Universum Prescription: Regularization Using Unlabeled Data

Xiang Zhang, Yann LeCun

Courant Institute of Mathematical Sciences, New York University 719 Broadway, 12th Floor, New York, NY 10003 {xiang, yann}@cs.nyu.edu

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, STL-10 and ImageNet 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.

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 regularization effect on supervised task and reduces the generalization gap, the expected difference between test and training errors. 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), STL-10 (Coates, Ng, and Lee 2011), and ImageNet (Russakovsky et al. 2015). The method is theoretically justified using Radamacher complexity (Bartlett and Mendelson 2003).

Here we briefly describe our three universum prescription methods. Uniform prescription forces a discrete uniform distribution for unlabeled samples. Dustbin class simply adds an extra class and prescribe all unlabeled data to this class. Background class also adds an extra class, but it uses a constant threshold to avoid parameterization.

Our work is a direct extension to learning in the presence of universum (Weston et al. 2006) (Chapelle et al. 2007),

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. We extended the idea of using universum from support vector machines to deep learning.

Most deep learning approaches utilizing unlabeled data belong to the scope of representation learning (reviewed by (Bengio, Courville, and Vincent 2013) and (Bengio and Le-Cun 2007)) and transfer learning (Thrun and Pratt 1998). They include ideas like pretraining (Erhan et al. 2010) (Hinton, Osindero, and Teh 2006) (Ranzato et al. 2006) and semisupervised training (Rasmus et al. 2015) (Zhao et al. 2015). Universum prescription incoporates unlabeled data without imposing priors such as sparsity or reconstruction.

Regularization – techniques for the control of generalization gap – has been studied extensively. Most approaches implement a secondary optimization objective, such as an L_2 norm. Some other methods such as dropout (Srivastava et al. 2014) cheaply simulate model averaging to control the model variance. As part of general statistical learning theory (Vapnik 1995), (Vapnik 1998), the justification for regularization is well-developed. We qualitatively justify the methods using Radamacher complexity (Bartlett and Mendelson 2003), similar to (Wan et al. 2013).

Universum Prescription

In this section we attempt to formalize the trick of prescribing "none of the above" labels – universum prescription. Consider the problem of exclusive k-way classification. 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:

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

2. (Universum assumption) The proportion of unlabeled samples belonging to a supervised class is negligible.

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

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

Copyright © 2017, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

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, in the sense that one can change the type of loss and use the same principles to derive different universum learning techniques.

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 U, the assumption states that

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

The universum assumption is opposite to the assumptions of information regularization (Corduneanu and Jaakkola 2006) and transduction learning (Chapelle, Schlkopf, and Zien 2006) (Gammerman, Vovk, and Vapnik 1998). It has similarities with (Zhang and Zhou 2010) that encourages diversity of outputs for ensemble methods. All our 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.

Uniform Prescription

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

$$L(h, x, y) = -\sum_{i=1}^{k} Q[Y = i|x] \log \Pr[Y = i|x, h]$$
 (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. If h is parameterized as a deep neural network, uniform output is achieved when these parameters are constantly zero. Therefore, uniform prescription may have the effect of reducing the magnitude of parameters, similar to norm-based regularization.

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

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

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

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

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

and an unlabeled sample

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

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

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

and an unlabeled sample

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

Table 1: The 21-layer network

DESCRIPTION

LAYERS

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

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.

Theoretical Justification

In this part, we derive a qualitative justification for universum prescription using probably approximately correct (PAC) learning (Valiant 1984). By being "qualitative", the justification is in contrast with numerical bounds such as Vapnik-Chervonenkis dimension (Vapnik and Chervonenkis 1971) (VC-dim) and others. Our theory is based on Rademacher complexity (Bartlett and Mendelson 2003), similar to (Wan et al. 2013) where both dropout (Srivastava et al. 2014) and dropconnect (Wan et al. 2013) are justified. VC-dim is an upper-bound of Rademacher complexity, suggesting that the latter is more accurate. Previous results on unlabeled data (Oneto et al. 2011) (Oneto et al. 2015) assume the same distribution for labeled and unlabeled data, which is impossible under the universum assumption.

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

$$\hat{\mathfrak{R}}_{S}(\mathcal{F}) = \mathop{\mathrm{E}}_{\eta} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^{m} \eta_{i} f(x_{i}) \right]$$
(10)

where $\boldsymbol{\eta} = (\eta_1, \dots, \eta_m)^T$, with η_i 's independent random variables uniformly distributed on $\{-1, 1\}$.

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

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

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

$$\hat{\mathfrak{R}}_{S}(\mathcal{F}) = \mathop{\mathrm{E}}_{\boldsymbol{\eta}} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^{m} \eta_{i} f(x_{i}) \right]$$

$$\geq \sup_{f \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^{m} \mathop{\mathrm{E}}_{\eta_{i}}[\eta_{i}] f(x_{i}) = 0.$$
(12)

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

$$\hat{\mathfrak{R}}_{S}(\mathcal{F}) = \mathop{\mathrm{E}}_{\boldsymbol{\eta}} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^{m} \eta_{i} f(x_{i}) \right]$$
$$= \mathop{\mathrm{E}}_{\boldsymbol{\eta}} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^{m} \eta_{i} f_{0}(x) \right]$$
$$= \frac{1}{m} \sum_{i=1}^{m} \mathop{\mathrm{E}}_{\boldsymbol{\eta}_{i}} [\eta_{i}] f_{0}(x) = 0.$$
(13)

Therefore, one way to 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 is a bound for functions that are finite in outputs. We use the formulation by (Zhang 2013), but anterior results exist (Bartlett, Boucheron, and Lugosi 2002) (Bartlett and Mendelson 2003) (Koltchinskii 2001) (Koltchinskii and Panchenko 2000).

Theorem 1 (Approximation bound with finite bound on output). For an energy function (LeCun et al. 2006) $\mathcal{E}(h, x, y)$ over hypothesis class \mathcal{H} , input set \mathcal{X} and output set \mathcal{Y} , if it has lower bound 0 and upper bound M > 0, then with probability at least $1 - \delta$, the following holds for all $h \in \mathcal{H}$:

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

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

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

D is the distribution for (x, y), and S is a sample set of size m drawn indentically and independently from **D**.

The meaning of the theorem is two-fold. 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 Rademacher complexity is written as $\Re_m(\mathcal{F}, (1-p)\mathbf{D} + p\mathbf{U})$, which is reduced when we prescribe constant outputs to unlabeled data.

The second fold is that when the theorem applies to the supervised distribution **D**, we would hope that $\Re_n(\mathcal{F}, \mathbf{D})$ can be bounded by $\Re_m(\mathcal{F}, (1-p)\mathbf{D} + p\mathbf{U})$, where *n* is the number of supervised samples randomly chosen by the joint problem. Note that the number *n* follows a binomial distribution with mean (1-p)m. Such a bound can be achieved in a probable and approximate sense.

Theorem 2 (Rademacher complexity bound on distribution mixture). Assume we have a joint problem where $p \le 0.5$ and there are m random training samples from the joint distribution $(1-p)\mathbf{D} + p\mathbf{U}$. With probability at least $1-\delta$, the following holds

$$\mathfrak{R}_{n}(\mathcal{F}, \mathbf{D}) \leq \frac{2-p}{(1-p)\left(1-p-\sqrt{\frac{\log(1/\delta)}{2m}}\right)} \mathfrak{R}_{m}(\mathcal{F}, (1-p)\mathbf{D}+p\mathbf{U}),$$
(16)

where n is a random number indicating the number of supervised samples in the total joint samples, and m is large enough such that

$$1 - p - \sqrt{\frac{\log(1/\delta)}{2m}} > 0.$$
 (17)

We present the proof of theorem 2 in the supplemental material, which utilizes Hoeffding's inequality (Hoeffding 1963) (Serfling 1974). The theorem tells us that the Rademacher complexity of the supervised problem can be bounded by that of the joint problem. The universum prescription algorithm attempts to make the Rademacher complexity of the joint problem small. Therefore, universum prescription improves generalization by incorporating unlabeled data.

However, theorem 2 has a requirement that $p \leq 0.5$, otherwise the bound is not achievable. Also, the value of $(2-p)/(1-p)^2$ – the asymptotic constant factor in inequality 16 when m is large – is monitonally increasing with respect to p with a range of [2, 6] when $p \leq 0.5$. These facts indicate that we need to keep p small. The following sections show that there is improvement if p is small, but training and testing errors became worse when p is large.

LAYERS DESCRIPTION

1-3	Conv 128x3x3
4	Pool 2x2
5-7	Conv 256x3x3
8	Pool 2x2
9-11	Conv 512x3x3
12	Pool 2x2
13-15	Conv 1024x3x3
16	Pool 2x2
17-19	Conv 2048x3x3
20	Pool 2x2
21-22	Full 4096

Finally, in terms of numerical asymptotics, theorem 2 suggests that $\Re_n(\mathcal{F}, \mathbf{D}) \leq \mathbf{O}(1/\sqrt{m})$, instead of the commonly known result $\Re_n(\mathcal{F}, \mathbf{D}) \leq \mathbf{O}(1/\sqrt{n})$. This bounds the supervised problem with a tighter asymptotical factor because there are more joint samples than supervised samples.

Experiments on Image Classification

In this section we test the methods on some image classification tasks. Three series of datasets – CIFAR-10/100 (Krizhevsky 2009), STL-10 (Coates, Ng, and Lee 2011) and ImageNet (Russakovsky et al. 2015) – are chosen due to the availability of unlabeled data. For CIFAR-10/100 and STL-10 datasets, we used a 21-layer convolutional network (ConvNet) (LeCun et al. 1989) (LeCun et al. 1998), in which the inputs are 32-by-32 images and all convolutional layers are 3-by-3 and fully padded. For ImageNet, the model is a 17-layer ConvNet with 64-by-64 images as inputs. These models are inspired by (Simonyan and Zisserman 2014), in which all pooling layers are max-pooling, and ReLUs (Nair and Hinton 2010) are used as the non-linearity. Two dropout (Srivastava et al. 2014) layers of probability 0.5 are inserted before the final two linear layers.

The algorithm used is stochastic 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 for CIFAR-10/100

Table 3: Result for baseline and uniform prescription. The numbers are percentages.

DATASET	BASELINE			UNIFORM		
	Train	Test	Gap	Train	Test	Gap
CIFAR-10	0.00	7.02	7.02	0.72	7.59	6.87
CIFAR-100 F.	0.09	37.58	37.49	4.91	36.23	31.32
CIFAR-100 C.	0.04	22.74	22.70	0.67	23.42	22.45
STL-10	0.00	31.16	31.16	2.02	36.54	34.52
STL-10 Tiny	0.00	31.16	31.16	0.62	30.15	29.47
ImageNet-1	10.19	34.39	24.20	13.84	34.61	20.77
ImageNet-5	1.62	13.68	12.06	3.02	13.70	10.68

Table 4: Result for dustbin class and background class. Continuation of table 3

DATASET	DUSTBIN			BACKGROUND		
	Train	Test	Gap	Train	Test	Gap
CIFAR-10	0.07	6.66	6.59	1.35	8.38	7.03
CIFAR-100 F.	2.52	32.84	30.32	8.56	40.57	42.01
CIFAR-100 C.	0.40	20.45	20.05	3.73	24.97	21.24
STL-10	3.03	36.58	33.55	14.89	38.95	24.06
STL-10 Tiny	0.00	27.96	27.96	0.11	30.38	30.27
ImageNet-1	13.80	33.67	19.87	13.43	34.69	21.26
ImageNet-5	2.83	13.35	10.52	2.74	13.84	11.10

and every 600,000 minibatch steps for ImageNet. The training stops at 400,000 steps for CIFAR-10/100 and STL10, and 2,500,000 steps for ImageNet. Table 1 and 2 summarize the configurations. The weights are initialized in the same way as (He et al. 2015). The following data augmentation steps are used.

- 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 (or 64-by-64 for ImageNet) region in the scaled image.

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, Fergus, and Freeman 2008). Each dataset contains 60,000 samples, consitituting 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 3 and 4 contain the results. The three numbers in each tabular indicate training error, testing error and generalization gap. Bold numbers are the best ones for each case. The generalization gap is approximated by the difference between testing and training errors. All the 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 5. It shows that our network is competitive to the state of the art. Although (Graham 2014) has the best results, we believe that by applying out universum prescription methods to their model design could also improve the results further.

Table 5: 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.

STL-10

The STL-10 dataset (Coates, Ng, and Lee 2011) has size 96-by-96 for each image. We downsampled them to 32-by-32 in order to use the same model. The dataset contains a very small number of training samples – 5000. The accompanying unlabeled data contain 100,000 samples. There is no guarantee that these unlabeled samples do not blong to the supervised classes (Coates, Ng, and Lee 2011), therefore universum prescription failed. To verify that the extra data is the problem, an experiment using the 80 million tiny images as the unlabeled dataset is shown in table 3 and 4. In this case the improvement is observed. Due to long training times of our models, we did not perform 10-fold training in the original paper (Coates, Ng, and Lee 2011).

One interesting observation is that the results on STL-10 became better with the use of 80 million tiny images instead of the original extra data. It indicates that dataset size and whether universum assumption is satisfied are affecting factors for the effectiveness of universum prescription.

ImageNet

The ImageNet dataset (Russakovsky et al. 2015) for classification task has in total 1,281,167 training images and 50,000 validation images. The reported testing errors are evaluated on this validation dataset. During training, we resize images to minimum dimension 64, and then feed a random 64-by-64 crop to the network. Same test-time augmentation technique as in (Szegedy et al. 2015) are applied, with size variants {64, 72, 80, 88}, where each image is viewed in 144 crops.

The extra data comes from the large ImageNet 2011 release¹, for which we only keep the classes whomself and whose children do not belong to the supervised classes. This is enabled by the super-subordinate (is-a) relation information provided with the WordNet distribution (Miller 1995) because all ImageNet classes are nouns of WordNet. Both top-1 and top-5 results are reported in tables 3 and 4.

In all experiments dustbin class provides best results. We believe that it is because the extra class is parameterized, which makes it adapt better on the unlabeled samples.

Table 6: ConvNet for the study of p

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

Effect of the Regularization Parameter

It is natural to ask how would the 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 6, which has 9 parameterized layers in total. 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 descrease 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. This confirms the assumptions and conclusions from theorem 2.

Comparing between CIFAR-10/100 and STL-10, one conclusion is that that the model variance is affected by the combined size of labeled and unlabeled datasets. The variance on training and testing errors are extremely small on CIFAR-10/100 datasets because the extra data we used is almost unlimited (in total 80 million), but on STL-10 the variance seems to be large with much smaller combined size of training and extra datasets. This suggests that using universum prescription with a large abundance of extra data could improve the stability of supervised learning algorithms.

Finally, the comparison between using and not using dropout does not show a difference. This suggests that the regularization effect of universum prescription alone is comparable to that of dropout.

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. One is that loss

¹http://www.image-net.org/releases



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

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. Theorem 2 suggests that asymptotically the generalization ability of the supervised problem could be bounded by the joint problem, which has more samples due to the addition of unlabeled data. Experiments are done using CIFAR-10, CIFAR-100, STL-10 and ImageNet datasets. The effect of the regularization parameter is also studied empirically.

These experiments show that all three universum prescrition methods provide certain improvement over the generalization gap, whereas dustbin class constantly performs the best because the parameterized extra class can adapt better to the unlabeled samples. Further conclusions include that additional unlabeled data can improve the variance of models during training, and that the results are comparable to data-agnostic regularization using dropout.

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

Acknowledgments

We gratefully acknowledge 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

Bartlett, P. L., and Mendelson, S. 2003. Rademacher and gaussian complexities: Risk bounds and structural results. *The Journal of Machine Learning Research* 3:463–482.

Bartlett, P. L.; Boucheron, S.; and Lugosi, G. 2002. Model selection and error estimation. *Machine Learning* 48(1-3):85–113.

Bengio, Y., and LeCun, Y. 2007. Scaling learning algorithms towards ai. In Bottou, L.; Chapelle, O.; DeCoste, D.; and Weston, J., eds., *Large-Scale Kernel Machines*. MIT Press.

Bengio, Y.; Courville, A.; and Vincent, P. 2013. Representation learning: A review and new perspectives. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 35(8):1798–1828.

Chapelle, O.; Agarwal, A.; Sinz, F. H.; and Schölkopf, B. 2007. An analysis of inference with the universum. In *Advances in neural information processing systems*, 1369–1376.

Chapelle, O.; Schlkopf, B.; and Zien, A. 2006. *A Discussion of Semi-Supervised Learning and Transduction*. MIT Press. 473–478.

Coates, A.; Ng, A. Y.; and Lee, H. 2011. An analysis of single-layer networks in unsupervised feature learning. In *International conference on artificial intelligence and statistics*, 215–223.

Corduneanu, A., and Jaakkola, T. 2006. Data dependent regularization. In Chapelle, O.; Schlkopf, B.; and Zien, A., eds., *Semi-supervised learning*. MIT Press.

Erhan, D.; Bengio, Y.; Courville, A.; Manzagol, P.-A.; Vincent, P.; and Bengio, S. 2010. Why does unsupervised pretraining help deep learning? *The Journal of Machine Learning Research* 11:625–660.

Gammerman, A.; Vovk, V.; and Vapnik, V. 1998. Learning by transduction. In *Proceedings of the Fourteenth conference on Uncertainty in artificial intelligence*, 148–155. Morgan Kaufmann.

Goodfellow, I.; Warde-farley, D.; Mirza, M.; Courville, A.; and Bengio, Y. 2013. Maxout networks. In *Proceedings* of the 30th International Conference on Machine Learning (ICML-13), 1319–1327.

Graham, B. 2014. Spatially-sparse convolutional neural networks. *CoRR* abs/1409.6070.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. *arXiv preprint arXiv:1502.01852*.

Hinton, G. E.; Osindero, S.; and Teh, Y.-W. 2006. A fast learning algorithm for deep belief nets. *Neural computation* 18(7):1527–1554.

Hoeffding, W. 1963. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association* 58(301):13–30.

Koltchinskii, V., and Panchenko, D. 2000. Rademacher processes and bounding the risk of function learning. In *High dimensional probability II*. Springer. 443–457.

Koltchinskii, V. 2001. Rademacher penalties and structural risk minimization. *Information Theory, IEEE Transactions on* 47(5):1902–1914.

Krizhevsky, A. 2009. Learning multiple layers of features from tiny images.

LeCun, Y.; Boser, B.; Denker, J. S.; Henderson, D.; Howard, R. E.; Hubbard, W.; and Jackel, L. D. 1989. Backpropaga-

tion applied to handwritten zip code recognition. *Neural Computation* 1(4):541–551.

LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86(11):2278–2324.

LeCun, Y.; Chopra, S.; Hadsell, R.; Ranzato, M.; and Huang, F. 2006. A tutorial on energy-based learning. In Bakir, G.; Hofman, T.; Schölkopf, B.; Smola, A.; and Taskar, B., eds., *Predicting Structured Data*. MIT Press.

Lee, C.-Y.; Xie, S.; Gallagher, P.; Zhang, Z.; and Tu, Z. 2015. Deeply-supervised nets. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, 562–570.

Lin, M.; Chen, Q.; and Yan, S. 2013. Network in network. *CoRR* abs/1312.4400.

Miller, G. A. 1995. Wordnet: a lexical database for english. *Communications of the ACM* 38(11):39–41.

Nair, V., and Hinton, G. E. 2010. Rectified linear units improve restricted boltzmann machines. In *Proceedings* of the 27th International Conference on Machine Learning (ICML-10), 807–814.

Oneto, L.; Anguita, D.; Ghio, A.; and Ridella, S. 2011. The impact of unlabeled patterns in rademacher complexity theory for kernel classifiers. In *Advances in Neural Information Processing Systems*, 585–593.

Oneto, L.; Ghio, A.; Ridella, S.; and Anguita, D. 2015. Local rademacher complexity: Sharper risk bounds with and without unlabeled samples. *Neural Networks* 65:115–125.

Polyak, B. 1964. Some methods of speeding up the convergence of iteration methods. USSR Computational Mathematics and Mathematical Physics 4(5):1 - 17.

Ranzato, M.; Poultney, C.; Chopra, S.; and LeCun, Y. 2006. Efficient learning of sparse representations with an energybased model. In et al., J. P., ed., *Advances in Neural Information Processing Systems (NIPS 2006)*, volume 19. MIT Press.

Rasmus, A.; Berglund, M.; Honkala, M.; Valpola, H.; and Raiko, T. 2015. Semi-supervised learning with ladder networks. In Cortes, C.; Lawrence, N. D.; Lee, D. D.; Sugiyama, M.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 28*. Curran Associates, Inc. 3546–3554.

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. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)* 115(3):211–252.

Serfling, R. J. 1974. Probability inequalities for the sum in sampling without replacement. *The Annals of Statistics* 2(1):39–48.

Sermanet, P.; Eigen, D.; Zhang, X.; Mathieu, M.; Fergus, R.; and LeCun, Y. 2013. Overfeat: Integrated recognition, localization and detection using convolutional networks. *CoRR* abs/1312.6229. Simonyan, K., and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *CoRR* abs/1409.1556.

Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research* 15(1):1929–1958.

Sutskever, I.; Martens, J.; Dahl, G.; and Hinton, G. 2013. On the importance of initialization and momentum in deep learning. In *Proceedings of the 30th international conference on machine learning (ICML-13)*, 1139–1147.

Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; and Rabinovich, A. 2015. Going deeper with convolutions. In *Computer Vision and Pattern Recognition (CVPR)*.

Thrun, S., and Pratt, L., eds. 1998. *Learning to Learn*. Norwell, MA, USA: Kluwer Academic Publishers.

Torralba, A.; Fergus, R.; and Freeman, W. T. 2008. 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.

Valiant, L. G. 1984. A theory of the learnable. *Communications of the ACM* 27(11):1134–1142.

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

Vapnik, V. N. 1995. *The Nature of Statistical Learning Theory*. New York, NY, USA: Springer-Verlag New York, Inc.

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

Vapnik, V. N. 2006. Estimation of Dependences Based on Empirical Data. New York, NY, USA: Springer New York.

Wan, L.; Zeiler, M.; Zhang, S.; LeCun, Y.; and Fergus, R. 2013. Regularization of neural networks using dropconnect. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, 1058–1066.

Weston, J.; Collobert, R.; Sinz, F.; Bottou, L.; and Vapnik, V. 2006. Inference with the universum. In *Proceedings of the 23rd international conference on Machine learning*, 1009–1016.

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

Zhang, M.-L., and Zhou, Z.-H. 2010. Exploiting unlabeled data to enhance ensemble diversity. In *2010 IEEE International Conference on Data Mining*, 619–628. IEEE.

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

Zhao, J.; Mathieu, M.; Goroshin, R.; and LeCun, Y. 2015. Stacked what-where auto-encoders. *arXiv preprint arXiv:1506.02351*.