

Contrastive Classification and Representation Learning with Probabilistic Interpretation

Rahaf Aljundi¹, Yash Patel², Milan Sulc², Nikolay Chumerin¹, Daniel Olmeda Reino¹

¹ Toyota Motor Europe

² Visual Recognition Group, Czech Technical University in Prague

rahaf.al.jundi@toyota-europe.com

Abstract

Cross entropy loss has served as the main objective function for classification-based tasks. Widely deployed for learning neural network classifiers, it shows both effectiveness and a probabilistic interpretation. Recently, after the success of self supervised contrastive representation learning methods, supervised contrastive methods have been proposed to learn representations and have shown superior and more robust performance, compared to solely training with cross entropy loss. However, cross entropy loss is still needed to train the final classification layer. In this work, we investigate the possibility of learning both the representation and the classifier using one objective function that combines the robustness of contrastive learning and the probabilistic interpretation of cross entropy loss. First, we revisit a previously proposed contrastive-based objective function that approximates cross entropy loss and present a simple extension to learn the classifier jointly. Second, we propose a new version of the supervised contrastive training that learns jointly the parameters of the classifier and the backbone of the network. We empirically show that our proposed objective functions show a significant improvement over the standard cross entropy loss with more training stability and robustness in various challenging settings. Supplementary materials can be found in ArXiv.

Introduction

Representation learning is a powerful tool to create an embedding space that is beneficial for performing downstream tasks e.g., classification or retrieval. Contrastive representation learning first proposed by (Chopra, Hadsell, and LeCun 2005) is a dominant successful line for representation learning. It divides the data into pairs of positive (similar) and negative (unrelated) samples with the objective of maximizing the similarity of positive pairs samples and minimize it for negative pairs.

More recently, contrastive learning has become a key component of methods for self-supervised learning (Chen et al. 2020a; Kalantidis et al. 2020; Chen et al. 2020b; Caron et al. 2021) and has shown impressive performance (Caron et al. 2020, 2021) that is very close to the supervised learning counterpart with cross entropy loss. Moreover, it was

shown that supervised contrastive learning marginally outperforms the *cross entropy loss* in fully supervised image classification (Khosla et al. 2020). Not only for standard supervised classification but it has been applied in continual learning (Davari et al. 2022), Out of Distribution Detection (Winkens et al. 2020), Domain Adaptation (Chen et al. 2022) and many more showing superior performance to cross entropy based counterpart.

Minimizing *cross entropy* (CE) loss is widely used in training deep neural network classifiers, derived as the maximum likelihood estimate (MLE) of classifier’s parameters θ to approximate posterior probabilities $\hat{p}(\text{class}|\text{observation})$. (Boudiaf et al. 2020) draw the connection among popular pairwise-distance losses and the cross entropy loss, showing that all of them are related to maximizing the mutual information (MI) between the learned embeddings and the corresponding samples’ labels.

We emphasize the advantages of the probabilistic interpretation of the CE loss in classification problems. Such explicit probabilistic interpretation is missing within the embedding spaces trained by popular contrastive learning methods. The posterior estimates $\hat{p}(\text{class}|\text{observation})$ can be utilized when combining classifiers (Kittler et al. 1998; Breiman 1996; Ju, Bibaut, and van der Laan 2018), in adaptation to prior shift (Saerens, Latinne, and Decaestecker 2002; Sulc and Matas 2019; Alexandari, Kundaje, and Shrikumar 2020; Sipka, Sulc, and Matas 2021), in knowledge distillation (Hinton, Vinyals, and Dean 2015); out-of-distribution detection (Hendrycks and Gimpel 2016) and in many other problems.

In this work, we suggest that one possible reason for the improved performance of supervised contrastive learning is the inherent access to a large number of samples pairs, while the “pairs” within the softmax CE loss are centered around the linear classifier weights. Here we draw an analogy with proxy based loss and consider the linear classifier weights optimized in the softmax CE loss as proxies for learning the samples representations. Proxy base training that utilizes proxies instead of the direct sample to sample relationship is simple and faster to converge, however, it doesn’t leverage the rich data to data similarities as the supervised contrastive loss. We refer to Figure 1 for an illustration on this assumption. We hypothesize that the access to more pairs during training might lead to a better convergence and




Figure 1: Illustration of the possible number of pairs that each loss accesses during training in the learned embedding space, N is the batch size and K is the number of classes. CE pairs are only defined through classes weights while in SupCon each sample forms positive pairs with its class samples and negative pairs with samples from other classes. For our ESUpCon in addition to the positive and negative samples pairs, class weights (prototypes) form positive pairs with corresponding class samples and negative pairs with other classes samples. Note that here we don't consider augmentations.

less overfitting resulting in the advantages hinted in recent works (Khosla et al. 2020; Graf et al. 2021).

Hence, to combine the advantages of contrastive representation learning via pairwise losses and the clear probabilistic interpretation of classifiers trained by cross entropy minimization, we present the following contributions: First, we consider the weights of the last linear classification layer as prototypes of each class. We show that adding a simple term corresponding to maximizing the similarity between the prototypes and their class samples, leads to an assignment of the prototypes to the mean of each class samples with momentum updates of representation. This is optimized during the representation training with a supervised contrastive loss (Khosla et al. 2020), resulting in a nearest prototype classifier (Wohlhart et al. 2013). Second, we propose an extension to the *supervised contrastive loss* (SupCon) (Khosla et al. 2020), where samples of a given class form positive pairs with their class prototype and other classes samples correspond to negative pairs.

We show that the resulting objective combines in its formulation the SupCon loss (Khosla et al. 2020) and the standard CE loss on prototypes related pairs, preserving the probabilistic interpretation of the predictions. We refer to this loss as ESUpCon (short for *Extended Supervised Contrastive loss*).

Third, we revisit the Simplified Pairwise Cross Entropy (SPCE) loss, proposed in the theoretical analysis of (Boudiaf et al. 2020), and compare it with standard CE loss and the supervised contrastive learning loss in an extensive experimental evaluation.

In our experimental evaluation, we not only consider the fully supervised setting but also for the first time a number of challenging settings (low sample regime, imbalanced data and noisy labels). To the best of our knowledge, this is the

first comprehensive evaluation of SupCon loss and the standard CE loss in addition to our proposed extensions.

We show ESUpCon is more powerful as a training objective than the standard CE loss while maintaining a probabilistic interpretation. and is more robust in challenging and low sample settings. Surprisingly, our simple prototypes similarity term is more robust than CE loss for learning a linear classifier after SupCon in most of the imbalanced and noisy data experiments.

In the following, we describe the closely Related Work, then provide a short Background on pairwise losses and the link to Cross Entropy loss, followed by our extension to Supervised Contrastive Loss. We validate and compare different studied losses in the Experiments, and summarize our contributions and limitations in Conclusion.

Related Work

CE loss is a standard and powerful training objective to optimize deep neural networks for classification-related problems. For long, the CE loss was believed to be more effective than representation learning losses e.g., metric learning based losses. For example, (Boudiaf et al. 2020) studied the relation of CE loss to contrastive metric learning losses and showed that the CE loss also has a contrastive and a tightness part. The authors suggested that CE “does it all” and that it is easier to optimize compared to its contrastive-learning counterparts. Recently, self-supervised learning losses have shown great success (Chen and He 2021; Grill et al. 2020; Chen et al. 2020a; He et al. 2020; Chen et al. 2020b; Grill et al. 2020; Caron et al. 2020, 2021) as pretraining methods with only a small performance gap to that of fully supervised learning. The core of the self-supervised methods is the use of rich data augmentation methods to construct positive pairs corresponding to augmented version of a given

sample. Closer to our work, (Li et al. 2020; Caron et al. 2020) construct clusters and establish cluster assignments through prototypes while learning the embedding space. It remains unclear how these losses can be extended to the supervised setting as in our case.

Inspired by the self-supervised SimCLR loss (Chen et al. 2020a), (Khosla et al. 2020) introduced a new supervised contrastive learning method called SupCon, which achieved superior results compared to the standard CE minimization, and which has been shown to be more generalizable and robust to noise. However, the method is only used to train the image representation and still relies on the CE loss to train the linear classifier afterwards. CE-based training suffers from known issues of noise sensitive, overfitting (Berrada, Zisserman, and Kumar 2018), and being less transferable than the representation learning counterparts (Chen et al. 2020a). Recently, (Graf et al. 2021) investigated the difference between the SupCon (Khosla et al. 2020) loss and the CE loss in the geometry of the targeted representation. It was shown that both losses target the same geometric solution, however, SupCon converges much closer to the target leading to a better generalization performance.

As such, starting from the nice suggested characteristics of the SupCon loss based training, we propose and study alternatives that can train the whole network (representation and classifier) end-to-end, while preserving both the performance improvements of contrastive representation learning and the clear probabilistic interpretation of the CE loss. We start by considering the classes weights as prototypes for each class samples. We learn these prototypes while maximizing positive pairs similarities and minimizing negative pairs similarities. Our work hence can be seen as combination of proxy (prototype) based and pairwise based contrastive representation learning. Proxy based losses resort to learning a set of proxies as representative of clusters or classes of samples and optimize the similarities to these proxies rather than the data to data similarities. Proxy NCA (Movshovitz-Attias et al. 2017) was the first proxy base metric learning method, it is an approximation of NCA (Neares Component Analysis) using proxies. We note that in the case of learning with class level labels the Proxy NCA matches learning with Softmax Cross Entropy loss when the last classification layer is without a bias term and its weights are normalized vectors. Proxy anchor loss (Kim et al. 2020), attempts to combine the benefits of both proxy-based and pairwise losses. While in the main loss formulation only similarities to proxies are considered, the magnitude of the loss gradient w.r.t. each sample is scaled by the corresponding proxy similarity proportional to other samples-proxies similarities. In general, proxy based losses do not use the proxy at test time and it is unknown how they perform for classification or whether there can exist any probabilistic interpretations. Circle loss (Sun et al. 2020) presents a unified framework for both pairwise and proxy based losses but it adopts an adaptive scaling of the loss depending on how much a given similarity is deviated from its optimum. In doing so, Circle loss abandons the probabilistic interpretation of a sample assignment to its prototype (proxy).

Background

In this section, we describe recent self-supervised and supervised contrastive losses and the connection with CE loss.

Pairwise Losses

Contrastive losses work with pairs of embeddings that are pulled together if a pair is positive (related embeddings) and pulled further apart otherwise (Chopra, Hadsell, and LeCun 2005). Consider the following: 1) a random data augmentation module that for each sample \mathbf{x} generates two differently augmented samples, 2) a neural network encoder f that maps an augmented input sample \mathbf{x} to its feature representation: $f(\mathbf{x}) = \mathbf{z}, \mathbf{z} \in \mathbb{R}^d$. We start by outlining SimCLR (Chen et al. 2020a), a popular, effective and simple self supervised contrastive loss to lay the ground for our work:

$$\ell_{\text{SimCLR}} = \frac{1}{2N} \sum_i^N \ell_{\text{SimCLR}}(\mathbf{z}_i, \mathbf{z}_{i+N}) + \ell_{\text{SimCLR}}(\mathbf{z}_{i+N}, \mathbf{z}_i), \quad (1)$$

$$\ell_{\text{SimCLR}}(\mathbf{z}_i, \mathbf{z}_j) = -\log \frac{\exp(\text{sim}(\mathbf{z}_i, \mathbf{z}_j)/\tau)}{\sum_{k \neq i} \exp(\text{sim}(\mathbf{z}_i, \mathbf{z}_k)/\tau)}, \quad (2)$$

where τ is the temperature scaling term, N is the mini batch size, and the pairs $(\mathbf{z}_i, \mathbf{z}_j)$ consist of features of two differently augmented views of the same data example and

$\text{sim}(\mathbf{z}_i, \mathbf{z}_j) = \frac{\mathbf{z}_i^\top \mathbf{z}_j}{\|\mathbf{z}_i\| \cdot \|\mathbf{z}_j\|}$ is the cosine similarity. Assuming normalized embedding vectors \mathbf{z}_i , this pairwise loss is:

$$\ell_{\text{SimCLR}}(\mathbf{z}_i, \mathbf{z}_j) = -\mathbf{z}_i^\top \mathbf{z}_j / \tau + \log \sum_{k \neq i} \exp(\mathbf{z}_i^\top \mathbf{z}_k / \tau). \quad (3)$$

Note that the first term corresponding to the positive pair is the tightness term and the second one is the contrastive term. The aforementioned self-supervised batch contrastive approach was extended in (Khosla et al. 2020) to the fully supervised setting with the Supervised Contrastive Loss:

$$\ell_{\text{SupCon}} = \frac{1}{2N} \sum_i^{2N} \ell_{\text{SupCon}}(\mathbf{z}_i, P_i), \quad (4)$$

$$\begin{aligned} \ell_{\text{SupCon}}(\mathbf{z}_i, P_i) = & -\frac{1}{|P_i|} \sum_{\mathbf{z}_p \in P_i} \log \frac{\exp(\text{sim}(\mathbf{z}_i, \mathbf{z}_p)/\tau)}{\sum_{j \neq i} \exp(\text{sim}(\mathbf{z}_i, \mathbf{z}_j)/\tau)} = \\ & \frac{1}{|P_i|} \sum_{\mathbf{z}_p \in P_i} \left(-(\mathbf{z}_i^\top \mathbf{z}_p)/\tau + \log \sum_{j \neq i} \exp((\mathbf{z}_i^\top \mathbf{z}_j)/\tau) \right), \end{aligned} \quad (5)$$

where P_i is the set of representations \mathbf{z}_p forming positive pairs for the i -th sample, and the index j iterates over all (original and augmented) samples. SupCon loss is expressed as the average of the loss defined on each positive pair where in this supervised setting, the positive pairs are formed of augmented views and other samples of the same class. The authors showed that the supervised contrastive learning achieves excellent results in image classification, improving ImageNet classification accuracy with ResNet-50 by 0.5% compared to the best results achieved by training with the CE loss.

Cross Entropy and Pairwise Cross Entropy

The cross entropy (CE) loss is a common choice for training classifiers, as its minimization leads to the maximum likelihood estimate of the classifier parameters for estimating the posterior probabilities $\hat{p}(\text{class}|\text{observation})$.

For N samples of K classes, and a single-label softmax classifier, the CE loss can be defined as follows:

$$\begin{aligned} \ell_{\text{CE}} &= \frac{1}{N} \sum_{i=1}^N \ell_{\text{CE}}(\mathbf{z}_i) = -\frac{1}{N} \sum_{i=1}^N \log \frac{\exp \boldsymbol{\theta}_{y_i}^\top \mathbf{z}_i}{\sum_{k=1}^K \exp \boldsymbol{\theta}_k^\top \mathbf{z}_i} \\ &= -\frac{1}{N} \sum_{i=1}^N \boldsymbol{\theta}_{y_i}^\top \mathbf{z}_i + \frac{1}{N} \sum_{i=1}^N \log \sum_{k=1}^K \exp \boldsymbol{\theta}_k^\top \mathbf{z}_i, \end{aligned} \quad (6)$$

where \mathbf{z}_i is sample feature for the i -th observation having label $y_i \in \{1, \dots, K\}$, and $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_K)$ are the parameters of the last fully connected layer, assuming that no bias term is used.

The *Simplified Pairwise Cross Entropy* (SPCE) loss was introduced in (Boudiaf et al. 2020) as a variant of the CE loss (6):

$$\ell_{\text{SPCE}} = -\frac{1}{N} \sum_{i=1}^N \log \frac{\exp \left(\frac{1}{N} \sum_{j: y_j = y_i} \mathbf{z}_j^\top \mathbf{z}_i \right)}{\sum_{k=1}^K \exp \left(\frac{1}{N} \sum_{j: y_j = k} \mathbf{z}_j^\top \mathbf{z}_i \right)}. \quad (7)$$

When training the feature encoder with the ℓ_{SPCE} loss, the classifier weights $\boldsymbol{\theta}$ can be estimated directly from the class feature means \mathbf{c}_k . Moreover, the class posterior probabilities $p(k|\mathbf{z}_i)$ also can be estimated explicitly:

$$p(k|\mathbf{z}_i) = \frac{\exp \left(\frac{1}{N} \sum_{j: y_j = k} \mathbf{z}_j^\top \mathbf{z}_i \right)}{\sum_{c=1}^K \exp \left(\frac{1}{N} \sum_{j: y_j = c} \mathbf{z}_j^\top \mathbf{z}_i \right)}. \quad (8)$$

In the experimental section, we will evaluate SPCE loss and compare it with SupCon. Differently from SPCE, with SupCon, one needs to train a classifier on top of the learned representation as a posthoc process. In the following we will discuss and propose alternatives to jointly learn the classifier and the feature extraction parameters.

Learning a Classifier Jointly with Representation Learning

Representation learning under SupCon or SPCE losses targets grouping one class samples together while pushing away samples of other classes. In fact, both losses contain tightness and contrastive terms and fulfill similar objectives to that of the cross entropy loss.

Assuming that forcing samples of different classes to lie far apart is achieved by the contrastive part of SupCon or SPCE, in order to learn the parameters of the classifier, one can consider the weight vectors of the linear classifier as prototypes and optimize these prototypes to be closest to the samples of the class they represent (with solely a tightness term). We assume that both the samples representations

and the classifier weights are normalized vectors and that the classifier is linear with no bias term. We define the following loss to learn the desired prototypes:

$$\ell_{\text{tt}} = \frac{1}{N} \sum_i \ell_{\text{tt}}(\mathbf{z}_i, \boldsymbol{\theta}_{y_i}) = \frac{1}{N} \sum_i -\mathbf{z}_i^\top \boldsymbol{\theta}_{y_i}. \quad (9)$$

Note that the number of samples in (9) might differ from N (e.g., due to augmentation), in which case N should be replaced by the corresponding number of samples. With that assumption, the classifier we use is a nearest prototype classifier i.e., assigning a test sample to the class of the nearest prototype. Note that ℓ_{tt} resembles only the tightness part of the CE loss (6). The gradient of the ℓ_{tt} loss w.r.t. the classifier weights can be directly derived from (9):

$$\frac{\partial \ell_{\text{tt}}}{\partial \boldsymbol{\theta}_k} = -\frac{1}{N} \sum_{i: y_i = k} \mathbf{z}_i. \quad (10)$$

Through minimizing this loss jointly with the representation learning loss, we update the classifier weights using the following iterative formula:

$$\boldsymbol{\theta}_k^0 = \eta \frac{1}{N} \sum_{i: y_i = k} \mathbf{z}_i^0, \quad \boldsymbol{\theta}_k^{t+1} = \boldsymbol{\theta}_k^t + \eta \frac{1}{N} \sum_{i: y_i = k} \mathbf{z}_i^{t+1}, \quad (11)$$

where t is the iteration index and η is the learning rate. Note that this is equivalent to setting (up to a constant) the class weights $\boldsymbol{\theta}_k$ to the class features mean \mathbf{c}_k with momentum updates, where the new prototype combines the new iteration representation mean with the previous iteration mean. We will compare the minimization of the ℓ_{tt} loss jointly with the the representation learning loss vs. simply setting the classifier weights $\boldsymbol{\theta}_k$ to the hard mean \mathbf{c}_k for each class k .

Extended Supervised Contrastive Learning

Here we aim at extending the SupCon loss to include the classes prototypes being learned. For this, we propose to consider an explicit linear classification layer with parameters $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_K)$ in the optimization of the supervised contrastive loss (5). Note that here we consider the embeddings \mathbf{z}_i and the class prototypes $\boldsymbol{\theta}_k$ in the same feature space. A class prototype $\boldsymbol{\theta}_k$ should represent as closely as possible its class features. Hence a prototype similarity with its class features should be maximized and minimized with other classes features. To achieve this we propose to construct the following prototype-feature pair $(\mathbf{z}_i, \boldsymbol{\theta}_{y_i})$ with sample representation \mathbf{z}_i ($y_i = k$) as a positive pair. Now we define the following loss on a positive prototype-feature pair:

$$\begin{aligned} \ell_{\text{pt}}(\mathbf{z}_i, \boldsymbol{\theta}_{y_i}) &= -\mathbf{z}_i^\top \boldsymbol{\theta}_{y_i} \\ &+ \log \left(\sum_{k=1}^K \exp(\mathbf{z}_i^\top \boldsymbol{\theta}_k) + \sum_{j=1: j \neq i}^{2N} \exp(\mathbf{z}_i^\top \mathbf{z}_j) \right). \end{aligned} \quad (12)$$

Note that SupCon loss on a positive pair of samples is defined as follows:

$$\ell_{\text{SupCon}}(\mathbf{z}_i, \mathbf{z}_p) = -\mathbf{z}_i^\top \mathbf{z}_p + \log \sum_{j \neq i} \exp(\mathbf{z}_i^\top \mathbf{z}_j). \quad (13)$$

Here we omit the temperature τ for clarity and for a better connection to the CE loss. In (12) we have extended the set of existing data representations \mathbf{z}_i with the class prototypes θ_l . Following the same analogy and constructing all positive prototype-feature pairs, the prototype loss for a class weight θ_k will be defined as follows.

$$\ell_{\text{pt}}(\theta_k) = \frac{1}{2N_k} \sum_{i:y_i=k} \ell_{\text{pt}}(\mathbf{z}_i, \theta_k). \quad (14)$$

Note that the number of summation terms in (14) is $2N_k$ (where N_k is the number of the non-augmented samples in k -th class), since the samples in SupCon are considered with their augmentations. Having the loss defined per prototype θ_k , we can define the full objective function that optimizes the encoder (representation backbone) parameters jointly with the classifier parameters θ as:

$$\ell_{\text{ESupCon}} = \frac{\left(\sum_{k=1}^K \ell_{\text{pt}}(\theta_k) + \sum_i^{2N} \ell_{\text{SupCon}}(\mathbf{z}_i, P_i) \right)}{2N + K}. \quad (15)$$

Next we show that our proposed prototype loss $\ell_{\text{pt}}(\mathbf{z}_i, \theta_k)$ for a given positive pair can be expressed in terms of SupCon loss on that positive pair and CE loss on the concerned sample. Let us define the following:

$$\begin{aligned} T &= \mathbf{z}_i^\top \theta_{y_i}, \quad C_1 = \sum_{k=1}^K \exp(\mathbf{z}_i^\top \theta_k), \quad C_2 = \sum_{j=1:j \neq i}^{2N} \exp(\mathbf{z}_i^\top \mathbf{z}_j), \\ \exp(\ell_{\text{CE}}(\mathbf{z}_i)) &= \exp(-T + \log(C_1)) = \exp(-T)C_1, \\ \exp(\ell_{\text{SupCon}}(\mathbf{z}_i, \theta_{y_i})) &= \exp(-T + \log(C_2 + \exp(T))) \\ &= \exp(-T)(C_2 + \exp(T)), \end{aligned} \quad (16)$$

where T is the tightness term, C_1 is the first contrastive term and C_2 is the second contrastive term, $\ell_{\text{CE}}(\mathbf{z}_i)$ is the CE loss for a sample \mathbf{z}_i , and the SupCon loss $\ell_{\text{SupCon}}(\mathbf{z}_i, \theta_{y_i})$ is estimated after including θ_{y_i} into the pool of representations.

Then our loss for the $(\mathbf{z}_i, \theta_{y_i})$ pair can be expressed as:

$$\begin{aligned} \ell_{\text{pt}}(\mathbf{z}_i, \theta_{y_i}) &= -T + \log(C_1 + C_2) \\ &= \log(\exp(-T + \log(C_1 + C_2))) \\ &= \log(\exp(-T)(C_1 + C_2)) \\ &= \log(\exp(-T)(C_1 + C_2 + \exp(T) - \exp(T))) \\ &= \log(\exp(-T)C_1 + \exp(-T)(C_2 + \exp(T)) \\ &\quad - \exp(-T)\exp(T)) \\ &= \log(\exp(\ell_{\text{CE}}(\mathbf{z}_i)) + \exp(\ell_{\text{SupCon}}(\mathbf{z}_i, \theta_{y_i})) - 1). \end{aligned} \quad (17)$$

As such, minimizing $\ell_{\text{pt}}(\mathbf{z}_i, \theta_{y_i})$ is minimizing the log sum exponential (LSE) of cross entropy loss and supervised contrastive loss for a given positive pair $(\mathbf{z}_i, \theta_{y_i})$, a smooth approximation to the max function. Note that $\ell_{\text{pt}}(\mathbf{z}_i, \theta_{y_i}) = 0 \iff \ell_{\text{CE}}(\mathbf{z}_i) = \ell_{\text{SupCon}}(\mathbf{z}_i, \theta_{y_i}) = 0$.

We refer to the loss in (15) as ESUPCon, short for Extended Supervised Contrastive learning. In the following, we will extensively compare the different studied loss functions.

Experiments

This section serves to compare the performance of deep models trained under the different objective functions discussed earlier including tightness loss term (9) and ESUPCon (15). Our goal is to perform an extensive evaluation of the different losses behaviour not only under fully supervised setting but also under more challenging yet more plausible settings, namely limited data, imbalanced data and noisy labels settings. For the purpose of this experimental validation, we focus on the object recognition problem.

Datasets

We consider Cifar-100, Cifar-10 (Krizhevsky and Hinton 2009), Tiny ImageNet (Le and Yang 2015) (a subset of 200 classes from ImageNet (Deng et al. 2009), rescaled to the 32×32) datasets and Caltech256 (Griffin, Holub, and Perona 2007).

Methods and Implementation Details

In all experiments we use ResNet50 as a main network and evaluate the following losses:

CE: we optimize the network parameters using the standard CE loss. For the SupCon loss (Khosla et al. 2020), we use the publicly available implementation, which uses L2-normalized outputs of a multi-layer head (FC, ReLU, FC), a projection head, on top of the embeddings used for classification. We learn the classifier parameters using: i) Cross entropy loss (SupCon+CE), on the linear layer after optimizing minimizing SupCon loss. ii) For the sake of fair comparison with other losses, we consider also cross entropy loss with no bias term, normalized embeddings and normalized classifier weights. We denote this variant by SupCon+CE (n). iii) Tightness loss (SupCon+Tt), where we optimize the parameters of a linear classifier using (9) during the optimization of the rest of the network (projection head + backbone) with SupCon loss. Note that the gradients of the tightness loss are not propagated to the rest of the network. SPCE: we optimize the backbone with SPCE loss (7) and the classifier weights with the tightness term (9). We also show the performance with directly assigning the weights to the mean of each class samples SPCE(M).

Our ESUPCon: with (15) we optimize jointly a linear classifier and the backbone parameters.

Note that SupCon+CE, SupCon+CE (n) and SupCon+Tt use a projection head, unlike CE, SPCE and ESUPCon. All studied variants benefit from the same type of data augmentations and hyper-parameters were estimated on Cifar-10 dataset and fixed for the rest.

Fully Supervised Classification

We first start by comparing the different studied methods on the standard classification setting while leveraging all the labelled training data of each dataset. Table 1 shows the average test accuracy at the end of the training on the three considered datasets.

First, ESUPCon outperforms CE training alone, using the same number of parameters. SupCon+CE improves over CE. SupCon+Tt is comparable to SupCon+CE (n).

Method	CIFAR-10	CIFAR-100	Tiny ImageNet	Caltech256	Avg.
CE	95.39	76.36	65.76	55.9	–
*SupCon+CE	95.50 +0.11	75.90 –0.46	65.56 –0.20	57.91 +2.01	+0.36
*SupCon+CE (n)	95.27 –0.12	74.57 –1.79	61.69 –4.07	52.92 –2.98	–1.52
*SupCon+Tt	95.20 –0.19	74.80 –1.56	59.66 –6.1	57.42 1.52	–2.24
SPCE	95.62 +0.23	78.15 +1.79	66.52 +0.76	48.46 –7.44	–1.16
SPCE(M)	95.30 –0.09	77.49 +1.13	66.28 +0.52	48.37 –7.52	–1.49
ESupCon	95.9 +0.51	76.92 +0.56	66.2 +0.44	58.27 +2.37	+0.97

Table 1: Accuracy (%) of the different studied and proposed losses on fully labelled datasets. * indicates the use of a projection head. Absolute gains and absolute declines over cross entropy are reported. The last column shows an average improvement or decline over CE, across the datasets.

Method	CIFAR-10			CIFAR-100			Tiny ImageNet			Avg.
	$N = 2K$	$N = 5K$	$N = 10K$	$N = 8K$	$N = 10K$	$N = 20K$	$N = 20K$	$N = 50K$	$N = 70K$	
CE	28.02	69.91	85.08	43.67	51.09	64.31	44.29	57.19	60.94	-
*SupCon+CE	72.27	82.37	88.03	50.96	54.49	64.39	44.00	59.24	62.88	+8.22
*SupCon+CE (n)	71.99	82.73	87.91	50.60	53.92	63.27	43.50	57.62	59.62	+7.41
*SupCon+Tt	72.17	82.97	87.37	51.23	54.49	64.28	43.82	52.21	57.88	+6.88
SPCE	31.81	78.60	86.15	50.09	53.78	64.82	40.94	55.70	55.49	+1.43
ESupCon	74.08	83.89	88.83	48.26	52.58	63.12	44.17	58.66	62.62	+7.97

Table 2: Accuracy (%) on CIFAR-10, CIFAR-100 and Tiny ImageNet for a low-sample training scenario, where N represents the number of samples used for the training. * indicates the use of a projection head. The last column shows an average improvement or decline over cross entropy (CE), across the datasets and the settings.

ESupCon shows the best performance on all four datasets. Except from Caltech dataset, SPCE achieves superior results to CE. When assigning the classifier weights directly to the mean of the features, SPCE(M), results are slightly inferior to the use of our tightness loss (9) for training the classifier parameters. For the rest the of experiments, we show only SPCE, using the suggested tightness term to train the classifier parameters.

Classification in Low-Sample Scenario

After studying the fully labelled scenario, here, we are interested in the performance under limited data setting. Our goal is to see how prone each method is to overfitting in low data regime and whether significant differences can be observed among the different alternatives. Table 2 reports the average test accuracy on Cifar-10, Cifar-100 and Tiny ImageNet using different numbers of training samples (N).

While CE performance is comparable to other losses on the full data scenario, here it is significantly lower than other competitors with a gap increasing as the sample size gets smaller. Except from Tiny ImageNet, SupCon+Tt shows comparable performance to SupCon+CE and is slightly inferior (0.5%) to SupCon+CE (n) on average. SPCE results are better than CE on Cifar-10 and Cifar-100. ESupCon improves significantly over CE while being comparable with SupCon+CE, however, with no projection head. ESupCon is much more robust than SPCE in this setting.

Classification under Imbalanced Data

Our goal is to compare the performance of a model trained by the different studied losses under various challenging settings beside the standard fully supervised setting. Here, we examine the scenario where training data are not uniformly distributed. Some classes are undersampled while others are oversampled. Specifically, we want to test the ability of the different losses to cope with this data nature and learn the underrepresented classes. We simulate this scenario by altering the training data in which half of the categories are underrepresented with a number of samples equals to the imbalance rate (IR) of other categories samples. The test set on which we report the average accuracy remains balanced.

Table 3 reports the average test accuracy of models trained to minimize the different losses on the three considered datasets. For each dataset we consider imbalance rates of 0.05, 0.1, and 0.5 where, for example, an imbalance rate of 0.1 means that the size of undersampled classes samples is 0.1 compared to the oversampled classes size.

Here it seems that SupCon+CE doesn’t improve over CE alone. SPCE results are marginally lower than CE. Our two proposed losses SupCon+Tt and ESupCon exhibit more robust and powerful performance compared to CE with ESupCon performing the best.

Classification under Noisy Data

We continue our investigation on the different losses performance under challenging setting and test another inter-

Method	CIFAR-10			CIFAR-100			Tiny ImageNet			Avg.
	IR = 0.05	IR = 0.1	IR = 0.5	IR = 0.05	IR = 0.1	IR = 0.5	IR = 0.05	IR = 0.1	IR = 0.5	
CE	82.85	87.83	93.99	48.57	54.44	71.19	40.65	46.16	60.30	–
*SupCon+CE	79.94	86.86	94.34	46.79	44.21	71.13	44.96	49.57	62.45	–0.64
*SupCon+CE (n)	47.77	47.64	90.14	40.00	40.32	55.94	35.69	35.61	37.70	–17.24
*SupCon+Tt	85.62	88.76	94.40	54.40	56.79	70.38	44.11	47.39	57.30	+1.46
SPCE	85.62	86.94	93.95	49.59	53.78	68.61	37.27	40.55	61.14	–0.95
ESupCon	86.00	89.26	94.77	52.74	58.08	71.37	45.55	50.90	63.08	+2.86

Table 3: Accuracy (%) on CIFAR-10, CIFAR-100 and Tiny ImageNet for an imbalanced training scenario, where IR represents the rate of imbalance. * indicates the use of a projection head. The last column shows an average improvement or decline over cross entropy (CE), across the datasets and the settings.

Method	CIFAR-10			CIFAR-100			Tiny ImageNet			Avg.
	NR = 0.5	NR = 0.3	NR = 0.2	NR = 0.5	NR = 0.3	NR = 0.2	NR = 0.5	NR = 0.3	NR = 0.2	
CE	60.88	87.08	88.93	35.47	56.57	64.93	31.55	49.62	55.40	–
*SupCon+CE	48.08	74.47	85.94	34.78	58.06	65.57	31.81	46.20	54.74	–3.42
*SupCon+CE (n)	46.35	77.42	87.45	33.55	62.44	67.87	28.88	54.64	58.62	–1.47
*SupCon+Tt	58.05	89.70	90.66	37.23	67.76	69.41	28.67	54.81	57.93	+2.64
SPCE	65.63	88.77	88.93	36.56	60.35	65.75	25.27	42.45	49.52	–0.80
ESupCon	59.18	88.15	90.92	37.04	62.94	65.81	32.555	52.80	56.79	+1.75

Table 4: Accuracy (%) on CIFAR-10, CIFAR-100 and Tiny ImageNet for a noisy training scenario, NR represents the rate of noise. * indicates the use of a projection head. The last column shows an average improvement or decline over cross entropy (CE), across the datasets and the settings.

esting scenario: classification with noisy labels. We want to test the ability of the different training regimes to learn generalizable decision boundaries in spite of the presence of wrongly labelled samples. To simulate this scenario, during training a percentage of the training data, denoted by noise rate (NR), is associated with wrong labels (shuffled labels). As in the previous experiments, we report the results on the standard, correctly labelled, test set. Table 4 reports the average test accuracy on Cifar-10, Cifar-100 and Tiny ImageNet with noise rates of (0.2, 0.3, 0.5). Here we obtained similar results to the imbalanced settings, SupCon+CE doesn't consistently improve over CE, same applies for SPCE. Our both proposed losses improve over CE with SupCon+Tt performing the best here.

General Remarks

We note the following on the shown results of the different losses: CE training after SupCon pretraining (SupCon+CE) improves over standard CE in full and low data regime. However, deploying CE to learn the classifier with or without SupCon pretraining is sensitive to noise and data imbalance. Interestingly, our proposed tightness term is more effective on these two scenarios, however inferior on the full and low data regime. In all studied settings, our proposed ESupCon loss improves over CE and over (SupCon+CE) on the challenging imbalanced and noisy settings.

Conclusion

In this work, we derive novel, robust objective functions, inspired by new evidence showing that contrastive losses improve performance over CE. Driven by the question of whether cross entropy loss is the best option to train jointly a good representation and powerful, generalizable, decision boundaries, we start from a recent approximation to cross entropy loss (SPCE) with pairwise training of representation where classifier weights can be assigned to the mean of each class features. We then suggest to learn the classifier weights under only a tightness term jointly with SupCon representation training or SPCE. Next, we propose an extension to SupCon, where the classifier weights are treated as learnable prototypes in the same space as the samples embeddings, and where data points form positive pairs with their classes prototypes. We show that the proposed loss for a given pair (z_i, θ_k) is a smooth approximation to the maximum of the CE and SupCon losses on that pair. To this point, we test the performance of models trained with the different discussed losses under different challenging settings. We show that the proposed extensions demonstrate more robust and stable performance across different settings and datasets. As a future work, we plan to extend the experiments to object detection and image segmentation problems, as well as to test the discussed losses on Out-Of-Distribution and Continual Learning benchmarks.

References

- Alexandari, A.; Kundaje, A.; and Shrikumar, A. 2020. Maximum likelihood with bias-corrected calibration is hard-to-beat at label shift adaptation. In *International Conference on Machine Learning*, 222–232. PMLR.
- Berrada, L.; Zisserman, A.; and Kumar, M. P. 2018. Smooth loss functions for deep top-k classification. *arXiv preprint arXiv:1802.07595*.
- Boudiaf, M.; Rony, J.; Ziko, I. M.; Granger, E.; Pedersoli, M.; Piantanida, P.; and Ayed, I. B. 2020. A unifying mutual information view of metric learning: cross-entropy vs. pairwise losses. In *European Conference on Computer Vision*, 548–564. Springer.
- Breiman, L. 1996. Bagging predictors. *Machine learning*, 24(2): 123–140.
- Caron, M.; Misra, I.; Mairal, J.; Goyal, P.; Bojanowski, P.; and Joulin, A. 2020. Unsupervised learning of visual features by contrasting cluster assignments. *arXiv preprint arXiv:2006.09882*.
- Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9650–9660.
- Chen, D.; Wang, D.; Darrell, T.; and Ebrahimi, S. 2022. Contrastive Test-Time Adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 295–305.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020a. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Chen, T.; Kornblith, S.; Swersky, K.; Norouzi, M.; and Hinton, G. 2020b. Big Self-Supervised Models are Strong Semi-Supervised Learners. *arXiv preprint arXiv:2006.10029*.
- Chen, X.; and He, K. 2021. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15750–15758.
- Chopra, S.; Hadsell, R.; and LeCun, Y. 2005. Learning a similarity metric discriminatively, with application to face verification. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 1, 539–546. IEEE.
- Davari, M.; Asadi, N.; Mudur, S.; Aljundi, R.; and Belilovsky, E. 2022. Probing Representation Forgetting in Supervised and Unsupervised Continual Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16712–16721.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Graf, F.; Hofer, C.; Niethammer, M.; and Kwitt, R. 2021. Dissecting supervised contrastive learning. In *International Conference on Machine Learning*, 3821–3830. PMLR.
- Griffin, G.; Holub, A.; and Perona, P. 2007. Caltech-256 object category dataset. Technical report.
- Grill, J.-B.; Strub, F.; Altché, F.; Tallec, C.; Richemond, P. H.; Buchatskaya, E.; Doersch, C.; Pires, B. A.; Guo, Z. D.; Azar, M. G.; et al. 2020. Bootstrap your own latent: A new approach to self-supervised learning. *arXiv preprint arXiv:2006.07733*.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9729–9738.
- Hendrycks, D.; and Gimpel, K. 2016. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *arXiv preprint arXiv:1610.02136*.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Ju, C.; Bibaut, A.; and van der Laan, M. 2018. The relative performance of ensemble methods with deep convolutional neural networks for image classification. *Journal of Applied Statistics*, 45(15): 2800–2818.
- Kalantidis, Y.; Sariyildiz, M. B.; Pion, N.; Weinzaepfel, P.; and Larlus, D. 2020. Hard negative mixing for contrastive learning. *arXiv preprint arXiv:2010.01028*.
- Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; and Krishnan, D. 2020. Supervised Contrastive Learning. *arXiv preprint arXiv:2004.11362*.
- Kim, S.; Kim, D.; Cho, M.; and Kwak, S. 2020. Proxy anchor loss for deep metric learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3238–3247.
- Kittler, J.; Hatef, M.; Duin, R. P.; and Matas, J. 1998. On combining classifiers. *IEEE transactions on pattern analysis and machine intelligence*, 20(3): 226–239.
- Krizhevsky, A.; and Hinton, G. 2009. Learning multiple layers of features from tiny images. Technical Report 0, University of Toronto, Toronto, Ontario.
- Le, Y.; and Yang, X. 2015. Tiny imagenet visual recognition challenge. *CS 231N course, Stanford Univ., Stanford, CA, USA*, 7(7): 3.
- Li, J.; Zhou, P.; Xiong, C.; and Hoi, S. C. 2020. Prototypical contrastive learning of unsupervised representations. *arXiv preprint arXiv:2005.04966*.
- Movshovitz-Attias, Y.; Toshev, A.; Leung, T. K.; Ioffe, S.; and Singh, S. 2017. No fuss distance metric learning using proxies. In *Proceedings of the IEEE International Conference on Computer Vision*, 360–368.
- Saerens, M.; Latinne, P.; and Decaestecker, C. 2002. Adjusting the outputs of a classifier to new a priori probabilities: a simple procedure. *Neural computation*, 14(1): 21–41.
- Sipka, T.; Sulc, M.; and Matas, J. 2021. The Hitchhiker’s Guide to Prior-Shift Adaptation. *arXiv preprint arXiv:2106.11695*.

Sulc, M.; and Matas, J. 2019. Improving cnn classifiers by estimating test-time priors. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, 0–0.

Sun, Y.; Cheng, C.; Zhang, Y.; Zhang, C.; Zheng, L.; Wang, Z.; and Wei, Y. 2020. Circle loss: A unified perspective of pair similarity optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6398–6407.

Winkens, J.; Bunel, R.; Roy, A. G.; Stanforth, R.; Natarajan, V.; Ledsam, J. R.; MacWilliams, P.; Kohli, P.; Karthikesalingam, A.; Kohl, S.; et al. 2020. Contrastive training for improved out-of-distribution detection. *arXiv preprint arXiv:2007.05566*.

Wohlhart, P.; Kostinger, M.; Donoser, M.; Roth, P. M.; and Bischof, H. 2013. Optimizing 1-nearest prototype classifiers. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 460–467.