatively minimize Rademacher complexity is to regularize functions in \mathcal{F} such that all functions tend to have the same output for a given input. Universum prescription precisely does that – the prescribed outputs for unlabeled data are all constantly the same.

The principal PAC-learning result from literature is an approximation bound for function spaces that has finite bounds for outputs. We use the formulation by Zhang (2013), but anterior results are in Bartlett et al. (2002), Bartlett & Mendelson (2003), Koltchinskii (2001) and Koltchinskii & Panchenko (2000). We refer the reader to these publications for proof.

Theorem 1 (Approximation bound with finite bound on output). *For a well-defined objective $\mathcal{E}(h, x, y)$ over hypothesis class \mathcal{H} , input set \mathcal{X} and output set \mathcal{Y} , if it has an upper bound $M > 0$, then with probability at least $1 - \delta$, the following holds for all hypothesis $h \in \mathcal{H}$:*

$$\mathbb{E}_{(x,y) \sim \mathbf{D}} [\mathcal{E}(h, x, y)] \leq \frac{1}{m} \sum_{(x,y) \in S} \mathcal{E}(h, x, y) + 2\mathfrak{R}_m(\mathcal{F}, \mathbf{D}) + M \sqrt{\frac{\log \frac{2}{\delta}}{2m}}, \quad (14)$$

where the function family \mathcal{F} is defined as

$$\mathcal{F} = \{\mathcal{E}(h, x, y) | h \in \mathcal{H}\}, \quad (15)$$

\mathbf{D} is a distribution on the samples (x, y) , and S is a set of samples of size m drawn indentially and independently from \mathbf{D} .

In the theorem above, the objective functional $\mathcal{E}(h, x, y)$ should be lower-bounded by 0, and it corresponds to a negatively correlated compatibility measurement between a hypothesis h and a sample (x, y) . It is similar to the definition of energy used by energy-based learning in LeCun et al. (2006). It could be the error function $\mathcal{E}(h, x, y) = 1 - \mathbf{1}\{y = \operatorname{argmin}_i(h_i(x))\}$, the exponential function $\mathcal{E}(h, x, y) = \exp(h_y(x))$, the negative probability function $\mathcal{E}(h, x, y) = 1 - \Pr[Y = y | x, h]$, or simply the loss $\mathcal{E}(h, x, y) = L(h, x, y)$.

Table 2: Result for universum prescription. The numbers are percentages. The three numbers in each tabular indicate training error, testing error and generalization gap. Bold numbers are the best ones for each case. CIFAR-100 F. and CIFAR-100 C. stand for fine-grained and coarse classification problems of CIFAR-100. STL-10 Tiny. stands for using 80 million images as the unlabeled dataset.

| DATASET | BASELINE | | | UNIFORM | | | DUSTBIN | | | BACKGROUND | | |
|--------------|-------------|--------------|-------|---------|-------|-------|---------|--------------|--------------|------------|-------|--------------|
| | Train | Test | Gap | Train | Test | Gap | Train | Test | Gap | Train | Test | Gap |
| CIFAR-10 | 0.00 | 7.02 | 7.02 | 0.72 | 7.59 | 6.87 | 0.07 | 6.66 | 6.59 | 1.35 | 8.38 | 7.03 |
| CIFAR-100 F. | 0.09 | 37.58 | 37.49 | 4.91 | 36.23 | 31.32 | 2.52 | 32.84 | 30.32 | 8.56 | 40.57 | 42.01 |
| CIFAR-100 C. | 0.04 | 22.74 | 22.70 | 0.67 | 23.42 | 22.45 | 0.40 | 20.45 | 20.05 | 3.73 | 24.97 | 21.24 |
| STL-10 | 0.00 | 31.16 | 31.16 | 2.02 | 36.54 | 34.52 | 3.03 | 36.58 | 33.55 | 14.89 | 38.95 | 24.06 |
| STL-10 Tiny. | 0.00 | 31.16 | 31.16 | 0.62 | 30.15 | 29.47 | 0.00 | 27.96 | 27.96 | 0.11 | 30.38 | 30.27 |

For some choices of \mathcal{E} we have a bound $\mathcal{E}(h, x, y) \leq M$ by design, whereas for some others it is more intricate to believe M exists. If the learning algorithm is an iterative optimization procedure such as gradient descent, at each step one could believe that a limit M exists relatively to the current hypothesis h_0 (Zhang (2013)). This is because of the dynamics of iterative optimization – the algorithm can only explore some sublevel hypothesis set \mathcal{H}_0 in later steps.

The meaning of the theorem is two-folded. When applying the theorem to the joint problem of training using both labeled and unlabeled data, the third term on the right hand of inequality 14 is reduced by the augmentation of the extra data. The joint problem can be written as $(x, y) \sim (1 - p)\mathbf{D} + p\mathbf{U}$. The value of the term $\mathfrak{R}_m(\mathcal{F}, (1 - p)\mathbf{D} + p\mathbf{U})$ is also reduced when we prescribe constant outputs, due to the qualitative properties of Rademacher complexity discussed before.

The second fold is that when the theorem applies to the supervised distribution \mathbf{D} , we would hope that $\mathfrak{R}_m(\mathcal{F}, \mathbf{D})$ can be bounded by $\mathfrak{R}_m(\mathcal{F}, (1 - p)\mathbf{D} + p\mathbf{U})$ and $\mathfrak{R}_m(\mathcal{F}, \mathbf{U})$. It turns out that a weighted sum of $\mathfrak{R}_i(\mathcal{F}, \mathbf{D}), i = 1, 2, \dots, m$ is bounded.

Theorem 2 (Rademacher complexity bound on distribution mixture). *Let $P(m, i) = \sum_{i=0}^m \binom{m}{i} (1 - p)^{m-i} p^i \frac{i}{m}$ and $Q(m, i) = \sum_{i=0}^m \binom{m}{i} (1 - p)^i p^{m-i} \frac{i}{m}$, we have*

$$\sum_{i=1}^m Q(m, i) \mathfrak{R}_i(\mathcal{F}, \mathbf{D}) \leq \mathfrak{R}_m(\mathcal{F}, (1 - p)\mathbf{D} + p\mathbf{U}) + \sum_{i=1}^m P(m, i) \mathfrak{R}_i(\mathcal{F}, \mathbf{U}) \quad (16)$$

The proof of theorem 2 is in supplemental material. We also show that $\sum_{i=0}^m P(m, i) \leq 1$ and $\sum_{i=0}^m Q(m, i) \leq 1$. The derivation above tells us that a weighted sum of the Rademacher complexity of the supervised problems for dataset size from 1 to m is bounded by the joint problem of size m and a weighted sum of unsupervised problems. Therefore, for different sample sizes of labeled and unlabeled data, universum prescription may bring improvement for generalization.

4 EXPERIMENTS ON IMAGE CLASSIFICATION

In this section we test the methods on some image classification tasks. Two series of datasets – CIFAR-10/100 (Krizhevsky (2009)) and STL-10 (Coates et al. (2011)) – are chosen due to the availability of unlabeled data. The model we used is a 21-layer convolutional network (ConvNet) (LeCun et al. (1989), LeCun et al. (1998)) inspired by Simonyan & Zisserman (2014), in which the inputs are 32-by-32 images and all convolutional layers are 3-by-3 and fully padded. All pooling layers are max-pooling, and ReLUs (Nair & Hinton (2010)) are used as the non-linearity after all convolutional and linear layers. Two dropout (Srivastava et al. (2014)) layers of probability 0.5 are inserted before the final two linear layers. The algorithm used is stochastic

Table 1: ConvNet for section 4

| LAYERS | DESCRIPTION |
|--------|---------------|
| 1-3 | Conv 256x3x3 |
| 4 | Pool 2x2 |
| 5-8 | Conv 512x3x3 |
| 9 | Pool 2x2 |
| 10-13 | Conv 1024x3x3 |
| 14 | Pool 2x2 |
| 15-18 | Conv 1024x3x3 |
| 19 | Pool 2x2 |
| 20-23 | Conv 2048x3x3 |
| 24 | Pool 2x2 |
| 25-26 | Full 2048 |

gradient descent with momentum (Polyak (1964), Sutskever et al. (2013)) 0.9 and a minibatch size of 32. The initial learning rate is 0.005 which is halved every 60,000 minibatch steps. The training stops at 400,000 minibatch steps. Table 1 summarizes the configurations. For a convolutional layer the number of output feature maps is shown, and for a linear layer the number of hidden units. The weights are initialized in the same way as He et al. (2015).

The initial motivation for choosing such a big network is to make sure it will have enough capacity for overfitting so that the effect of regularization is clearly shown. However, in practice such a large network already has a very good baseline even without universum prescription. This is probably due to the data augmentation steps below, which are used in all our experiments.

1. (Horizontal flip.) Flip the image horizontally with probability 0.5.
2. (Scale.) Randomly scale the image between $1/1.2$ and 1.2 times of its height and width.
3. (Crop.) Randomly crop a 32-by-32 region in the scaled image.
4. (Rotation.) Randomly rotate between $-\pi/6$ and $\pi/6$ radians.

4.1 CIFAR-10 AND CIFAR-100

The samples of CIFAR-10 and CIFAR-100 datasets (Krizhevsky (2009)) are from the 80 million tiny images dataset (Torralba et al. (2008)). Each dataset contains 60,000 samples, constituting a very small portion of 80 million. This is an ideal case for our methods, in which we can use the entire 80 million images as the unlabeled data. The CIFAR-10 dataset has 10 classes, and CIFAR-100 has 20 (coarse) or 100 (fine-grained) classes. Table 2 contains the results. The generalization gap is approximated by the difference between testing and training errors. All of the universum prescription models use unlabeled data with probability $p = 0.2$.

We compared other single-model results on CIFAR-10 and CIFAR-100 (fine-grained case) in table 3. It shows that our network is competitive to the state of the art.

Table 3: Comparison of single-model CIFAR-10 and CIFAR-100 results, in second and third columns. The fourth column indicates whether data augmentation is used for CIFAR-10. The numbers are percentages.

4.2 STL-10

The STL-10 dataset (Coates et al. (2011)) has size 96-by-96 for its images. We downsampled them to 32-by-32 so as to use the same model. The dataset contains a very small number of training samples – 5000 in total. The accompanying unlabeled dataset is larger with 100,000 samples. There is no guarantee that these

extra samples are outside of the supervised training classes. Both the dataset size and failure to comply with universum assumption might be the reason why universum prescription failed.

To verify that the extra data is the problem, we also performed an experiment using the 80 million tiny images as the unlabeled dataset, as shown in table 2. Due to long training times of our models, we did not perform 10-fold training as in the original paper by Coates et al. (2011), therefore our result is not comparable to those in the literature. We present them only to show the effectiveness of universum prescription influenced by the universum assumption on the unlabeled data.

5 EFFECT OF THE REGULARIZATION PARAMETER

One natural question to ask of our models is how would change of the probability p of sampling from unlabeled data affect the results. In this section we show the experiments. To prevent an exhaustive search on the regularization parameter from overfitting our models on the testing data, we use a different model for this section. It is described in table 4, which has 9 parameterized layers in total.




Figure 1: Experiments on regularization parameter. The four rows are CIFAR-10, CIFAR-100 fine-grained, CIFAR-100 coarse and STL-10 respectively.

The design is inspired by Sermanet et al. (2013). For each choice of p we conducted 6 experiments combining universum prescription models and dropout. The dropout layers are two ones added in between the fully-connected layers with dropout probability 0.5. Figure 1 shows the results.

From figure 1 we can conclude that increasing p will decrease generalization gap. However, we cannot make p too large since after a certain point the training collapses and both training and testing errors become worse. Comparing between CIFAR-10/100 and STL-10, the model variance is affected by the combined size of labeled and unlabeled datasets.

6 CONCLUSION AND OUTLOOK

This article shows that universum prescription can be used to regularize a multi-class classification problem using extra unlabeled data. Two assumptions are made, in which one is that loss used is negative log-likelihood and the other is negligible probability of a supervised sample existing in the unlabeled data. The loss assumption is a necessary detail rather than a limitation. The three universum prescription methods are uniform prescription, dustbin class and background class. We further provided a theoretical justification. Experiments are done using CIFAR-10, CIFAR-100 and STL-10 datasets. The effect of the regularization parameter is also studied empirically.

In the future, we hope to apply these methods to a broader range of problems.

Table 4: ConvNet for section 5

| LAYERS | DESCRIPTION |
|--------|---------------|
| 1 | Conv 1024x5x5 |
| 2 | Pool 2x2 |
| 3 | Conv 1024x5x5 |
| 4-7 | Conv 1024x3x3 |
| 8 | Pool 2x2 |
| 9-11 | Full 2048 |

ACKNOWLEDGMENTS

We gratefully acknowledge the support of NVIDIA Corporation with the donation of 2 Tesla K40 GPUs used for this research. Sainbayar Sukhbaatar offered many useful comments. Aditya Ramesh and Junbo Zhao helped cross-checking the proofs.

REFERENCES

- Abu-Mostafa, Yaser S. Learning from hints in neural networks. *Journal of complexity*, 6(2):192–198, 1990.
- Bartlett, Peter L and Mendelson, Shahar. Rademacher and gaussian complexities: Risk bounds and structural results. *The Journal of Machine Learning Research*, 3:463–482, 2003.
- Bartlett, Peter L, Boucheron, Stéphane, and Lugosi, Gábor. Model selection and error estimation. *Machine Learning*, 48(1-3):85–113, 2002.
- Baxter, Jonathan. A model of inductive bias learning. *J. Artif. Int. Res.*, 12(1):149–198, March 2000. ISSN 1076-9757.
- Bengio, Yoshua and LeCun, Yann. Scaling learning algorithms towards ai. In Bottou, L., Chapelle, O., DeCoste, D., and Weston, J. (eds.), *Large-Scale Kernel Machines*. MIT Press, 2007.
- Bengio, Yoshua, Courville, Aaron, and Vincent, Pierre. Representation learning: A review and new perspectives. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 35(8):1798–1828, 2013.
- Bucilu, Cristian, Caruana, Rich, and Niculescu-Mizil, Alexandru. Model compression. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 535–541. ACM, 2006.
- Caruana, Richard. Multitask learning: A knowledge-based source of inductive bias. In *Proceedings of the Tenth International Conference on Machine Learning*, pp. 41–48. Morgan Kaufmann, 1993.
- Chapelle, O., Schlkopf, B., and Zien, A. *A Discussion of Semi-Supervised Learning and Transduction*, pp. 473–478. MIT Press, 2006a. ISBN 9780262255899.
- Chapelle, Olivier, Schlkopf, Bernhard, and Zien, Alexander. *Semi-supervised learning*. Adaptive computation and machine learning. MIT Press, Cambridge (Mass.), 2006b. ISBN 0-262-03358-5.
- Coates, Adam, Ng, Andrew Y, and Lee, Honglak. An analysis of single-layer networks in unsupervised feature learning. In *International conference on artificial intelligence and statistics*, pp. 215–223, 2011.
- Corduneanu, Adrian and Jaakkola, Tommi. Data dependent regularization. In Chapelle, Olivier, Schlkopf, Bernhard, and Zien, Alexander (eds.), *Semi-supervised learning*. MIT Press, 2006.
- Dudley, Richard M. The sizes of compact subsets of hilbert space and continuity of gaussian processes. *Journal of Functional Analysis*, 1(3):290–330, 1967.
- Erhan, Dumitru, Bengio, Yoshua, Courville, Aaron, Manzagol, Pierre-Antoine, Vincent, Pascal, and Bengio, Samy. Why does unsupervised pre-training help deep learning? *The Journal of Machine Learning Research*, 11:625–660, 2010.
- Gamerman, Alexander, Vovk, Volodya, and Vapnik, Vladimir. Learning by transduction. In *Proceedings of the Fourteenth conference on Uncertainty in artificial intelligence*, pp. 148–155. Morgan Kaufmann Publishers Inc., 1998.
- Geman, Stuart, Bienenstock, Elie, and Doursat, René. Neural networks and the bias/variance dilemma. *Neural computation*, 4(1):1–58, 1992.
- Goodfellow, Ian, Warde-farley, David, Mirza, Mehdi, Courville, Aaron, and Bengio, Yoshua. Maxout networks. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pp. 1319–1327, 2013.
- Graham, Benjamin. Spatially-sparse convolutional neural networks. *CoRR*, abs/1409.6070, 2014.
- He, Kaiming, Zhang, Xiangyu, Ren, Shaoqing, and Sun, Jian. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. *arXiv preprint arXiv:1502.01852*, 2015.
- Hinton, Geoffrey, Vinyals, Oriol, and Dean, Jeff. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.

- Hinton, Geoffrey E, Osindero, Simon, and Teh, Yee-Whye. A fast learning algorithm for deep belief nets. *Neural computation*, 18(7):1527–1554, 2006.
- Koltchinskii, Vladimir. Rademacher penalties and structural risk minimization. *Information Theory, IEEE Transactions on*, 47(5):1902–1914, 2001.
- Koltchinskii, Vladimir and Panchenko, Dmitriy. Rademacher processes and bounding the risk of function learning. In *High dimensional probability II*, pp. 443–457. Springer, 2000.
- Krizhevsky, Alex. Learning multiple layers of features from tiny images, 2009.
- LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W., and Jackel, L. D. Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1(4):541–551, Winter 1989.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, November 1998.
- LeCun, Yann, Chopra, Sumit, Hadsell, Raia, Ranzato, Marc’Aurelio, and Huang, Fu-Jie. A tutorial on energy-based learning. In Bakir, G., Hofman, T., Schölkopf, B., Smola, A., and Taskar, B. (eds.), *Predicting Structured Data*. MIT Press, 2006.
- Lee, Chen-Yu, Xie, Saining, Gallagher, Patrick, Zhang, Zhengyou, and Tu, Zhuowen. Deeply-supervised nets. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, pp. 562–570, 2015.
- Lin, Min, Chen, Qiang, and Yan, Shuicheng. Network in network. *CoRR*, abs/1312.4400, 2013.
- Massart, Pascal. Some applications of concentration inequalities to statistics. In *Annales de la Faculté des sciences de Toulouse: Mathématiques*, volume 9, pp. 245–303, 2000.
- Mozer, Michael C and Smolensky, Paul. Skeletonization: A technique for trimming the fat from a network via relevance assessment. In *Advances in neural information processing systems*, pp. 107–115, 1989.
- Nair, Vinod and Hinton, Geoffrey E. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pp. 807–814, 2010.
- Polyak, B.T. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4(5):1 – 17, 1964. ISSN 0041-5553.
- Ranzato, Marc’Aurelio, Poultney, Christopher, Chopra, Sumit, and LeCun, Yann. Efficient learning of sparse representations with an energy-based model. In et al., J. Platt (ed.), *Advances in Neural Information Processing Systems (NIPS 2006)*, volume 19. MIT Press, 2006.
- Sermanet, Pierre, Eigen, David, Zhang, Xiang, Mathieu, Michaël, Fergus, Rob, and LeCun, Yann. Overfeat: Integrated recognition, localization and detection using convolutional networks. *CoRR*, abs/1312.6229, 2013.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014.
- Srivastava, Nitish, Hinton, Geoffrey, Krizhevsky, Alex, Sutskever, Ilya, and Salakhutdinov, Ruslan. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15 (1):1929–1958, 2014.
- Sudderth, Steven C and Holden, Alistair DC. Symbolic-neural systems and the use of hints for developing complex systems. *International Journal of Man-Machine Studies*, 35(3):291–311, 1991.
- Sutskever, Ilya, Martens, James, Dahl, George, and Hinton, Geoffrey. On the importance of initialization and momentum in deep learning. In *Proceedings of the 30th international conference on machine learning (ICML-13)*, pp. 1139–1147, 2013.
- Thrun, Sebastian and Pratt, Lorien (eds.). *Learning to Learn*. Kluwer Academic Publishers, Norwell, MA, USA, 1998. ISBN 0-7923-8047-9.
- Torralba, Antonio, Fergus, Rob, and Freeman, William T. 80 million tiny images: A large data set for nonparametric object and scene recognition. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 30 (11):1958–1970, 2008.
- Valiant, Leslie G. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.

Vapnik, Vladimir N. *The Nature of Statistical Learning Theory*. Springer-Verlag New York, Inc., New York, NY, USA, 1995. ISBN 0-387-94559-8.

Vapnik, Vladimir N. *Statistical Learning Theory*. Wiley-Interscience, 1998.

Vapnik, Vladimir N. *Estimation of Dependences Based on Empirical Data*. Springer New York, New York, NY, USA, 2006. ISBN 978-0-387-30865-4.

Vapnik, Vladimir N and Chervonenkis, A Ya. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability & Its Applications*, 16(2):264–280, 1971.

Wan, Li, Zeiler, Matthew, Zhang, Sixin, LeCun, Yann, and Fergus, Rob. Regularization of neural networks using dropconnect. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pp. 1058–1066, 2013.

Weston, Jason, Collobert, Ronan, Sinz, Fabian, Bottou, Léon, and Vapnik, Vladimir. Inference with the universum. In *Proceedings of the 23rd international conference on Machine learning*, pp. 1009–1016, 2006.

Zeiler, Matthew D. and Fergus, Rob. Stochastic pooling for regularization of deep convolutional neural networks. *CoRR*, abs/1301.3557, 2013.

Zhang, Xiang. Pac-learning for energy-based models. Master’s thesis, Computer Science Department, Courant Institute of Mathematical Sciences, New York University, 2013.

Zhu, Xiaojin and Goldberg, Andrew B. Introduction to semi-supervised learning. *Synthesis lectures on artificial intelligence and machine learning*, 3(1):1–130, 2009.

APPENDIX: PROOF OF THEOREM 2

Lemma 1 (Separation of dataset on empirical Rademacher complexity). *Let S be a dataset of size m . If S_1 and S_2 are two non-overlap subset of S such that $|S_1| = m - i$, $|S_2| = i$ and $S_1 \cup S_2 = S$, then*

$$\hat{\mathfrak{R}}_S(\mathcal{F}) \geq \frac{m-i}{m} \hat{\mathfrak{R}}_{S_1}(\mathcal{F}) - \frac{i}{m} \hat{\mathfrak{R}}_{S_2}(\mathcal{F}) \quad (17)$$

Proof. Let $(x_j, y_j) \in S_1$ for $j = 1, 2, \dots, m - i$ and $(x_j, y_j) \in S_2$ for $j = m - i + 1, m - i + 2, \dots, m$. Denote \mathbf{N} as the discrete uniform distribution on $\{1, -1\}$. We can derive by the convexity of supremum and symmetry of \mathbf{N}

$$\begin{aligned} \hat{\mathfrak{R}}_{S_1}(\mathcal{F}) &= \mathbb{E}_{\eta \sim \mathbf{N}^{m-i}} \left[\sup_{f \in \mathcal{F}} \frac{1}{m-i} \sum_{j=1}^{m-i} \eta_j f(x_j) \right] \\ &= \frac{2}{m-i} \mathbb{E}_{\eta \sim \mathbf{N}^{m-i}} \left[\sup_{f \in \mathcal{F}} \frac{1}{2} \sum_{j=1}^{m-i} \eta_j f(x_j) \right] \\ &= \frac{2}{m-i} \mathbb{E}_{\eta \sim \mathbf{N}^m} \left[\sup_{f \in \mathcal{F}} \left(\frac{1}{2} \sum_{j=1}^{m-i} \eta_j f(x_j) + \frac{1}{2} \sum_{j=m-i+1}^m \eta_j f(x_j) - \frac{1}{2} \sum_{j=m-i+1}^m \eta_j f(x_j) \right) \right] \\ &= \frac{2}{m-i} \mathbb{E}_{\eta \sim \mathbf{N}^m} \left[\sup_{f \in \mathcal{F}} \left(\frac{1}{2} \sum_{j=1}^m \eta_j f(x_j) - \frac{1}{2} \sum_{j=m-i+1}^m \eta_j f(x_j) \right) \right] \\ &\leq \frac{2}{m-i} \mathbb{E}_{\eta \sim \mathbf{N}^m} \left[\frac{1}{2} \sup_{f \in \mathcal{F}} \left(\sum_{j=1}^m \eta_j f(x_j) \right) + \frac{1}{2} \sup_{f \in \mathcal{F}} \left(\sum_{j=m-i+1}^m -\eta_j f(x_j) \right) \right] \\ &= \frac{m}{m-i} \mathbb{E}_{\eta \sim \mathbf{N}^m} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \sum_{j=1}^m \eta_j f(x_j) \right] + \frac{i}{m-i} \mathbb{E}_{\eta \sim \mathbf{N}^i} \left[\sup_{f \in \mathcal{F}} \frac{1}{i} \sum_{j=m-i+1}^m -\eta_j f(x_j) \right] \\ &= \frac{m}{m-i} \mathbb{E}_{\eta \sim \mathbf{N}^m} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \sum_{j=1}^m \eta_j f(x_j) \right] + \frac{i}{m-i} \mathbb{E}_{\eta \sim \mathbf{N}^i} \left[\sup_{f \in \mathcal{F}} \frac{1}{i} \sum_{j=m-i+1}^m \eta_j f(x_j) \right] \\ &= \frac{m}{m-i} \hat{\mathfrak{R}}_S(\mathcal{F}) + \frac{i}{m-i} \hat{\mathfrak{R}}_{S_2}(\mathcal{F}). \end{aligned}$$

The lemma is obtained by directly transforming the inequality above. \square

For any function space \mathcal{F} and distribution \mathbf{D} , denote $\mathfrak{R}_0(\mathcal{F}, \mathbf{D}) = 0$ and $\hat{\mathfrak{R}}_\emptyset(\mathcal{F}) = 0$. Define $P(m, i) = \sum_{i=0}^m \binom{m}{i} (1-p)^{m-i} p^i \frac{i}{m}$ and $Q(m, i) = \sum_{i=0}^m \binom{m}{i} (1-p)^i p^{m-i} \frac{i}{m}$. By definition of Rademacher complexity and lemma 1, we get

$$\begin{aligned}
\mathfrak{R}_m(\mathcal{F}, (1-p)\mathbf{D} + p\mathbf{U}) &= \mathbb{E}_{S \sim ((1-p)\mathbf{D} + p\mathbf{U})^m} [\hat{\mathfrak{R}}_S(\mathcal{F})] \\
&= \sum_{i=0}^m \binom{m}{i} (1-p)^{m-i} p^i \mathbb{E}_{S_1 \sim \mathbf{D}^{m-i}} \left[\mathbb{E}_{S_2 \sim \mathbf{U}^i} [\hat{\mathfrak{R}}_{S_1 \cup S_2}(\mathcal{F})] \right] \\
&\geq \sum_{i=0}^m \binom{m}{i} (1-p)^{m-i} p^i \mathbb{E}_{S_1 \sim \mathbf{D}^{m-i}} \left[\mathbb{E}_{S_2 \sim \mathbf{U}^i} \left[\frac{m-i}{m} \hat{\mathfrak{R}}_{S_1}(\mathcal{F}) - \frac{i}{m} \hat{\mathfrak{R}}_{S_2}(\mathcal{F}) \right] \right] \\
&= \sum_{i=0}^m \binom{m}{i} (1-p)^{m-i} p^i \left[\frac{m-i}{m} \mathbb{E}_{S_1 \sim \mathbf{D}^{m-i}} [\hat{\mathfrak{R}}_{S_1}(\mathcal{F})] - \frac{i}{m} \mathbb{E}_{S_2 \sim \mathbf{U}^i} [\hat{\mathfrak{R}}_{S_2}(\mathcal{F})] \right] \\
&= \sum_{i=0}^m \binom{m}{i} (1-p)^{m-i} p^i \left[\frac{m-i}{m} \mathfrak{R}_{m-i}(\mathcal{F}, \mathbf{D}) - \frac{i}{m} \mathfrak{R}_i(\mathcal{F}, \mathbf{U}) \right] \\
&= \left[\sum_{i=0}^m Q(m, m-i) \mathfrak{R}_{m-i}(\mathcal{F}, \mathbf{D}) \right] - \left[\sum_{i=0}^m P(m, i) \mathfrak{R}_i(\mathcal{F}, \mathbf{U}) \right] \\
&= \left[\sum_{i=1}^m Q(m, i) \mathfrak{R}_i(\mathcal{F}, \mathbf{D}) \right] - \left[\sum_{i=1}^m P(m, i) \mathfrak{R}_i(\mathcal{F}, \mathbf{U}) \right].
\end{aligned}$$

Theorem 2 is therefore established. The fact that $\sum_{i=0}^m Q(m, i) \leq 1$ and $\sum_{i=0}^m P(m, i) \leq 1$ is evident from the proof by each term being less than a term of the binomial expansion.