

Decoupled Contrastive Learning

Chun-Hsiao Yeh^{1,2}, Cheng-Yao Hong¹, Yen-Chi Hsu^{1,3}, Tyng-Luh Liu^{*1},
Yubei Chen⁴, and Yann LeCun^{4,5}

¹ IIS, Academia Sinica, Taiwan

² UC Berkeley

³ National Taiwan University

⁴ Meta AI Research

⁵ New York University

{sensible,yenchi,liutyng}@iis.sinica.edu.tw,
daniel_yeh@berkeley.edu, {yubeic,yann}@fb.com

Abstract. Contrastive learning (CL) is one of the most successful paradigms for self-supervised learning (SSL). In a principled way, it considers two augmented “views” of the same image as *positive* to be pulled closer, and all other images as *negative* to be pushed further apart. However, behind the impressive success of CL-based techniques, their formulation often relies on heavy-computation settings, including large sample batches, extensive training epochs, etc. We are thus motivated to tackle these issues and establish a simple, efficient, yet competitive baseline of contrastive learning. Specifically, we identify, from theoretical and empirical studies, a noticeable *negative-positive-coupling* (NPC) effect in the widely used InfoNCE loss, leading to unsuitable learning efficiency concerning the batch size. By removing the NPC effect, we propose decoupled contrastive learning (DCL) loss, which removes the positive term from the denominator and significantly improves the learning efficiency. DCL achieves competitive performance with less sensitivity to sub-optimal hyperparameters, requiring neither large batches in SimCLR, momentum encoding in MoCo, or large epochs. We demonstrate with various benchmarks while manifesting robustness as much less sensitive to suboptimal hyperparameters. Notably, SimCLR with DCL achieves 68.2% ImageNet-1K top-1 accuracy using batch size 256 within 200 epochs pre-training, outperforming its SimCLR baseline by 6.4%. Further, DCL can be combined with the SOTA contrastive learning method, NNCLR, to achieve 72.3% ImageNet-1K top-1 accuracy with 512 batch size in 400 epochs, which represents a new SOTA in contrastive learning. We believe DCL provides a valuable baseline for future contrastive SSL studies.

Keywords: Contrastive learning, self-supervised learning

* Corresponding author. E-mail:liutyng@iis.sinica.edu.tw




Fig. 1. An overview of the batch size issue is that general contrastive approaches need large batch sizes to perform better: (a) shows the NPC multiplier q_B in different batch sizes. As the batch size gradually increases, the q_B will approach to 1 with a small coefficient of variation ($C_v = \sigma/\mu$); and (b) illustrates the distribution of q_B with various batch sizes and indicates that the mode value of q_B will shift towards 1 when the batch size increases. Note that the σ and μ are the standard deviation and mean of q_B , respectively. The coefficient of variation, C_v , measures the dispersion of a frequency distribution.

1 Introduction

As a fundamental task in machine learning, representation learning aims to extract useful information from the raw data for the downstream tasks. It has been regarded as a long-acting goal over the past decades. Recent progress on representation learning has achieved a significant milestone over self-supervised learning (SSL), facilitating feature learning with its competence in exploiting massive raw data without any annotated supervision. In the early stage of SSL, representation learning has focused on exploiting pretext tasks, which are addressed by generating pseudo-labels to the unlabeled data through different transformations, such as solving jigsaw puzzles [20], colorization [36] and rotation prediction [12]. Though these approaches succeed in computer vision, there is a large gap between these methods and supervised learning. Recently, there has been a significant advancement in using contrastive learning [31, 21, 26, 15, 7] for self-supervised pre-training, which significantly closes the gap between the SSL method and supervised learning. Contrastive SSL methods, e.g., SimCLR [7], in general, try to pull different views of the same instance close and push different instances far apart in the representation space.

Despite the evident progress of the state-of-the-art contrastive SSL methods, there have been facing several challenges into future development in this direction, including 1) The SOTA models, e.g., [15] may require specific structures such as the momentum encoder and large memory queues, which may complicate the underlying representation learning. 2) The contrastive SSL models, e.g., [7] often depend on large batch size and huge epoch numbers to achieve competitive

performance, posing a computational challenge for academia to explore this direction. 3) They tend to be sensitive to hyperparameters and optimizers, introducing additional difficulty reproducing the results on various benchmarks.

Through the analysis of the widely adopted InfoNCE loss in contrastive learning, we identified a negative-positive-coupling (NPC) multiplier q_B in the gradient as shown in Proposition 1. The NPC multiplier modulates the gradient of each sample, and it reduces the learning efficiency due to easy SSL classification tasks: 1) when a positive sample is very close to the anchor; 2) when negative samples are far away from the anchor; and 3) when there is only a small number of negative samples (i.e., a small batch size). A less-informative (nearby) positive view would reduce the gradient from a batch of informative negative samples or vice versa. Such a coupling exacerbates when smaller batch sizes are used.

Meanwhile, we also investigate the relationship between q_B and batch size through the baseline, SimCLR. As can be seen in Figure 1, the distribution of q_B has a strong positive correlation with the batch size. Figure 1(a) shows that when batch size gradually increases, q_B not only approaches 1 but also reduces the coefficient of variation C_v . The distribution with larger C_v has low statistical dispersion and vice versa. Figure 1(b) indicates that the mode value of q_B will also shift from 0 to 1 when the batch size becomes larger. Hence, it is reasonable to fix the value of q_B , alleviating the influence of batch size.

By removing the coupling term from the Info-NCE loss, we reach a new formulation, the *decoupled contrastive learning* (DCL). The new objective function significantly improves the training efficiency with less sensitivity to sub-optimal hyper-parameters requires neither large batches, momentum encoding, or large epochs to achieve competitive performance on various benchmarks. The main contributions of the proposed DCL can be characterized as follows:

- 1) We provide both theoretical analysis and empirical evidence to show the NPC effect in the InfoNCE-based contrastive learning;
- 2) We introduce DCL objective, which casts off the NPC coupling phenomenon, significantly improves the training efficiency, and it is less sensitive to sub-optimal hyper-parameters;
- 3) Extensive experiments are provided to show the effectiveness of the proposed method that DCL achieves competitive performance **without** large batch sizes, large training epochs, momentum encoding, or additional tricks such as stop-gradient and multi-cropping, etc. This leads to a plug-and-play improvement to the widely adopted InfoNCE-based contrastive learning;
- 4) We show that DCL can be easily combined with the SOTA contrastive methods, e.g. NNCLR [10], to achieve further improvements.

2 Related Work

Contrastive Learning. Contrastive learning (CL) constructs positive and negative sample pairs to extract information from the data itself. In CL, each anchor image in a batch has only one positive sample to construct a positive sample




Fig. 2. Contrastive learning and negative-positive coupling (NPC). (a) In SimCLR, each sample \mathbf{x}_i has two augmented views $\{\mathbf{x}_i^{(1)}, \mathbf{x}_i^{(2)}\}$. They are encoded by the same encoder f and further projected to $\{\mathbf{z}_i^{(1)}, \mathbf{z}_i^{(2)}\}$ by a normalized MLP. (b) According to Equation 4. For the view $\mathbf{x}_i^{(1)}$, the cross-entropy loss $L_i^{(1)}$ leads to a positive force $\mathbf{z}_i^{(2)}$, which comes from the other view $\mathbf{x}_i^{(2)}$ of \mathbf{x} and a negative force, which is a weighted average of all the negative samples, i.e. $\{\mathbf{z}_j^{(l)} | l \in \{1, 2\}, j \neq i\}$. However, the gradient $-\nabla_{\mathbf{z}_i^{(2)}} L_i^{(1)}$ is proportional to the NPC multiplier. (c) We show two cases when the NPC term affects learning efficiency. The positive sample is close to the anchor and less informative on the top. However, the gradient from the negative samples is also reduced. On the bottom, when the negative samples are far away and less informative, the learning rate from the positive sample is mistakenly reduced. In general, the NPC multiplier from the InfoNCE loss makes the SSL task simpler to solve, leading to reduced learning efficiency.

pair [14,7,15]. CPC [21] predicts the future output of sequential data by using current output as prior knowledge, which can improve the feature representing the ability of the model. Instance discrimination [31] proposes a non-parametric cross-entropy loss to optimize the model at the instance level. Inv. spread [32] makes use of data augmentation invariant and the spread-out property of instance to learn features. MoCo [15] proposes a dictionary to maintain a negative sample set, thus increasing the number of negative sample pairs. Different from the aforementioned self-supervised CL approaches, [19] proposes a supervised CL that considers all the same categories as positive pairs to increase the utility of images.

Collapsing Issue on the Number of Negatives. In CL, the objective is to maximize the mutual information between the positive pairs. However, to avoid the “*collapsing output*”, vast quantities of negative samples are needed so that the learning objectives obtain the maximum similarity and have the minimum similarity with negative samples. For instance, in SimCLR [7], training requires many negative samples, leading to a large batch size (i.e., 4096). Furthermore, to optimize such a huge batch, a specially designed optimizer LARS [33] is used. Similarly, MoCo [15] needs a vast queue (i.e., 65536) to achieve competitive performance. BYOL [13] does not collapse output without using any negative samples by considering all the images are positive and to maximize the similarity of “projection” and “prediction” features. On the other hand, SimSiam [9] leverages

the Siamese network to introduce inductive biases for modeling invariance. With the small batch size (i.e., 256), SimSiam is a rival to BYOL (i.e., 4096). Unlike both approaches that achieved their success through empirical studies, this paper tackles from a theoretical perspective, proving that an intertwined multiplier q_B of positive and negative is the main issue to contrastive learning.

Batch Size Sensitivity on InfoNCE. Several works of literature focus on batch size sensitivity concerning the InfoNCE objective function. [27] proposes an objective based on relative predictive coding that maintains the balance between training stability and batch size sensitivity. [17] follows the [3] and extends the idea between the local and global features. [22] proposes a Wasserstein distance to prevent the encoder from learning any other differences between unpaired samples. [18] and [25] learn better representation by sampling hard negatives, particularly for small batches. Other recent works [37, 11] aim to mitigate the issue of small batch size in InfoNCE loss. Although the basic principle of recent works and DCL is derived from InfoNCE objective function, we provide a novel perspective to support the decoupling between positive and negative terms in InfoNCE loss is essential. Simply removing the term from the denominator pre-training to positive pairs can drastically improve the performance and keep the objective function invariant to batch size sensitivity.

3 Decouple Negative and Positive Samples in Contrastive Learning

We choose to start from SimCLR because of its conceptual simplicity. Given a batch of N samples (e.g. images), $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$, let $\mathbf{x}_i^{(1)}, \mathbf{x}_i^{(2)}$ be two augmented views of the sample x_i and B be the set of all of the augmented views in the batch, i.e. $B = \{\mathbf{x}_i^{(k)} | k \in \{1, 2\}, i \in [1, N]\}$. As shown by Figure 2(a), each of the views $\mathbf{x}_i^{(k)}$ is sent into the same encoder network f and the output $\mathbf{h}_i^{(k)} = f(\mathbf{x}_i^{(k)})$ is then projected by a normalized MLP projector that $\mathbf{z}_i^{(k)} = g(\mathbf{h}_i^{(k)}) / \|g(\mathbf{h}_i^{(k)})\|$. For each augmented view $\mathbf{x}_i^{(k)}$, SimCLR solves a classification problem by using the rest of the views in B as targets, and assigns the only positive label to $\mathbf{x}_i^{(u)}$, where $u \neq k$. So SimCLR creates a cross-entropy loss function $L_i^{(k)}$ for each view $\mathbf{x}_i^{(k)}$, and the overall loss function is $L = \sum_{k \in \{1, 2\}, i \in [1, N]} L_i^{(k)}$.

$$L_i^{(k)} = -\log \frac{\exp(\langle \mathbf{z}_i^{(1)}, \mathbf{z}_i^{(2)} \rangle / \tau)}{\exp(\langle \mathbf{z}_i^{(1)}, \mathbf{z}_i^{(2)} \rangle / \tau) + U_{i,k}}, \quad (1)$$

where

$$U_{i,k} = \sum_{l \in \{1, 2\}, j \in [1, N], j \neq i} \exp(\langle \mathbf{z}_i^{(k)}, \mathbf{z}_j^{(l)} \rangle / \tau) \quad (2)$$

means the summation of negative terms for the view k of the sample i .

Proposition 1 : *There exists a negative-positive coupling (NPC) multiplier $q_{B,i}^{(1)}$ in the gradient of $L_i^{(1)}$:*

$$\begin{cases} -\nabla_{\mathbf{z}_i^{(1)}} L_i^{(1)} = \frac{q_{B,i}^{(1)}}{\tau} \left(\mathbf{z}_i^{(2)} - \sum_{l \in \{1,2\}, j \in \llbracket 1, N \rrbracket, j \neq i} \frac{\exp(\langle \mathbf{z}_i^{(1)}, \mathbf{z}_j^{(l)} \rangle / \tau)}{U_{i,1}} \cdot \mathbf{z}_j^{(l)} \right) \\ -\nabla_{\mathbf{z}_i^{(2)}} L_i^{(1)} = \frac{q_{B,i}^{(1)}}{\tau} \cdot \mathbf{z}_i^{(1)} \\ -\nabla_{\mathbf{z}_j^{(l)}} L_i^{(1)} = -\frac{q_{B,i}^{(1)}}{\tau} \frac{\exp(\langle \mathbf{z}_i^{(1)}, \mathbf{z}_j^{(l)} \rangle / \tau)}{U_{i,1}} \cdot \mathbf{z}_i^{(1)} \end{cases} \quad (3)$$

where the NPC multiplier $q_{B,i}^{(1)}$ is:

$$q_{B,i}^{(1)} = 1 - \frac{\exp(\langle \mathbf{z}_i^{(1)}, \mathbf{z}_i^{(2)} \rangle / \tau)}{\exp(\langle \mathbf{z}_i^{(1)}, \mathbf{z}_i^{(2)} \rangle / \tau) + U_{i,1}} \quad (4)$$

and $U_{i,1} = \sum_{l \in \{1,2\}, j \in \llbracket 1, N \rrbracket, j \neq i} \exp(\langle \mathbf{z}_i^{(1)}, \mathbf{z}_j^{(l)} \rangle / \tau)$. Due to the symmetry, a similar NPC multiplier $q_{B,i}^{(k)}$ exists in the gradient of $L_i^{(k)}$, $k \in \{1, 2\}$, $i \in \llbracket 1, N \rrbracket$.

As we can see, all of the partial gradients in Equation 3 are modified by the common NPC multiplier $q_{B,i}^{(k)}$ in Equation 4. Equation 4 makes intuitive sense: when the SSL classification task is easy, the gradient would be reduced by the NPC term. However, the positive samples and negative samples are strongly coupled. When the negative samples are far away and less informative (easy negatives), the gradient from an informative, positive sample would be reduced by the NPC multiplier $q_{B,i}^{(1)}$. On the other hand, when the positive sample is close (easy positive) and less informative, the gradient from a batch of informative negative samples would also be reduced by the NPC multiplier. When the batch size is smaller, the SSL classification problem can be significantly simpler to solve. As a result, the learning efficiency can be significantly reduced with a small batch size setting.

Figure 1(b) shows the NPC multiplier q_B distribution shift w.r.t. different batch sizes for a pre-trained SimCLR baseline model. While all of the shown distributions have prominent fluctuation, the smaller batch size makes q_B cluster towards 0, while the larger batch size pushes the distribution towards $\delta(1)$. Figure 1(a) shows the averaged NPC multiplier $\langle q_B \rangle$ changes w.r.t. the batch size and the relative fluctuation. The small batch sizes introduce significant NPC fluctuation. Based on this observation, we propose to remove the NPC multipliers from the gradients, which corresponds to the case $q_{B,N \rightarrow \infty}$. This leads to the decoupled contrastive learning formulation. [29] also proposes an alignment & uniformity loss which does not have the NPC. However, a similar analysis introduces negative-negative coupling from different positive samples. In other words, [29] considers all the negative samples in the batch together, which may cause the gradient to be dominated by a specific negative pair. In Appendix 5, we provide a thorough discussion and demonstrate the advantage of DCL loss against [29].

Proposition 2 the DCL Loss: *Removing the positive pair from the denominator of Equation 1 leads to a decoupled contrastive learning loss. If we remove the NPC multiplier $q_{B,i}^{(k)}$ from Equation 3, we reach a decoupled contrastive learning loss $L_{DC} = \sum_{k \in \{1,2\}, i \in \llbracket 1, N \rrbracket} L_{DC,i}^{(k)}$, where $L_{DC,i}^{(k)}$ is:*

$$L_{DC,i}^{(k)} = -\log \frac{\exp(\langle \mathbf{z}_i^{(1)}, \mathbf{z}_i^{(2)} \rangle / \tau)}{\exp(\langle \mathbf{z}_i^{(1)}, \mathbf{z}_i^{(2)} \rangle / \tau) + U_{i,k}} \quad (5)$$

$$= -\langle \mathbf{z}_i^{(1)}, \mathbf{z}_i^{(2)} \rangle / \tau + \log U_{i,k} \quad (6)$$

The proofs of Proposition 1 and 2 are given in Appendix. Further, we can generalize the loss function L_{DC} to L_{DCW} by introducing a weighting function for the positive pairs i.e. $L_{DCW} = \sum_{k \in \{1,2\}, i \in \llbracket 1, N \rrbracket} L_{DCW,i}^{(k)}$.

$$L_{DCW,i}^{(k)} = -w(\mathbf{z}_i^{(1)}, \mathbf{z}_i^{(2)}) (\langle \mathbf{z}_i^{(1)}, \mathbf{z}_i^{(2)} \rangle / \tau) + \log U_{i,k} \quad (7)$$

where we can intuitively choose w to be a negative von Mises-Fisher weighting function that $w(\mathbf{z}_i^{(1)}, \mathbf{z}_i^{(2)}) = 2 - \frac{\exp(\langle \mathbf{z}_i^{(1)}, \mathbf{z}_i^{(2)} \rangle / \sigma)}{\mathbb{E}_i [\exp(\langle \mathbf{z}_i^{(1)}, \mathbf{z}_i^{(2)} \rangle / \sigma)]}$ and $\mathbb{E} [w] = 1$. L_{DC} is a special case of L_{DCW} and we can see that $\lim_{\sigma \rightarrow \infty} L_{DCW} = L_{DC}$. The intuition behind $w(\mathbf{z}_i^{(1)}, \mathbf{z}_i^{(2)})$ is that there is more learning signal when a positive pair of samples are far from each other, and $\mathbb{E} [w(\mathbf{z}_i^{(1)}, \mathbf{z}_i^{(2)}) \langle \mathbf{z}_i^{(1)}, \mathbf{z}_i^{(2)} \rangle] \approx \mathbb{E} [\langle \mathbf{z}_i^{(1)}, \mathbf{z}_i^{(2)} \rangle]$. Other similar weight functions also provide similar results. In general, we find such a weighting function, which gives a larger weight to the hard positives tend to increase the representation quality.

4 Experiments

This section empirically evaluates the proposed decoupled contrastive learning (DCL) and compares it to general contrastive learning methods. We summarize the experiments and analysis as the following: (1) the proposed work significantly outperforms the general InfoNCE-based contrastive learning on both large-scale and small-scale vision benchmarks; (2) we show that the enhanced version of DCL, DCLW, could further improve the representation quality; and (3) we further analyze DCL with ablation studies on ImageNet-1K, hyperparameters, and few learning epochs, which shows fast convergence of the proposed DCL. Note that all the experiments are conducted with 8 Nvidia V100 GPUs on a single machine.

4.1 Implementation Details

ImageNet. For a fair comparison on ImageNet data, we implement the proposed decoupled structure, DCL, by following SimCLR [7] with ResNet-50 [16] as the encoder backbone and use cosine annealing schedule with SGD optimizer. We set the temperature τ to 0.1 and the latent vector dimension to 128. Following




Fig. 3. Comparisons on ImageNet-1K with/without DCL under different numbers of (a): batch sizes for SimCLR and (b): queues for MoCo. Without DCL, the top-1 accuracy significantly drops when batch size (SimCLR) or queues (MoCo) becomes very small. Note that the temperature τ is 0.1 for SimCLR and 0.07 for MoCo in the comparison.

the OpenSelfSup benchmark [35], we evaluate the pre-trained models by training a linear classifier with frozen learned embedding on ImageNet data. We further consider evaluating DCL on ImageNet-100, a selected subset of 100 classes of ImageNet-1K. Note that all models on ImageNet are trained for 200 epochs.

CIFAR and STL10. For CIFAR10, CIFAR100, and STL10, ResNet-18 [16] is used as the encoder architecture. Following the small-scale benchmark [30], we set the temperature τ to 0.07. All models are trained for 200 epochs with SGD optimizer, a base $lr = 0.03 * batchsize / 256$, and evaluated by k nearest neighbor (kNN) classifier. Note that on STL10, we include both the *train* and *unlabeled* set for model pre-training. We further use ResNet-50 as a stronger backbone by following the implementation [24], using the same backbone and hyperparameters.

4.2 Experiments and Analysis

DCL on ImageNet. This section illustrates the effect of DCL against InfoNCE-based approaches under different batch sizes and queues. The initial setup is to have 1024 batch size (SimCLR) and 65536 queues (MoCo [15]) and gradually reduce the batch size (SimCLR) and queue (MoCo) to show the corresponding top-1 accuracy by linear evaluation. Figure 3 indicates that without DCL, the top-1 accuracy drastically drops when batch size (SimCLR) or queue (MoCo) becomes very small. While with DCL, the performance keeps steadier than baselines (SimCLR: -4.1% vs. -8.3% , MoCo: -0.4% vs. -5.9%).

Specifically, Figure 3 further shows that in SimCLR, the performance with DCL improves from 61.8% to 65.9% under 256 batch size; MoCo with DCL

Table 1. Comparisons with/without DCL under different batch sizes from 32 to 512. Results show the effectiveness of DCL on five widely used benchmarks. The performance of DCL keeps steadier than the SimCLR baseline while the batch size is varied.

Batch Size	32	64	128	256	512
Dataset	ImageNet-1K (kNN / Linear)				
Baseline (ResNet-50)	40.2/56.8	42.9/58.9	45.1/60.6	46.3/61.8	49.4/64.0
w/ DCL (ResNet-50)	43.7/61.5	46.3/63.4	48.5/64.3	49.8/65.9	50.1/65.8
Dataset	ImageNet-100 (kNN / Linear)				
Baseline (ResNet-50)	67.8/74.2	71.9/77.6	73.2/79.3	74.6/80.7	75.4/81.3
w/ DCL (ResNet-50)	74.9/80.8	76.3/82.0	76.5/81.9	76.9/83.1	76.8/82.8
Dataset	CIFAR-10 (kNN / Linear)				
Baseline (ResNet-18)	78.9/79.8	80.4/81.3	81.1/82.8	81.4/83.0	81.3/83.3
w/ DCL (ResNet-18)	83.7/85.1	84.4/85.9	84.4/85.7	84.2/85.3	83.5/84.7
Dataset	CIFAR-100 (kNN / Linear)				
Baseline (ResNet-18)	49.4/51.3	50.3/53.8	51.8/55.3	52.0/56.3	52.4/56.8
w/ DCL (ResNet-18)	51.1/55.4	54.3/58.3	54.6/58.9	54.9/58.5	55.0/58.4
Dataset	STL-10 (kNN / Linear)				
Baseline (ResNet-18)	74.1/76.2	77.6/77.8	79.3/80.0	80.7/81.3	81.3/81.5
w/ DCL (ResNet-18)	82.0/85.2	82.8/86.3	81.8/86.1	81.2/85.7	81.0/85.6

improves from 54.7% to 60.8% under 256 queues. The comparison fully demonstrates the necessity of DCL, especially when the number of negatives is small. Although batch size increases to 1024, DCL (66.1%) still improves over the SimCLR baseline (65.1%).

We further observe the same phenomenon on ImageNet-100 data. Table 1 shows that, with DCL, the top-1 linear performance only drops 2.3% compared to the InfoNCE baseline (SimCLR) of 7.1% when the batch size is varied.

In summary, it is worth noting that, while the batch size is small, the strength of $q_{B,i}$, which is used to push the negative samples away from the positive sample, is also relatively weak. This phenomenon tends to reduce the efficiency of learning representation. While taking advantage of DCL alleviates the performance gap between small and large batch sizes. Hence, through the analysis, we find out DCL can simply tackle the batch size issue in contrastive learning. With this considerable advantage given by DCL, general SSL approaches can be implemented with fewer computational resources or lower standard platforms. Compared to InfoNCE, DCL is more applicable across all large-scale SSL applications.

DCL on CIFAR and STL10. For STL10, CIFAR10, and CIFAR100, we implement DCL with ResNet-18 as encoder backbone. In Table 1, it is observed that DCL also demonstrates its strong effectiveness on small-scale benchmarks. In the evaluation (kNN / Linear) summary, DCL outperforms its baseline by 4.8% / 5.3% (CIFAR10) and 1.7% / 4.4% (CIFAR100) under a small batch size 32.

Table 2. Comparisons between SimCLR baseline, DCL, and DCLW. The linear and kNN top-1 (%) results indicate that DCL improves baseline performance, and DCLW further provides an extra boost. Note that results are under batch size 256 and epoch 200. All models are both trained and evaluated with the same experimental settings. The backbones are ResNet-18 and ResNet-50 for CIFAR and ImageNet, respectively.

Dataset	CIFAR10 (kNN)	CIFAR100 (kNN)	ImageNet-100 (linear)	ImageNet-1K (linear)
SimCLR	81.4	52.0	80.7	61.8
DCL	84.2 (+2.8)	54.9 (+2.9)	83.1 (+2.4)	65.9 (+4.1)
DCLW	84.8 (+3.4)	55.2 (+3.2)	84.2 (+3.5)	66.9 (+5.1)

Table 3. Improve the DCL model performance on ImageNet-1K with tuned hyperparameters: temperature and learning rate, and stronger image augmentation. Note that models are trained with 256 batch size and 200 epochs.

ImageNet-1K (256 Batch size; 200 epoch)	Linear Top-1 Accuracy (%)
DCL	65.9
+ optimal $(\tau, l_r) = (0.2, 0.07)$	67.8 (+1.9)
+ asymmetric augmentation [13]	68.2 (+0.4)

The accuracy (kNN / Linear) of the SimCLR baseline on STL10 is also improved significantly by 7.9% / 9.0%.

Decoupled Objective with Re-Weighting DCLW. We only replace L_{DC} with L_{DCW} with no possible advantage from additional tricks. Both DCL and the baselines apply the same training instruction of the OpenSelfSup benchmark for fairness. Note that we empirically choose $\sigma = 0.5$ in the experiments. Results in Table 2 indicates that, DCLW achieves extra 5.1% (ImageNet-1K), 3.5% (ImageNet-100) gains compared to the baseline. For CIFAR data, an extra 3.4% (CIFAR10) 3.2% is gained from the addition of DCLW. It is worth noting that, trained with 200 epochs, DCLW reaches 66.9% with batch size 256, surpassing the SimCLR baseline: 66.2% with batch size 8192.

4.3 Ablations

We perform extensive ablations on the hyperparameters of DCL on both ImageNet data and other small-scale data, i.e., CIFAR and STL10. By seeking better configurations empirically, we see that DCL gives consistent gains over the standard InfoNCE baselines (SimCLR and MoCo-v2). In other ablations, we see that DCL achieves more gains over both SimCLR and MoCo-v2, i.e., InfoNCE-based baselines, also when training for 100 epochs only.

DCL Ablations on ImageNet. In Table 3, we have slightly improved the DCL model performance on ImageNet-1K: 1) tuned hyperparameters, temperature τ and learning rate ; 2) asymmetric image augmentation (e.g., BYOL). To obtain a stronger baseline, we conduct an empirical hyperparameter search with batch size 256 and 200 epochs. This improves DCL from 65.9% to 67.8% top-1 accuracy

Table 4. The comparisons with/without DCL under various batch sizes from 32 to 512 on ResNet-50.

Architecture@epoch	ResNet-50@500 epoch									
Dataset	CIFAR10 (kNN)					CIFAR100 (kNN)				
Batch Size	32	64	128	256	512	32	64	128	256	512
SimCLR	82.2	85.9	88.5	88.9	89.1	49.8	55.3	59.9	60.6	61.1
SimCLR w/ DCL	86.1	88.3	89.9	90.1	90.3	54.3	58.4	61.6	62.0	62.2

Table 5. Linear top-1 accuracy (%) comparison with MoCo-V2 on ImageNet-1K and ImageNet-100.

Queue Size	32	64	128	256	8192	64	256	65536
Dataset	ImageNet-100 (Linear)					ImageNet-1K (Linear)		
MoCo-v2 Baseline (ResNet-50)	73.7	76.4	78.7	78.7	79.8	63.9	67.1	67.5
MoCo-v2 w/DCL (ResNet-50)	76.2	78.3	79.6	79.6	80.5	65.8	67.6	67.7

on ImageNet-1K. We further adopt the asymmetric augmentation policy from BYOL and improve DCL from 67.8% to 68.2% top-1 accuracy on ImageNet-1K.

DCL Ablations on CIFAR. Further experiments are conducted based on the ResNet-50 backbone and large learning epochs (i.e., 500 epochs). The DCL model with kNN eval, batch size 32, and 500 epochs of training could reach 86.1% compared to 82.2%. For the following experiments in Table 4, we show DCL ResNet-50 performance on CIFAR10 and CIFAR100. In these comparisons, we vary the batch size to show the effectiveness of DCL.

MoCo-v2 with DCL. We are aware that it is more convincing to compare the proposed DCL against a more compelling version, MoCo-v2. Comparisons on both ImageNet-1K and ImageNet-100 in Table 5 indicate that DCL becomes significantly more effective than MoCo-v2 when the queue size gets smaller.

Few Learning Epochs. DCL can alleviate the shortcoming of the traditional contrastive learning framework, which needs a large batch size long learning epochs to achieve higher performance. The previous state-of-the-art, SimCLR, heavily relies on large quantities of learning epochs to obtain high top-1 accuracy. (e.g., 69.3% with up to 1000 epochs). DCL aims to achieve higher learning efficiency with few learning epochs. We demonstrate the effectiveness of DCL in InfoNCE-based frameworks SimCLR and MoCo-v2 [8]. We choose the batch size of 256 (queue of 65536) as the baseline and train the model with only 100 epochs. We make sure other parameter settings are the same for a fair comparison. Table 6 shows the result on ImageNet-1K using linear evaluation. With DCL, SimCLR can achieve 64.6% top-1 accuracy with only 100 epochs compared to SimCLR baseline: 57.5%; MoCo-v2 with DCL reaches 64.4% compared to MoCo-v2 baseline: 63.6% with 100 epochs pre-training.

Table 6. ImageNet-1K top-1 accuracy (%) on SimCLR and MoCo-v2 with/without DCL under few training epochs. We further list results under 200 epochs for clear comparison. With DCL, the performance of SimCLR trained under 100 epochs nearly reaches its performance under 200 epochs. The MoCo-v2 with DCL also reaches higher accuracy than the baseline under 100 epochs.

	SimCLR	SimCLR w/ DCL	MoCo-v2	MoCo-v2 w/ DCL
100 Epoch	57.5	64.6	63.6	64.4
200 Epoch	61.8	65.9	67.5	67.7

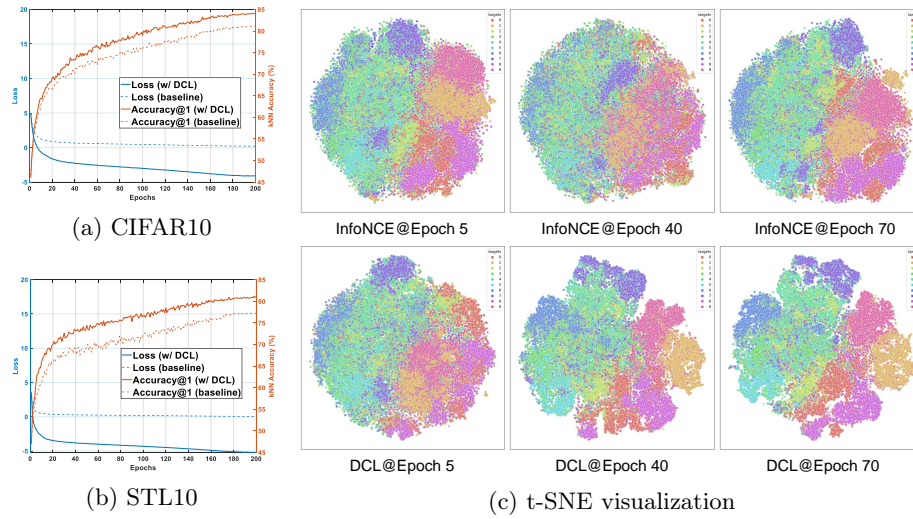


Fig. 4. Comparisons between DCL and InfoNCE-based baseline (SimCLR) on (a) CIFAR10 and (b) STL10 data. DCL speeds up the model convergence during the SSL pre-training and provides better performance than the baseline on CIFAR and STL10 data. (c) t-SNE visualization of CIFAR-10 with 32 batch size. DCL shows a stronger separation force between the features than SimCLR.

We further demonstrate that, with DCL, learning representation becomes faster during the early stage of training compared to the InfoNCE-based learning scheme. The reason is that DCL successfully solves the decoupled issue between positive and negative pairs. Figure 4 on (a) CIFAR10 and (b) STL10 shows that DCL improves the speed of convergence and reaches higher performance than the baseline on CIFAR and STL10 data. The t-SNE visualization in Figure 4 (c) also supports the proposed theoretical derivation that removing the batch-size dependent impact (i.e., NPC multiplier) should improve representation learning abilities over the InfoNCE-based learning scheme.

Table 7. Linear top-1 accuracy (%) comparison of SSL approaches on ImageNet-1K. Given lower computational budget, DCL model are better than recent SOTA approaches. Its effectiveness **does not rely on** large batch size and epochs (SimCLR [7], NNCLR [10]), momentum encoding (BYOL [13], MoCo-v2 [8]), or other tricks such as stop-gradient (SimSiam [9]) and multi-cropping (SwAV [5]).

ResNet-50 w/	SimCLR BYOL SwAV		MoCo-v2	SimSiam	Barlow Twins	NNCLR	NNCLR +DCL
Epoch	400		400		300	1000	400
Batch Size	4096		256		256	256 / 512	256 / 512
ImageNet-1K (Linear)	69.8	73.2 70.7	71.0 70.8		70.7	68.7 / 71.7	71.1 / 72.3

5 Discussion

Comparison with other SOTA SSL Approaches. The primary goal of this work is to provide an efficient and effective improvement to the widely used InfoNCE-based contrastive learning, where we decouple the positive and negative terms to achieve better representation quality. DCL is less sensitive to suboptimal hyperparameters and achieves competitive results with minimal requirements. Its effectiveness does not rely on large batch size and learning epochs, momentum encoding, negative sample queues, or additional tactics (e.g., stop-gradient and multi-cropping). Overall, DCL provides a more robust baseline for the contrastive-based SSL approaches. Though this work aims not to provide a SOTA SSL approach, DCL can be combined with the SOTA contrastive learning methods, such as NNCLR [10], to achieve better performance without large batch size and learning epochs. In Table 7, we provide extensive comparisons to SOTA SSL approaches on ImageNet-1K to validate the effectiveness of DCL. In Table 8, we further show that DCL achieves competitive results compared to VICReg [2], Barlow Twins [34], SimSiam [9], SwAV [4], and DINO [6] on ImageNet-100 and CIFAR-10.

Generalization of DCL to Different Domains. DCL can be easily adapted to different domains (e.g., speech and language models) to achieve competitive performance. We demonstrate that DCL can be combined with SOTA SSL speech models, e.g., wav2vec 2.0 [1] which uses transformer backbone and requires enormous computation resources. We evaluate wav2vec 2.0 on its downstream tasks and perform better by applying the DCL method. Detailed results and discussion can be found in Appendix. To the best of our knowledge, DCL can be potentially combined with a transformer-based language model, CLIP [23], which uses a very large batch size of 32768. With DCL, CLIP shall maintain its complexity and achieve huge learning efficiency when the batch size becomes smaller. Note that it has been implemented by [28].

DCL Convergence for Large Batch Sizes. The performance of DCL appears to have less gain compared to InfoNCE-based baseline when the batch size is large. According to Figure 1 and the theoretical analysis, the reason is that the NPC multiplier $q_B \rightarrow 0$ when the batch size is large (e.g., 1024). As shown in the analysis, InfoNCE loss converges to the DCL loss as the batch size approaches

Table 8. kNN & linear top-1 accuracy (%) comparison of SSL approaches on CIFAR10 and ImageNet-100.

ResNet-18 @ 256 Batch Size	DINO	SwAV	SimSiam	VICReg	Barlow	Twins	NNCLR	NNCLR+DCL
CIFAR-10, 1000 Epoch (kNN)	89.5	89.2	90.5	92.1	92.1	91.8	92.3	
ImageNet-100, 400 Epoch (Linear)	74.9	74.0	74.5	79.2	80.2	79.8	80.6	

Table 9. Results of DCL and SimCLR with large batch size and learning epochs.

ImageNet-1K (ResNet-50)	Batch Size	Epoch	Top-1 Accuracy (%)
SimCLR	256	200	61.8
SimCLR	256	400	64.8
SimCLR	1024	400	67.3
SimCLR w/ DCL	256	200	67.8 (+ 6.0)
SimCLR w/ DCL	256	400	69.5 (+ 4.7)
SimCLR w/ DCL	1024	400	69.9 (+ 2.6)

infinity. With 400 training epochs, the ImageNet-1K top-1 accuracy slightly increases from 69.5% to 69.9% when the batch size increases from 256 to 1024. Please refer to Table 9.

6 Conclusion

This paper identifies the negative-positive-coupling (NPC) effect in the widely used InfoNCE loss, making the SSL task significantly easier to solve with smaller batch size. By removing the NPC effect, we reach a new objective function, *decoupled contrastive learning* (DCL). The proposed DCL loss function requires minimal modification to the SimCLR baseline and provides efficient, reliable, and nontrivial performance improvement on various benchmarks. Given the conceptual simplicity of DCL and that it requires neither momentum encoding, large batch size, or long epochs to reach competitive performance. Notably, DCL can be combined with the SOTA contrastive learning method, NNCLR, to achieve 72.3% ImageNet-1K top-1 accuracy with 512 batch size in 400 epochs. We wish that DCL can serve as a strong baseline for the contrastive-based SSL methods. Further, an important lesson from the DCL loss is that a more efficient SSL task shall maintain its complexity when the batch size becomes smaller.

Acknowledgements. This work was supported in part by the MOST grants 110-2634-F-007-027, 110-2221-E-001-017 and 111-2221-E-001-015 of Taiwan. We are grateful to National Center for High-performance Computing and Meta AI Research for providing computational resources and facilities.

References

1. Baevski, A., Zhou, Y., Mohamed, A., Auli, M.: wav2vec 2.0: A framework for self-supervised learning of speech representations. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2020) 13
2. Bardes, A., Ponce, J., LeCun, Y.: Vicreg: Variance-invariance-covariance regularization for self-supervised learning. *CoRR* **abs/2105.04906** (2021) 13
3. Belghazi, M.I., Baratin, A., Rajeswar, S., Ozair, S., Bengio, Y., Hjelm, R.D., Courville, A.C.: Mutual information neural estimation. In: *Proceedings of the International Conference on Machine Learning (ICML)* (2018) 5
4. Caron, M., Bojanowski, P., Joulin, A., Douze, M.: Deep clustering for unsupervised learning of visual features. In: *European Conference on Computer Vision (ECCV)* (2018) 13
5. Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A.: Unsupervised learning of visual features by contrasting cluster assignments. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2020) 13
6. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. *CoRR* **abs/2104.14294** (2021) 13
7. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.E.: A simple framework for contrastive learning of visual representations. In: *Proceedings of the International Conference on Machine Learning (ICML)* (2020) 2, 4, 7, 13
8. Chen, X., Fan, H., Girshick, R.B., He, K.: Improved baselines with momentum contrastive learning. *CoRR* **abs/2003.04297** (2020) 11, 13
9. Chen, X., He, K.: Exploring simple siamese representation learning. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2021) 4, 13
10. Dwivedi, D., Aytar, Y., Tompson, J., Sermanet, P., Zisserman, A.: With a little help from my friends: Nearest-neighbor contrastive learning of visual representations. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 9588–9597 (2021) 3, 13
11. Ermolov, A., Siarohin, A., Sangineto, E., Sebe, N.: Whitening for self-supervised representation learning. In: *International Conference on Machine Learning (ICML)* (2021) 5
12. Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. In: *International Conference on Learning Representations (ICLR)* (2018) 2
13. Grill, J., Strub, F., Altché, F., Tallec, C., Richemond, P.H., Buchatskaya, E., Doersch, C., Pires, B.Á., Guo, Z., Azar, M.G., Piot, B., Kavukcuoglu, K., Munos, R., Valko, M.: Bootstrap your own latent - A new approach to self-supervised learning. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2020) 4, 10, 13
14. Hadsell, R., Chopra, S., LeCun, Y.: Dimensionality reduction by learning an invariant mapping. In: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)* (2006) 4
15. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2020) 2, 4, 8
16. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016) 7, 8

17. Hjelm, R.D., Fedorov, A., Lavoie-Marchildon, S., Grewal, K., Bachman, P., Trischler, A., Bengio, Y.: Learning deep representations by mutual information estimation and maximization. In: International Conference on Learning Representations (ICLR) (2019) 5
18. Kalantidis, Y., Sariyildiz, M.B., Pion, N., Weinzaepfel, P., Larlus, D.: Hard negative mixing for contrastive learning. In: Advances in Neural Information Processing Systems (NeurIPS) (2020) 5
19. Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., Krishnan, D.: Supervised contrastive learning. In: Advances in Neural Information Processing Systems (NeurIPS) (2020) 4
20. Noroozi, M., Favaro, P.: Unsupervised learning of visual representations by solving jigsaw puzzles. In: European Conference on Computer Vision (ECCV) (2016) 2
21. van den Oord, A., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. CoRR **abs/1807.03748** (2018) 2, 4
22. Ozair, S., Lynch, C., Bengio, Y., van den Oord, A., Levine, S., Sermanet, P.: Wasserstein dependency measure for representation learning. In: Advances in Neural Information Processing Systems (NeurIPS) (2019) 5
23. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision. In: Meila, M., Zhang, T. (eds.) Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event. Proceedings of Machine Learning Research, vol. 139, pp. 8748–8763. PMLR (2021) 13
24. Ren, H.: A pytorch implementation of simclr. <https://github.com/leftthomas/SimCLR> (2020) 8
25. Robinson, J.D., Chuang, C., Sra, S., Jegelka, S.: Contrastive learning with hard negative samples. In: International Conference on Learning Representations (ICLR) (2021) 5
26. Tian, Y., Krishnan, D., Isola, P.: Contrastive multiview coding. In: European Conference on Computer Vision (ECCV) (2020) 2
27. Tsai, Y.H., Ma, M.Q., Yang, M., Zhao, H., Morency, L., Salakhutdinov, R.: Self-supervised representation learning with relative predictive coding. In: International Conference on Learning Representations (ICLR) (2021) 5
28. Wang, P.: x-clip. <https://github.com/lucidrains/x-clip> (2021) 13
29. Wang, T., Isola, P.: Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In: International Conference on Machine Learning (ICML) (2020) 6
30. Wang, X., Liu, Z., Yu, S.X.: Unsupervised feature learning by cross-level instance-group discrimination. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) 8
31. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via non-parametric instance discrimination. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018) 2, 4
32. Ye, M., Zhang, X., Yuen, P.C., Chang, S.F.: Unsupervised embedding learning via invariant and spreading instance feature. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019) 4
33. You, Y., Gitman, I., Ginsburg, B.: Large batch training of convolutional networks. arXiv preprint arXiv:1708.03888 (2017) 4
34. Zbontar, J., Jing, L., Misra, I., LeCun, Y., Deny, S.: Barlow twins: Self-supervised learning via redundancy reduction. In: International Conference on Machine Learning. pp. 12310–12320. PMLR (2021) 13

- 35. Zhan, X., Xie, J., Liu, Z., Lin, D., Change Loy, C.: OpenSelfSup: Open mmlab self-supervised learning toolbox and benchmark. <https://github.com/open-mmlab/openselfsup> (2020) 8
- 36. Zhang, R., Isola, P., Efros, A.A.: Colorful image colorization. In: European Conference on Computer Vision (ECCV) (2016) 2
- 37. Zhu, B., Huang, J., Li, Z., Zhang, X., Sun, J.: Eqco: Equivalent rules for self-supervised contrastive learning. arXiv preprint arXiv:2010.01929 (2020) 5