Loss Function Entropy Regularization for Diverse Decision Boundaries

Sue Sin Chong

Institute for Interdisciplinary Information Sciences, Tsinghua University
z18516937829@163.com

Abstract—Is it possible to train several classifiers to perform meaningful crowd-sourcing to produce a better prediction label set without any ground-truth annotation? In this paper, we will attempt to modify the contrastive learning objectives to automatically train a self-complementing ensemble to produce a state-of-the-art prediction on the CIFAR10 and CIFAR100-20 task. This paper will present a remarkably simple method to modify a single unsupervised classification pipeline to automatically generate an ensemble of neural networks with varied decision boundaries to learn a larger feature set of classes. Loss Function Entropy Regularization (LFER), are regularization terms to be added upon the pre-training and contrastive learning objective functions, gives us a gear to modify the entropy state of the output space of unsupervised learning, thereby diversifying the latent representation of decision boundaries of neural networks. Ensemble trained with LFER have higher successful prediction accuracy for samples near decision boundaries. LFER is a effective gear to perturb decision boundaries, and has proven to be able to produce classifiers that beat state-of-the-art at contrastive learning stage. Experiments show that LFER can produce an ensemble where each have accuracy comparable to the state-of-the-art, yet have each have varied latent decision boundaries. It allows us to essence perform meaningful verification for samples near decision boundaries, encouraging correct classification of near-boundary samples. By compounding the probability of correct prediction of a single sample amongst an ensemble of neural network trained, our method is able to improve upon a single classifier by denoising and affirming correct feature mappings.

Index Terms—Unsupervised Learning, Diversified Feature Set, Entropy Regularization, Decision Boundaries, Loss Function

I. INTRODUCTION

How to make unsupervised algorithms to converge and learn different features at will at the same time has always been difficult. Recently, a combination of representation learning method and end-to-end learning has given promising results on the coarse classification of reasonably large datasets such as CIFAR100 [1]. [1]'s three-step process involves 2 contrastive learning stages and a fine-tuning stage.

Recently, [2] works on feature selection in unsupervised learning. [3], [4] works on making the pretext training phase of unsupervised learning more robust. The properties of similarity search vectors and application in unsupervised image classification task, has also been much studied in [5], [6] . The integration of noise and feature into the unsupervised learning pipeline has also proven to be useful [7]. [8] [9]

has studied the possibility of searching in a quantized sparse space representation. [10] has studied momentum of the unsupervised clustering process. Most recently, [11] introduced a simple contrastive learning framework. Whereas [12] makes the contrastive learning stage robust. And [13] is a fine-tuning technique that has proven to be very effective in the final stages of unsupervised learning. Whereas [14] attempts to find new objects or features as an optimization objective.

Entropy regularization is an important technique in machine learning, and has many applications. [15], [16], [17] all utilized entropy regularization enhance unsupervised learning, entropy regularization can either maximize marginal entropy of bits or speed-up classification. In searching, [18] also utilized entropy regularization on one-hot codes. Whereas [19] has studied the mass-spring-damper system without a singular kernel, which might be able to play the role of controller in unsupervised learning.

A. Related Work

Out of Distribution (OOD) Detection is an increasing prominent field in machine learning. [20] is a method for generating state-of-the-art OOD detector with adversarial method. [21] used a method based on Gaussian to extract features from a trained neural network to train an OOD detector with a relatively simple tools. Via [21], one is able to link the certainty of features learnt to OOD classifier accuracy.

B. Motivation

The ability to check for feature that has not been detected by an unsupervised learner but also an important feature of the dataset is important. It serves as a fill-in-the-blank check for an unsupervised classifier. It allows us to paint a class as a set of more specific and finer and at the same time correct features, instead of having to accept a class as a very generic template of features.

Secondly, the ability to separate weakly-classified samples from well-classified samples only by using similarly unsupervised-trained neural network or pipeline, allows the unsupervised learner to self-check, and improve. When it is hard to differentiate, it will be very helpful to be able to have neural networks trained with varied decision boundaries to vote on a particular sample. This motivates us to modify entropy configuration of output space to train neural networks which has varied decision boundaries to help partition the dataset into easy-to-classify and possibly-hard-to-classify samples.

II. Preliminaries

We attempt to introduce the possibility of searching in the space of decision boundaries in the unsupervised image classification task. This work attempts to combine the best of representation learning, end-to-end learning and the properties of spreading vector and entropy control to give users a simple yet intuitive way to explore and exploit along the entropy dimension in output space of contrastive learning.

In this paper, we present a remarkably simple method for training reasonably good ensembles from a single pipeline with diverse decision boundary and is able to learn varied latent representations of the dataset. We view the problem of improving unsupervised classification as the unsupervised and automatic exploration of decision boundaries under varied entropy configuration. We present a framework where we can generate an ensemble which have varied yet complementing decision boundaries simply by changing constants. Our objective is to maximize the number of possibilities in output space entropy distribution of converging neural networks, which allows for a large variety of decision boundaries and therefore a highly diversified set of learnt latent representation of class features.

To encourage entropy exploration and exploitation, we developed an approach of adding entropy regularization terms in objective functions in pretext and contrastive learning stage. LFER is a set of regularization terms to be added on unsupervised learning objective functions. We offer a simple-to-implement yet sensitive gear for contrastive learning, which allows for unsupervised learning which automatically exploits differences in latent representation of decision boundaries. We show that generating an ensemble from a arbitrary machine learning architecture and a dataset (pipeline) is simple, and bound have improved accuracy as compared to the best neural network possibly trained from the same pipeline.

The main contributions of our paper are as follows:

- We identified a method to exploit entropy in decision boundary formation in unsupervised classification problem, the LFER method.
- LFER is simple to implement and use, and almost always guarantee to train neural network with varied decision boundaries from the same dataset from a single architecture. Without LFER, unsupervised learning always converges on to the same set of feature with the same decision boundaries, reproducably.
- LFER method can produce networks that have classification accuracy which compares to state-of-the-art, but with varied latent decision boundaries. It can also serve as a uniqueness of convergence check on an unsupervised classification pipeline.
- Ensemble trained with LFER can meaningfully encourage prediction for samples which lies near decision boundaries.
- On CIFAR100-20 task, our ensemble is able to capture a larger feature set of super-classes.
- LFER can mine for neural networks which can be better trained as out-of-distribution detectors.

In the following section, we will discuss the definition and implications of solving the problem of diversified unsupervised learning with LFER.

III. Problem Definition

Given a unsupervised classification pipeline $Pipeline$, how to automatically train an ensemble which can improve prediction accuracy where each neural network in the ensemble independently a reasonable classifier.

A. Applications

LFER is a natural filter for weakly-classified samples in unsupervised learning. With a series of neural network which each converge to a different set of correct features, it can immediately identify weakly-classified samples. Secondly, LFER checks for convergence in features learnt in unsupervised classification pipelines. When the dataset has multiple ways of matching different features to the same set of classes, it is possible for LFER to list the different mappings between features and classes. It serves as a neighborhood exploration tool which allows for searching with constants on some arbitrary entropy structure in the output space.

B. Implications

When a $Pipeline$ (machine learning architecture + dataset) is sufficient to produce a reasonable classifier by unsupervised learning on a dataset. By using LFER, we can exhaust the feature discovery possibilities of the target $Pipeline$. If there exist any subclass within a super-class, the LFER will effectively mine for sub-class representations in the original dataset as there will very likely exist a neural network which converge to an alternative set of features, it serves as a checker for the uniqueness of feature convergence in a unsupervised classification process. This is particularly useful when the actual number of classes is unknown.

The first section will discuss the LFER method on Unsupervised Semantic Clustering pipeline. Whereas the second section will discuss reasoning about entropy regularization, and the convergence of neural network trained with LFER. The third section will discuss the implications of finding complementing neural networks which has learnt different features of the same dataset. The fourth section will discuss the application of combinations of neural network with different latent decision boundaries. This paper will conclude with applications and implications of using LFER as a output space entropy controlling tool.

IV. Loss Function Entropy Regularization (LFER)

Loss Function Entropy Regularization (LFER) is a series of entropy regularization terms to be added to contrastive learning optimization functions, as follows. LFER is to be added on the contrastive learning and pre-training stages of unsupervised classification to encourage different decision boundary and entropy distribution in the output space, thereby resulting neural networks with similar accuracy but different latent representation of features. Experiments show that training

neural networks with different confusion matrix is not possible without LFER. The implementation of LFER merely requires an additional tens of lines of code to the loss function.

$$
(-\lambda_0, -\lambda_1, +\lambda_2, -\lambda_3)
$$
$$
(\sum \Phi_\eta^{'c} \log \Phi_\eta^{'c}, \sum_{c\in\mathcal{C}} \Phi_\eta^{'c} \log \Phi_\eta^{'c},
$$
$$
\sum_{c\in\mathcal{C}} \sum_{k\in\mathcal{N}_X} \Phi_\eta^c \Phi_\eta^k \log \Phi_\eta^c \Phi_\eta^k,
$$
$$
\sum_{c\in\mathcal{C}} \sum_{k\in\mathcal{N}_X} \Phi_\eta^{'c} \Phi_\eta^k \log \Phi_\eta^{'c} \Phi_\eta^k)
$$
$$
, \Phi_\eta^{'c} = \frac{1}{|\mathcal{D}|} \sum_{X\in\mathcal{D}} \Phi_\eta^c(X)) \tag{1}
$$

Entropy terms were previously added in semantic clustering to encourage uniform prediction amongst classes in the contrastive learning stage. Our experiment shows that regularization in unsupervised classification objective functions are indispensable for identification of larger set of features, ensemble trained with LFER demonstrates an improvement in latent representation of features of the neural network.

Higher order entropy terms play the role of controller in controlling the distances between clusters in the output space, hence it is possible to produce neural networks with all sorts of varying confusion matrices. It is possible to both maximize or minimize for between-cluster spaces and the smoothing of the clustering process.

A. LFER Acting Upon Unsupervised Semantic Clustering

There are 3 portions to unsupervised classification in the [1], SimCLR, Semantic Clustering(SCAN), and Selflabel(SLL).SCAN being contrastive learning stages. Adding entropy terms in the objective functions of contrastive learning stages can improve the diversity of neural network trained. Regularization term in SimCLR is as follows.

$$
\min_\theta d(\Phi_\theta(X_i), \Phi_\theta(T[X_i]))
$$
$$
- \lambda_0 \langle \Phi_\theta(X_i), \Phi_\theta(T[X_i]) \rangle \log \langle \Phi_\theta(X_i), \Phi_\theta(T[X_i]) \rangle \tag{2}
$$

Regularization term in SCAN portion of the pipeline is as follows. λ_2 and λ_3 plays the role of a spring term and damper term for between-clusters entropy.

$$
\Lambda = -\frac{1}{|\mathcal{D}|} \sum_{X\in\mathcal{D}} \sum_{k\in\mathcal{N}_X} \log \langle \Phi_\eta(X), \Phi_\eta(k) \rangle
$$
$$
+ \lambda_1 \sum \Phi_\eta^{'c} \log \Phi_\eta^{'c}
$$
$$
- \lambda_2 \sum_{c\in\mathcal{C}} \sum_{k\in\mathcal{N}_X} \Phi_\eta^c \Phi_\eta^k \log \Phi_\eta^c \Phi_\eta^k
$$
$$
+ \lambda_3 \sum_{c\in\mathcal{C}} \sum_{k\in\mathcal{N}_X} \Phi_\eta^{'c} \Phi_\eta^k \log \Phi_\eta^{'c} \Phi_\eta^k \tag{3}
$$
$$
, \Phi_\eta^{'c} = \frac{1}{|\mathcal{D}|} \sum_{X\in\mathcal{D}} \Phi_\eta^c(X)
$$

The scalar terms $(\lambda_i)_{i\geq 0} = f(classes)$, is a function of number of classes. Below is a table of scalar terms and the classification accuracy of different lambda combinations. Together, the new optimization function presents a control function related to number of classes to modify the entropy in the input space. Hence we can specify a specific behavior desired of output neural network, and then simply by modifying relative values of $(\lambda_i)_{i\geq 1}$ train a converging network with the desired property.

TABLE I: Experiment Results

Desc	λ_0	λ_1	λ_2	λ_3	Scan	SLL
SCAN	0	5	0	0	0.44	0.507
1	0	5	4	0	0.439	0.503
2	2	5	4	8	0.4335	0.5012
3	2	5	4	4	0.4005	0.4408
4	2	5	$4\sqrt{n}$	$-8/n$	0.456	0.4767

Reflecting sensitivity of spring constant and dampening ratio in a physical spring system, neural networks trained with objective functions with regularization terms that are scalar multiples of each other can also have significantly different decision boundaries. We can also choose to improve the performance of a single classifier in the contrastive learning stage.

1) Reasoning About Entropy State: By optimizing for entropy exploration, we are adding the constraint of number of classes to the output space, forcing different values and structures of classes in the trained neural network. The regularization of entropy throughout the pipeline, encourages us to introduce an abridged notation to reason about the state of entropy that the training has result in. It will allow for the convenience of solving for deducing about the stability of training process and the properties of the output neural network. The 3-step unsupervised classification process is essentially a cascading function of the entropy exploration and exploitation of the dataset.

TABLE II: Entropy Regularization State Notation

$(\lambda_i)_{i\geq 1}$	Value	Notation
λ_0	0	$g(x) = x$
λ_0	> 0	$g(x) = \hat{x}$

Let $g(x)$ be the pretext regularization function. Similarly, we define the following for SCAN.

$$
\Lambda = -\frac{1}{|\mathcal{D}|} \sum_{X\in\mathcal{D}} \sum_{k\in\mathcal{N}_X} \log \langle \Phi_\eta(X), \Phi_\eta(k) \rangle
$$
$$
+ \lambda_2 x - \lambda_3 x' + \lambda_4 x'' \tag{4}
$$

Let $\lambda_2 x - \lambda_3 x' + \lambda_4 x'' = h(x)$ be the SCAN regularization function. Then compounding the two regularization functions, the state of entropy as a result of the training process can be expressed as $h(g(x))$. Hence, the contrastive training process is maximizing similarity together with a second order differential equation of entropy, offering fine-grained knobs to the desired output entropy state, and the decision boundary formation within clusters. LFER serves as a control function for entropy in the contrastive learning stage.

B. Pseudo Entropy Control Functions

$(\lambda_i)_{i\geq 1}$ is a function of number of classes, aka $(\lambda_i)_{i\geq 1} = f_{i\geq 1}(n)$. Expressing the entropy control function as an inner

product between a series of functions of number of classes, we have the following.

$$
\begin{aligned}
h(g(x)) &= h(n, g(x)) \\
&= \lambda_1 g(x) - \lambda_2 g(x)' + \lambda_3 g(x)'' \\
&= f_1(n)g(x) - f_2(n)g(x)' + f_3(n)g(x)'' \\
&= \langle f_{i \geq 1}(n), g(x) \rangle
\end{aligned}
\tag{5}
$$

Self-labelling which attempts to correctly classify noisy neighbors near clusters by minimizing cross-entropy loss, is the only step in the pipeline which only optimizes for cross-entropy. LFER makes it possible to systematically produce converging neural networks with different decision boundaries simply by modifying $(f_i(n))_{n \geq 1}$. Experiments show that neural networks trained with different objective functions are most confident about a wide variety of different prototype images for each super-class. This offers a special edge to the task of CIFAR100-20, as it allows for the mining of more subclass prototypes within a super-class, thereby improving overall accuracy of classification.

C. Grid-Searching for Constants to Train the Best Performing Neural Network

In contrastive learning stages, LFER has consistently been able to mine for neural networks that beat state-of-the-art by $1-2\%$ at the end of contrastive learning stage on the CIFAR10 and CIFAR100-20 classification task. However, the slight edge gained doesn't persist through the self-labelling stage. LFER is also able to produce neural networks which in turn train better out-of-distribution classifiers on the same dataset.

V. ENTROPY PERTURBATION FOR DECISION BOUNDARIES

LFER is necessary to train neural networks to learn different features. Experiment shows that repeated training on the model without any regularization term, results in similar confusion matrix. The terms forces different structure on the neural network decision boundaries, which results in differences in confusion matrix.

LFER is a sensitive gear for modifying the entropy environment / configuration in output space. In this section, we will discuss an example ensemble, trained with $\lambda_3 = \{4, 8, 16, 32\}$. To demonstrate the fact that LFER is able to produce neural networks which learn different set of feature at ease, we introduce the notion of n guess accuracy.

The accuracy of N guesses is calculated as follows. Given n neural networks, if any of the n neural network predicts the label correctly, then it is calculated as a correct prediction. There is no hierarchy or preference to the set of prediction label produced by the set of n neural networks.

A. Self-complementing Ensemble

The λ_3 term is an effective knob for controlling the fine-grained cluster formation process. It is easy to train a set of neural networks simply by multiplying the λ_3 with a geometric series. By changing λ_3, we are almost certain to find a neural network which has learnt a different set of correct features.

TABLE III: 2 Guess Of a λ_3 Series

2 guess ACC	Agreement	λ_0	λ_1	λ_2	λ_3
63.79	44.90	2	5	4	4
		2	5	4	8
61.82	55.81	2	5	4	8
		2	5	4	16
61.8	50.59	2	5	4	16
		2	5	4	32
61.68	42.64	2	5	4	4
		2	5	4	32

The 2-ensemble agrees samples in which they are both certain, and only disagrees on samples where there are confusion or lie near the decision boundaries. With LFER, we are able to identify classes and samples where there are second opinions. Therefore it can serve as a filtering method for noisy samples and noisy classes.

TABLE IV: 3/4 Guess Of a λ_3 Series

3/4 guess ACC	2-agreement	λ_0	λ_1	λ_2	λ_3
69.35	75.86	2	5	4	4
		2	5	4	8
		2	5	4	32
67.80	83.97	2	5	4	4
		2	5	4	8
		2	5	4	16
67.49	82.50	2	5	4	4
		2	5	4	16
		2	5	4	32
71.92	NA	2	5	4	4
		2	5	4	8
		2	5	4	16
		2	5	4	32

We sieve for confident samples with majority votes from the ensemble, for at least three quarters of the samples within the ensemble. We can further clamp down on samples where all of the networks do not agree on, aka samples with very high confusion. Below is an image of the most confident prototypes of an LFER ensemble. This ensemble has learnt the many sub-classes within the 20 super-classes of the CIFAR100 dataset. For instance, in the large carnivores super class, neural networks in the ensemble has learnt bear, tiger, leopard and lion respectively as their most confident prototype for the same superclass.
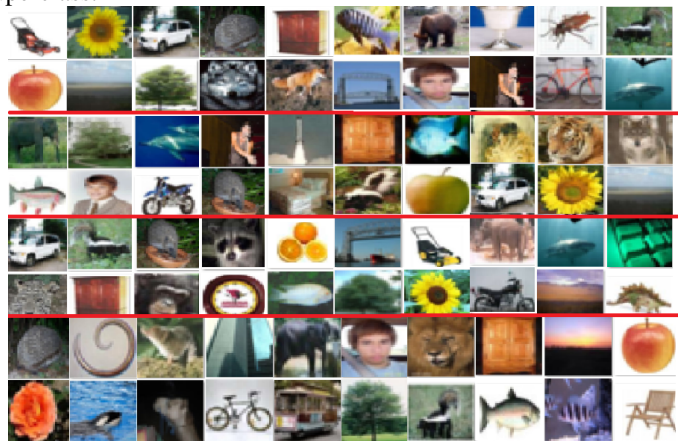

Fig. 1: **Figure 1. Different Prototype Images of Superclasses**

LFER can mine for a wider form of representation of features of classes, when there are many sub-classes within a super-class, and when sub-classes differ drastically, the advantage will be prominent. 4 neural networks is sufficient to achieve for verification check for weakly-classified samples.

VI. MORE ENTROPY CONFIGURATIONS

It is also possible to explore the entropy space further with even larger classes of LFER functions. Lambdas templates which result in converging neural networks correspond to spring damper constants which result in spring-damper systems with forced harmonic motion.

A. Extended Stability

It has been proven via experimentation that the following lambda bounds will produce optimization functions that can produce a converging neural network. The number of possible LFER functions which can train converging neural networks in the contrastive learning stage is very large.

TABLE V: Big O Bounds for $f(n)$

$(\lambda_i)_{i>1}$	Big O Bound
λ_0	$O(1)$
λ_1	$O(1)$
λ_2	$O(n)$
λ_3	$O(n\sqrt{(n)})$

TABLE VI: Stable Templates

λ_0	λ_1	λ_2	λ_3
0	$O(1)$	0	0
0	$O(1)$	$O(1)$	0
$O(1)$	$O(1)$	$O(1)$	0
$O(1)$	$O(1)$	$O(1)$	$O(1)$
$O(1)$	$O(1)$	$O(\sqrt{n})$	0
$O(1)$	$O(1)$	0	$O(n)$
$O(1)$	$O(1)$	$O(\sqrt{n})$	$O(\sqrt{n})$
$O(1)$	$O(1)$	$O(1)$	$O(\frac{1}{n})$
$O(1)$	$O(n)$	$O(\frac{1}{n})$	$O(n\sqrt{n})$

Decision boundaries of neural network trained is very sensitive to small changes in $(\lambda_i)_{i\geq1}$. It is reasonable to expect to find neural networks that can identify all other subclass prototypes in the super-class by enumerating the set of λs. With LFER, it is possible to have multiple neural networks which is confident about more sub-classes within the super-class to learn the sub-classes within a super-class.

We observed that accuracy of a single classifier in CIFAR20 can hardly surpass 0.5 due to both limit of number of classes and lack of labels. Having multiple networks learning the 20 classes in CIFAR100-20, makes it possible to have different neural networks learn the features of smaller classes, which result in a better superclass classification accuracy.

B. Combined Ensembles Trained with LFER

On the CIFAR100 dataset, we have trained 25 Res18Nets for classification into 20 coarse classes. Each of the 25 Res18Nets has been trained with a different lambda constant set. Results in the following table is obtained by combining 2 to 4 of the 25 neural networks.

TABLE VII: N-guess Performances

N Guess	Best ACC	Mean ACC	Median ACC
1 nn	50.7	NA	NA
2 guess	64.50	58.20	59.24
3 guess	70.57	63.70	64.70
4 guess	73.99	68.90	69.10

Each of the classifier is independently an at least > 0.40 accuracy classifier. And > 0.85 of the classifiers show at least 0.1 improvement in classification guess. When we compound LFER of different classes, it is much more likely to make the n guess of the neural networks more robust. Below is a listing of some of the best combinations of neural networks. Results show that some of the best combinations often involve neural networks with non-zero $\lambda_0, \lambda_2, \lambda_3$, demonstrating the indispensability of the entropy regularization terms in training neural networks to learn different features.

TABLE VIII: 2 Guess of Combined Ensembles

2 guess ACC		λ_0	λ_1	λ_2	λ_3
1	64.50	0	5	0	0
		2	5	4	0
	64.12	0	5	0	0
		2	5	$1/4n^2$	$1/2n$
	63.79	2	5	4	8
		2	5	4	4

TABLE IX: 3 Guess of Combined Ensembles

3 guess ACC		λ_0	λ_1	λ_2	λ_3
1	70.79	0	5	0	0
		2	5	4	0
		2	5	$1/4n$	$0.5/\sqrt{n}$
	70.66	0	5	0	0
		2	5	4	0
		2	5	4	4
	70.60	0	5	0	0
		2	5	$1/4n$	$0.5/\sqrt{n}$
		2	5	$1/4n^2$	$1/2n$

In the CIFAR100-20 task, there are 5 sub classes in each super-class, there are multiple ways of mapping sub-classes to a super class. LFER mines the input dataset for multiple sub-classes within the super-class. By compounding 2 to 3 neural networks whereby each network learns a subclass within a super-class, the n guess predictions will in expectation make good guesses which will include each subclass.
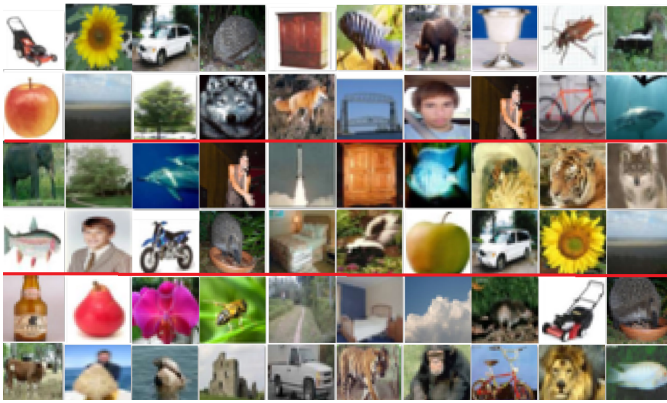
Fig. 2: **Figure 2. Prototype Images for 2 Guess Combined Ensembles**

Best 3 combinations often have confident prototypes drawn from a larger variety of sub-classes.
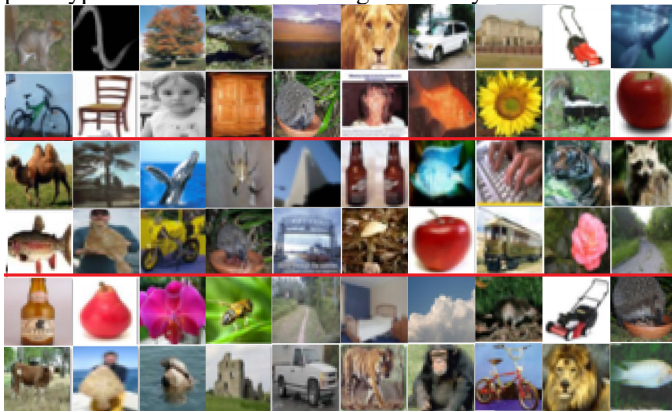

Fig. 3: **Figure 3. Prototype Images for 3 Guess Combined Ensembles**

While there may be confusion across super-classes, each subclass is only represented once in each classifier confident prototypes. Combinations in aggregation learn to be confident about a large set of sub-classes.

1) Learning Sub-classes Within Super-class: With the LFER, we are able to mine for different prototypical images within a super-class in the CIFAR100-20 task. Instead of being forced to view a super-class as a set non-descriptive generic feature, LFER is able to mine for decision boundaries which respect specific subclass features within a super-class.

VII. MAJORITY VOTE

By training several neural network, and have them voting on a classes of a sample can produce state-of-the-art prediction accuracy. N guess accuracy is the upper bound of majority vote accuracy. We are able to improve upon the accuracy of a single prediction set by compounding the accuracy of several classifiers. Our best combinations which involves 27 different neural networks voting for the label of every single sample, reaches a new state-of-the-art accuracy rate of 0.58.

This demonstrates that the fuzziness induced by LFER allows for meaningful verification check across neural network trained from the same architecture. We also observed that by dividing the classifiers into three tiers according to its accuracy rate, and organize a voting amongst neural networks which are

TABLE X: State-of-the-Art Results on CIFAR100-20 by Majority Voting

Number of Classifiers	Accuracy by Majority Vote
State-of-the-Art	50.7
3	55.9
4	56.1
27	58.1

drawn from each of the tiers more easily produce an accuracy rate which is higher than randomly selected neural networks.

VIII. CONCLUSION

Neural networks trained with LFER have different latent representation of decision boundaries. The knobs presented by $(\lambda_i)_{i\geq 1}$ is fine-grained decision boundaries gear for modifying entropy environment of output space which influences decision boundary formation. This presents the possibility of using simple search techniques on λs to automatically mine-train neural ensembles with varied decision boundaries by exploiting the differences in a wide array of of latent representation of decision boundaries about a single dataset. Ensembles trained with LFER can often beat the accuracy of the best neural network possible for a particular architecture on a dataset.

REFERENCES

[1] W. Van Gansbeke, S. Vandenhende, S. Georgoulis, M. Proesmans, and L. Van Gool, "Scan: Learning to classify images without labels," in *European Conference on Computer Vision (ECCV)*, 2020.

[2] J. Guo, H. Chang, and W. Zhu, "Preserving ordinal consensus: Towards feature selection for unlabeled data," in *Proc. AAAI Conf. Artificial Intell. (AAAI)*, New York City, New York, Feb. 2020, pp. 75–82.

[3] I. Misra and L. v. d. Maaten, "Self-supervised learning of pretext-invariant representations," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.

[4] M. Caron, P. Bojanowski, J. Mairal, and A. Joulin, "Leveraging large-scale uncurated data for unsupervised pre-training of visual features," *CoRR*, vol. abs/1905.01278, 2019. [Online]. Available: http://arxiv.org/abs/1905.01278

[5] A. Sablayrolles, M. Douze, C. Schmid, and H. Jégou, "Spreading vectors for similarity search," 2018.

[6] X. Zhang, F. X. Yu, S. Kumar, and S. Chang, "Learning spread-out local feature descriptors," *CoRR*, vol. abs/1708.06320, 2017. [Online]. Available: http://arxiv.org/abs/1708.06320

[7] P. Bojanowski and A. Joulin, "Unsupervised learning by predicting noise," 2017.

[8] H. Jain, P. Pérez, R. Gribonval, J. Zepeda, and H. Jégou, "Approximate search with quantized sparse representations," 2016.

[9] T. Zhang, G.-J. Qi, J. Tang, and J. Wang, "Sparse composite quantization," *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4548–4556, 2015.

[10] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," 2019.

[11] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," 2020.

[12] C.-Y. Chuang, J. Robinson, L. Yen-Chen, A. Torralba, and S. Jegelka, "Debiased contrastive learning," *arXiv preprint arXiv:2007.00224*, 2020.

[13] A. YM., R. C., and V. A., "Self-labelling via simultaneous clustering and representation learning," in *International Conference on Learning Representations (ICLR)*, 2020.

[14] H. V. Vo, F. Bach, M. Cho, K. Han, Y. LeCun, P. Perez, and J. Ponce, "Unsupervised image matching and object discovery as optimization," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 8287–8296.

[15] G. Pereyra, G. Tucker, J. Chorowski, Łukasz Kaiser, and G. Hinton, "Regularizing neural networks by penalizing confident output distributions," 2017.

[16] M. Cuturi, "Sinkhorn distances: Lightspeed computation of optimal transportation distances," 2013.

[17] V. E. Liong, J. Lu, G. Wang, P. Moulin, and J. Zhou, "Deep hashing for compact binary codes learning," *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2475–2483, 2015.

[18] H. Jain, J. Zepeda, P. Pérez, and R. Gribonval, "SUBIC: A supervised, structured binary code for image search," *CoRR*, vol. abs/1708.02932, 2017. [Online]. Available: http://arxiv.org/abs/1708.02932

[19] J. Gomez-Aguilar, "Modeling of a mass-spring-damper system by fractional derivatives with and without a singular kernel," pp. 6289–6303, 2015.

[20] S. Pidhorskyi, R. Almohsen, and G. Doretto, "Generative probabilistic novelty detection with adversarial autoencoders," in *Advances in neural information processing systems*, 2018, pp. 6822–6833.

[21] K. Lee, K. Lee, H. Lee, and J. Shin, "A simple unified framework for detecting out-of-distribution samples and adversarial attacks," 2018.