

# Learning Decorrelated Representations Efficiently Using Fast Fourier Transform

Yutaro Shigeto\* Masashi Shimbo\* Yuya Yoshikawa Akikazu Takeuchi  
 {shigeto, shimbo, yoshikawa, takeuchi}@stair.center  
 STAIR Lab, Chiba Institute of Technology, Narashino, Chiba, Japan

## Abstract

*Barlow Twins and VICReg are self-supervised representation learning models that use regularizers to decorrelate features. Although these models are as effective as conventional representation learning models, their training can be computationally demanding if the dimension  $d$  of the projected embeddings is high. As the regularizers are defined in terms of individual elements of a cross-correlation or covariance matrix, computing the loss for  $n$  samples takes  $O(nd^2)$  time. In this paper, we propose a relaxed decorrelating regularizer that can be computed in  $O(nd \log d)$  time by Fast Fourier Transform. We also propose an inexpensive technique to mitigate undesirable local minima that develop with the relaxation. The proposed regularizer exhibits accuracy comparable to that of existing regularizers in downstream tasks, whereas their training requires less memory and is faster for large  $d$ . The source code is available.<sup>1</sup>*

## 1. Introduction

Self-supervised learning (SSL) of representations [1, 3–5, 10, 13, 15, 27–30] has become an integral part of deep learning applications in computer vision. Most SSL models for visual representations employ a multi-view, Siamese network architecture. First, an input image is altered by different transformations that preserve its original semantics. The two augmented examples are then fed to a neural network (or in some cases, two different networks) that consists of a backbone network cascaded with a small projection network (usually a multi-layer perceptron) to produce “twin” projected embeddings, or views, of the original. Finally, the network weights are trained so that the twin embeddings (referred to as a “positive pair”) are similar, reflecting the fact that they represent the same original image. After training, the projection network is discarded, while the backbone network is reused for down-

stream tasks; the idea is that the pretrained backbone should produce generic representations of images that also benefit downstream tasks.

One major issue in SSL is to sidestep *collapsed embeddings*, or the presence of trivial solutions such that all examples are projected to a single embedding vector. Contrastive approaches [3, 4, 15, 27, 28] eliminate such solutions by a loss term that repels embeddings of different original images (known as “negative pairs”) from each other. Considering all possible negative pairs is infeasible, and negative sampling is usually performed.

Recent studies have explored non-contrastive SSL models. Among these models, Barlow Twins [29] and VICReg [1] use loss functions to penalize features (components of embedding vectors) with a small variance, which are characteristic of collapsed embeddings. Also notable in these loss functions are regularization terms for feature decorrelation. They reduce redundancy in the learned features and make Barlow Twins and VICReg perform as well as contrastive models. However, since the regularizers are defined in terms of elements in a sample cross-correlation or covariance matrix, they require  $O(nd^2)$  time to compute, where  $n$  is the number of examples in a batch, and  $d$  is the dimensionality of the projected embeddings. This is problematic, as an increased  $d$  has been reported to improve the performance of both Barlow Twins and VICReg [1, 29].

**Contributions.** We propose a relaxed decorrelating regularizer to address the inefficiency of Barlow Twins and VICReg. The proposed regularizer does not require the explicit calculation of cross-correlation or covariance matrices and can be computed in  $O(nd \log d)$  time by means of Fast Fourier Transform (FFT) [7]. Although undesirable local minima develop as a result of relaxation, they can be mitigated by feature permutation; we give an account of why this simple technique works. The SSL models using the proposed regularizer achieve competitive performance with Barlow Twins and VICReg in downstream tasks, with substantially less computation time for large  $d$ . With  $d = 8192$ , training is 1.2 (with a ResNet-50 backbone) or 2.2 (with a

\*Equal contribution.

<sup>1</sup><https://github.com/yutaro-s/scalable-decorrelation-ssl.git>

lightweight ResNet-18 backbone) times as fast as Barlow Twins. The proposed method also reduces memory consumption, which allows for larger batch size.

**Notation.** We use zero-based indexing for vector and matrix components unless stated otherwise; thus, for a vector  $\mathbf{x} \in \mathbb{R}^d$ , the component index ranges from 0 to  $d-1$ . We use  $[\mathbf{x}]_i$  to denote the  $i$ th component of vector  $\mathbf{x}$ , and  $[\mathbf{M}]_{ij}$  to denote the  $(i, j)$ th element of matrix  $\mathbf{M}$ . For a complex vector  $\mathbf{c}$ ,  $\bar{\mathbf{c}}$  denotes its componentwise complex conjugate. For vectors  $\mathbf{x}$  and  $\mathbf{y}$ ,  $\mathbf{x} \circ \mathbf{y}$  denotes their componentwise product.

## 2. Related Work

**Contrastive SSL.** Contrastive SSL uses positive and negative pairs of augmented samples [3, 4, 6, 9, 15, 24, 27, 28]. The commonly used InfoNCE loss [27] consists of an alignment term, which maximizes the similarity between positive pairs, and a uniformity term, which minimizes the similarity between negative pairs [28]. SimCLR [3] is a state-of-the-art contrastive SSL method. However, to obtain effective representations, SimCLR requires a large number of negative pairs [3], or, in other words, a large batch (or memory bank) size  $n$ . This can be a computational bottleneck, as the loss computation of SimCLR takes  $O(n^2d)$  time where  $d$  is the dimensionality of the projected embeddings.

**Non-contrastive SSL Using Asymmetric Architecture.** Recently, researchers have started exploring non-contrastive approaches to SSL, i.e., those that do not use negative pairs for training. To overcome collapsed embeddings, models such as BYOL [13] and SimSiam [5] introduce asymmetry in the architecture, e.g., by suppressing gradient updates and/or using the moving average of network parameters for one branch of the Siamese network. These methods are heuristically motivated, as collapsed embeddings are not explicitly penalized; however, they are effective in practice.

**Non-contrastive SSL by Decorrelating Regularization.** A line of work exists that maintains a standard (symmetric) Siamese network but introduces loss functions to suppress collapsed embeddings. These loss functions also have regularization terms to promote feature decorrelation. Barlow Twins [29] is the first method in this line. It uses a decorrelating regularizer based on cross-correlation across two views. VICReg [1] uses regularizers that are defined in terms of covariance matrices of individual views. We review these methods in detail in Sec. 3.

**Non-contrastive SSL by Whitening.** Some authors [10, 18] have used whitening to explicitly decorrelate features during training, as opposed to performing regularization.

Subsequent work [30] whitened both features and instances. Because the whitening procedures used in these approaches require the computation of all the eigenvalues of the covariance matrices, a training epoch takes  $O(\min(dn^2, nd^2))$  time, which is inefficient with large  $d$  or  $n$ .

**Use of Convolution in Machine Learning.** Convolution is the fundamental building block of convolution neural networks (CNNs). CNNs take (linear) convolution of input vectors with small learnable kernels to extract local features. Although FFT reduces the asymptotic complexity of convolution computation, it is seldom used with CNNs, because the size of kernels is typically too small to benefit from speed-up by FFT. In contrast to CNNs, we use (circular) convolution to compute summary statistics of covariance and cross-correlation matrices.

In other areas of machine learning, circular convolution and its non-commutative analogue, circular correlation, have been used for implementing associative memory [2, 20, 22]. The idea has recently been applied to knowledge graph embeddings (KGEs) [19], with the resulting model later shown [14] to be isomorphic to complex-valued KGEs [25] by means of the convolution theorem.

## 3. SSL Using Decorrelating Regularizers

This section reviews Barlow Twins [29] and VICReg [1]. Given a batch of  $n$  original training examples, these models apply two randomly chosen transformations to each example. The transformed examples are then processed independently by a neural network to produce twin embeddings (views) of the original. Let  $\mathbf{a}^{(k)}, \mathbf{b}^{(k)} \in \mathbb{R}^d$  denote the twin embeddings thus obtained for the  $k$ th example, and let  $\mathbf{A} = \{\mathbf{a}^{(k)}\}_{k=1}^n, \mathbf{B} = \{\mathbf{b}^{(k)}\}_{k=1}^n$  be the sets of embeddings for individual views. Given these embeddings, the network is trained to minimize the loss function specific to each model.

**Barlow Twins.** Barlow Twins optimizes the loss function defined in terms of the cross-correlation matrix  $\mathbf{C}(\mathbf{A}, \mathbf{B}) \in \mathbb{R}^{d \times d}$  between views  $\mathbf{A}$  and  $\mathbf{B}$ :

$$L_{\text{BT}} = \sum_{i=0}^{d-1} (1 - [\mathbf{C}(\mathbf{A}, \mathbf{B})]_{ii})^2 + \lambda R_{\text{off}}(\mathbf{C}(\mathbf{A}, \mathbf{B})), \quad (1)$$

where hyperparameter  $\lambda \geq 0$  controls the strength of regularization, and  $R_{\text{off}} : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}$  is a regularizer function defined as

$$R_{\text{off}}(\mathbf{M}) = \sum_{i=0}^{d-1} \sum_{\substack{j=0 \\ j \neq i}}^{d-1} [\mathbf{M}]_{ij}^2. \quad (2)$$

The first term in Eq. (1) is minimized when the corresponding features of two views are fully correlated, i.e.,

$[\mathbf{C}(A, B)]_{ii} = 1$  for  $i = 0, \dots, d - 1$ . This term can be efficiently computed in  $O(nd)$  time. Regularizer  $R_{\text{off}}(\mathbf{C}(A, B))$  in the second term is responsible for feature decorrelation, as it is minimized when all the off-diagonal elements in  $\mathbf{C}(A, B)$  are zero. For a large  $d$ ,  $R_{\text{off}}(\mathbf{C}(A, B))$  can be a computational burden, as it takes  $O(nd^2)$  time to compute due to the  $d \times d$  matrix  $\mathbf{C}(A, B)$ .

**VICReg.** Let  $\mathbf{K}(A), \mathbf{K}(B) \in \mathbb{R}^{d \times d}$  be the covariance matrices of  $A$  and  $B$ , respectively. The VICReg loss function is

$$L_{\text{VIC}} = \frac{\alpha}{n} \sum_{k=1}^n \|\mathbf{a}^{(k)} - \mathbf{b}^{(k)}\|_2^2 + \frac{\mu}{d} (R_{\text{var}}(\mathbf{K}(A)) + R_{\text{var}}(\mathbf{K}(B))) + \frac{\nu}{d} (R_{\text{off}}(\mathbf{K}(A)) + R_{\text{off}}(\mathbf{K}(B))), \quad (3)$$

where hyperparameters  $\alpha, \mu, \nu \geq 0$  control the importance of individual terms, and  $R_{\text{var}}$  is the regularizer defined as

$$R_{\text{var}}(\mathbf{M}) = \sum_{i=0}^{d-1} \max(0, \gamma - \sqrt{[\mathbf{M}]_{ii}}), \quad (4)$$

with the target standard deviation  $\gamma > 0$ . Function  $R_{\text{off}}$  is the same regularizer as used in Barlow Twins (Eq. (2)), but here, it is applied to  $\mathbf{K}(A)$  and  $\mathbf{K}(B)$  instead of  $\mathbf{C}(A, B)$ .

The first term in Eq. (3) brings two embeddings of the same example closer. Regularizer  $R_{\text{var}}$  penalizes collapsed embeddings with zero variance, whereas  $R_{\text{off}}$  promotes the diversity of features by encouraging the feature covariance to be zero. The time complexity of calculating  $R_{\text{off}}(\mathbf{K}(A))$  and  $R_{\text{off}}(\mathbf{K}(B))$  is  $O(nd^2)$ , identical to that of  $R_{\text{off}}(\mathbf{C}(A, B))$  in Barlow Twins.

## 4. Proposed Method

We propose a weaker but efficiently computable alternative to the regularizer function  $R_{\text{off}}$ . In the following, our regularizer is presented in terms of cross-correlation matrix  $\mathbf{C}(A, B)$ , similarly to Barlow Twins. However, if applied instead to  $\mathbf{K}(A)$  and  $\mathbf{K}(B)$ , it produces a relaxed version of the corresponding regularization terms in the VICReg loss.

### 4.1. Regularizer Based on Sums of Cross-correlations

Recall that the Barlow Twins loss is based on the cross-correlation matrix  $\mathbf{C} = \mathbf{C}(A, B)$  of two views  $A = \{\mathbf{a}^{(k)}\}_{k=1}^n$  and  $B = \{\mathbf{b}^{(k)}\}_{k=1}^n$ . For brevity, assume that both  $A$  and  $B$  are standardized. Then, we can simply write  $\mathbf{C} = (1/(n - 1)) \sum_{k=1}^n \mathbf{a}^{(k)} \mathbf{b}^{(k)\top}$ .




Figure 1. A  $3 \times 3$  cross-correlation matrix  $\mathbf{C} = [c_{ij}]$  ( $i, j = 0, 1, 2$ ) and  $\text{sumvec}(\mathbf{C}) = [v_0 \ v_1 \ v_2]^\top$ .

Our regularizer is defined in terms of a  $d$ -dimensional “summary” vector of the  $d \times d$  matrix  $\mathbf{C}$ . This vector, denoted by  $\text{sumvec}(\mathbf{C})$ , is given componentwise by

$$[\text{sumvec}(\mathbf{C})]_i = \sum_{j=0}^{d-1} [\mathbf{C}]_{j, (i+j) \bmod d}. \quad (5)$$

Note the zero-based component indices. The zeroth component  $[\text{sumvec}(\mathbf{C})]_0$  is the trace of  $\mathbf{C}$ . Each of the remaining  $d - 1$  components corresponds to a sum of  $d$  different off-diagonal elements of  $\mathbf{C}$ , with no single element appearing in two distinct sums. Thus, every element in  $\mathbf{C}$  appears exactly once in the summations in Eq. (5). The calculation of a summary vector for a  $3 \times 3$  covariance matrix is illustrated in Fig. 1.

Now, we define a regularizer in terms of all but the zeroth component of  $\text{sumvec}(\mathbf{C})$ :

$$R_{\text{sum}}(\mathbf{C}) = \sum_{i=1}^{d-1} \|[\text{sumvec}(\mathbf{C})]_i\|_q^q, \quad (6)$$

where hyperparameter  $q \in \{1, 2\}$ . This function  $R_{\text{sum}}$  can be used as a drop-in replacement for  $R_{\text{off}}$  in the Barlow Twins loss (Eq. (1)). The zeroth component  $[\text{sumvec}(\mathbf{C})]_0$  is excluded in Eq. (6), because it concerns the diagonal elements of  $\mathbf{C}$  that are irrelevant to feature decorrelation; they do not appear in Barlow Twin’s regularizer  $R_{\text{off}}(\mathbf{C})$  either.

The regularizer  $R_{\text{sum}}$  is weaker than  $R_{\text{off}}$  in that it imposes constraints on the components of the summary vector, or the sums of  $d$  elements of  $\mathbf{C}$ , whereas  $R_{\text{off}}$  constrains individual elements. Indeed, the minimizers of  $R_{\text{off}}(\mathbf{C})$  also minimize  $R_{\text{sum}}(\mathbf{C})$ , but the converse does not necessarily hold. However, as we discuss in Sec. 4.2,  $R_{\text{sum}}$  allows faster computation. Furthermore, in Sec. 4.3, we provide a simple technique to mitigate the weakness of  $R_{\text{sum}}$ .

### 4.2. Efficient Computation

Computing  $\text{sumvec}(\mathbf{C})$  by Eq. (5) requires the cross-correlation matrix  $\mathbf{C}(A, B)$ , whose calculation incurs the same computational cost as that of the Barlow Twins loss. However, by means of the Fourier transform,  $\text{sumvec}(\mathbf{C})$

can be calculated directly from the vectors in A and B without explicit calculation of their cross-correlation matrix. To this end, we use involution and circular convolution.

For a vector  $\mathbf{x} \in \mathbb{R}^d$ , its *involution* [22] (also called *flipping* [23])  $\text{inv}(\mathbf{x})$  is the vector obtained by reversing the order of its first (not the zeroth) to  $(d-1)$ st components; i.e.,  $[\text{inv}(\mathbf{x})]_i = [\mathbf{x}]_{(d-i) \bmod d}$  for  $i = 0, \dots, d-1$ .

For vectors  $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ , their *circular convolution*  $\mathbf{x} * \mathbf{y}$  is a  $d$ -dimensional vector with components given as

$$[\mathbf{x} * \mathbf{y}]_i = \sum_{j=0}^{d-1} [\mathbf{x}\mathbf{y}^T]_{j, (i-j) \bmod d}. \quad (7)$$

As a result, circular convolution is known as the ‘‘compressed outer product.’’

Now, for each twin embedding pair  $\mathbf{a}^{(k)} \in \mathbf{A}$  and  $\mathbf{b}^{(k)} \in \mathbf{B}$  ( $k = 1, \dots, n$ ), let us consider vector  $\text{inv}(\mathbf{a}^{(k)}) * \mathbf{b}^{(k)} \in \mathbb{R}^d$ .<sup>2</sup> Noting the indices altered by involution, we see that this vector is given componentwise by

$$\left[ \text{inv}(\mathbf{a}^{(k)}) * \mathbf{b}^{(k)} \right]_i = \sum_{j=0}^{d-1} \left[ \mathbf{a}^{(k)} \mathbf{b}^{(k)T} \right]_{j, (i+j) \bmod d}. \quad (8)$$

Substituting  $\mathbf{C} = (1/(n-1)) \sum_{k=1}^n \mathbf{a}^{(k)} \mathbf{b}^{(k)T}$  into Eq. (5) and using Eq. (8), we have

$$\begin{aligned} [\text{sumvec}(\mathbf{C})]_i &= \sum_{j=0}^{d-1} \left[ \overbrace{\frac{1}{n-1} \sum_{k=1}^n \mathbf{a}^{(k)} \mathbf{b}^{(k)T}}^{\mathbf{C}} \right]_{j, (i+j) \bmod d} \\ &= \frac{1}{n-1} \sum_{k=1}^n \sum_{j=0}^{d-1} \left[ \mathbf{a}^{(k)} \mathbf{b}^{(k)T} \right]_{j, (i+j) \bmod d} \\ &= \frac{1}{n-1} \sum_{k=1}^n \left[ \text{inv}(\mathbf{a}^{(k)}) * \mathbf{b}^{(k)} \right]_i, \end{aligned} \quad (9)$$

or, as a vector,

$$\text{sumvec}(\mathbf{C}) = \frac{1}{n-1} \sum_{k=1}^n \text{inv}(\mathbf{a}^{(k)}) * \mathbf{b}^{(k)}. \quad (10)$$

Let  $\mathcal{F}$  and  $\mathcal{F}^{-1}$  denote the discrete Fourier transform and inverse discrete Fourier transform, respectively. Noting that  $\mathcal{F}(\text{inv}(\mathbf{x})) = \overline{\mathcal{F}(\mathbf{x})}$  for any  $\mathbf{x} \in \mathbb{R}^d$  (see e.g., [23], Section 7.4.2) and using the convolution theorem  $\mathcal{F}(\mathbf{x} * \mathbf{y}) = \mathcal{F}(\mathbf{x}) \circ \mathcal{F}(\mathbf{y})$ , we have

$$\text{inv}(\mathbf{a}^{(k)}) * \mathbf{b}^{(k)} = \mathcal{F}^{-1} \left( \overline{\mathcal{F}(\mathbf{a}^{(k)})} \circ \mathcal{F}(\mathbf{b}^{(k)}) \right). \quad (11)$$

<sup>2</sup>The vector  $\text{inv}(\mathbf{x}) * \mathbf{y}$  is known as the *circular (cross-)correlation* of  $\mathbf{x}$  and  $\mathbf{y}$  [20, 22, 23]. We opt not to use this term in this paper to avoid confusion with the cross-correlation of random vectors, which is used in Barlow Twins.

Substituting Eq. (11) into Eq. (10), we obtain

$$\begin{aligned} \text{sumvec}(\mathbf{C}) &= \frac{1}{n-1} \sum_{k=1}^n \overbrace{\mathcal{F}^{-1} \left( \overline{\mathcal{F}(\mathbf{a}^{(k)})} \circ \mathcal{F}(\mathbf{b}^{(k)}) \right)}^{\text{inv}(\mathbf{a}^{(k)}) * \mathbf{b}^{(k)}} \\ &= \frac{1}{n-1} \mathcal{F}^{-1} \left( \sum_{k=1}^n \overline{\mathcal{F}(\mathbf{a}^{(k)})} \circ \mathcal{F}(\mathbf{b}^{(k)}) \right). \end{aligned} \quad (12)$$

Using this equation, we can compute  $\text{sumvec}(\mathbf{C})$  directly from the embedding vectors  $\{\mathbf{a}^{(k)}, \mathbf{b}^{(k)}\}_{k=1}^n$ , bypassing the cumbersome calculation of  $\mathbf{C}$ . Noting that the (inverse) Fourier transform of a  $d$ -dimensional vector can be performed in  $O(d \log d)$  time by the FFT algorithm, while the calculation of complex conjugates, component products, and the sum of  $n$  vectors takes  $O(nd)$  time, we can see that the overall time to obtain  $\text{sumvec}(\mathbf{C})$  is  $O(nd \log d)$ , which is also the time needed to compute  $R_{\text{sum}}(\mathbf{C})$ . This is a substantial improvement over the  $O(nd^2)$  computation time of  $R_{\text{off}}(\mathbf{C})$  in the Barlow Twins loss.

Furthermore, the space complexity of computing  $R_{\text{sum}}(\mathbf{C})$  is  $O(nd)$ , which is optimal if the  $O(nd)$  space needed to store input vectors A and B is considered as part of the complexity. In contrast, Barlow Twins requires extra  $O(d^2)$  space to store  $\mathbf{C}$ .

### 4.3. Feature Permutation to Mitigate Undesirable Local Minima

As seen from Eq. (5), the components of  $\text{sumvec}(\mathbf{C})$  are the sums of  $d$  elements in  $\mathbf{C}$ , and the proposed regularizer  $R_{\text{sum}}$  (Eq. (6)) encourages these sums to be close to zero. This is weaker than the regularizer  $R_{\text{off}}$  (Eq. (2)) in the Barlow Twins loss, which pushes individual elements of  $\mathbf{C}$  towards zero. Indeed,  $R_{\text{sum}}(\mathbf{C})$  can be close to zero even if individual elements in  $\mathbf{C}$  are not; the summands in Eq. (5) can cancel each other, since they can be either positive or negative. As a result, undesirable local minima develop in the parameter space, rendering our regularizer ineffective.

Here, we propose a technique to eliminate these local minima: we randomly permute feature indices during training so that sums of different cross-correlation terms constitute  $\text{sumvec}(\mathbf{C})$ . To understand why this simple technique is effective, consider minimizing  $R_{\text{sum}}(\mathbf{C})$ , regarding the elements of  $\mathbf{C}$  as independent<sup>3</sup> variables. It is easy to see that the minimum is attained by the solutions to a homogeneous system of linear equations:

$$\overbrace{\sum_{j=0}^{d-1} [\mathbf{C}]_{j, (i+j) \bmod d}}^{[\text{sumvec}(\mathbf{C})]_i} = 0, \quad \text{for } i = 1, \dots, d-1.$$

<sup>3</sup>This assumption is not unrealistic given the high capacity of modern neural networks.



This is an underdetermined system, with only  $d - 1$  equations but with  $d(d - 1)$  unknowns, namely,  $[\mathbf{C}]_{j\ell}$ ,  $j, \ell = 0, \dots, d - 1, j \neq \ell$ . This underdeterminacy is the cause of nontrivial solutions such that  $[\mathbf{C}]_{j\ell} \neq 0$ , that is, undesirable solutions in which summands with opposite signs cancel each other in an equation.

Now, by repeatedly permuting the feature indices and minimizing the loss, we effectively introduce additional equations to the system, since permutation can produce different sets of linear equations over the unknowns, and these new constraints eventually render nontrivial solutions inadmissible.

For ease of implementation, we permute feature indices randomly during training, instead of generating all permutations systematically at once, similarly to the way that stochastic gradient descent is a modification to full gradient descent. Note that the permuted feature indices need not be identical across mini-batches, even within a single epoch; indeed, in the experiments presented in Sec. 5, we use a different random permutation of features in every mini-batch in every epoch.

#### 4.4. Feature Grouping to Control the Degree of Relaxation

Instead of computing a summary vector for an entire cross-correlation matrix  $\mathbf{C}$ , we can compute summaries at a more fine-grained level. Specifically, we partition  $d$  features into groups of size  $b$  each<sup>4</sup>. This partitioning induces in  $\mathbf{C}$  a total of  $\lceil d/b \rceil^2$  block submatrices of size  $b \times b$ , i.e.,  $\mathbf{C} = [\mathbf{C}_{ij}]$  ( $i, j = 1, \dots, \lceil d/b \rceil$ ) with submatrices  $\mathbf{C}_{ij} \in \mathbb{R}^{b \times b}$ . We then define the regularizer as

$$R_{\text{sum}}^{(b)}(\mathbf{C}) = \sum_{i=1}^{\lceil d/b \rceil} \sum_{\ell=1}^{b-1} \|\text{sumvec}(\mathbf{C}_{i\ell})\|_q^q + \sum_{\substack{i,j=1 \\ i \neq j}}^{\lceil d/b \rceil} \sum_{\ell=0}^{b-1} \|\text{sumvec}(\mathbf{C}_{ij})\|_q^q. \quad (13)$$

As before,  $\text{sumvec}(\mathbf{C}_{ij})$  can be computed without explicitly computing  $\mathbf{C}_{ij}$  by means of involution, circular convolution (of subvectors of embeddings), and the Fourier transform. Calculating a single  $\text{sumvec}(\mathbf{C}_{ij})$  takes  $O(nb \log b)$  time using FFT, and since there are  $\lceil d/b \rceil^2$  blocks, the total time needed to compute  $R_{\text{sum}}^{(b)}$  is  $O((nd^2/b) \log b)$ .

The block size hyperparameter  $b$  controls the granularity of the summary computation and interpolates the proposed and existing regularizers. On one hand, the regularizer  $R_{\text{sum}}^{(1)}(\mathbf{C})$  reduces  $R_{\text{off}}(\mathbf{C})$  of Barlow Twins when  $b = 1$ , provided that  $q = 2$ . On the other hand, when  $b = d$ , we

<sup>4</sup>If  $d$  is not divisible by  $b$ , pad dummy features that are constantly 0 in the last group.

recover  $R_{\text{sum}}^{(d)}(\mathbf{C}) = R_{\text{sum}}(\mathbf{C})$  given in Eq. (6). Thus, this grouping formulation provides a generalization of Barlow Twins, with parameter  $b$  controlling the trade-off between the computational efficiency and the degree of relaxed regularization. Empirically, the performance can be slightly improved by the use of a feature group of moderate size, with no substantial degradation observed in training time and memory usage; see Sec. 5.

It should be noted that the permutation and grouping of features are compatible and can be combined.

#### 4.5. Regularizer Based on Sums of Feature Covariances

The function  $R_{\text{sum}}$  can also be used to define a VICReg-style regularizer based on covariance, simply by replacing  $R_{\text{off}}$  with  $R_{\text{sum}}$  in Eq. (3), and passing correlation matrices  $\mathbf{K}(\mathbf{A})$  or  $\mathbf{K}(\mathbf{B})$  instead of  $\mathbf{C}(\mathbf{A}, \mathbf{B})$  as the argument. Fast computation is possible with FFT, and the grouping version is also straightforward; this is described in greater detail in Supplementary Material.

#### 4.6. Summary of Proposed Models

The loss functions of the proposed models are summarized below. Let the function  $R = R_{\text{sum}}^{(b)}$  if feature grouping is used with block size  $b$ , or let  $R = R_{\text{sum}}$  if grouping is not used.

For Barlow Twins-style cross-correlation regularization, the loss function is

$$L = \sum_i (1 - [\mathbf{C}(\mathbf{A}, \mathbf{B})]_{ii})^2 + \lambda R(\mathbf{C}(\mathbf{A}, \mathbf{B})), \quad (14)$$

whereas for VICReg-style covariance regularization, we use

$$L = \frac{\alpha}{n} \sum_i \|\mathbf{a}^{(i)} - \mathbf{b}^{(i)}\|_2^2 + \frac{\mu}{d} (R_{\text{var}}(\mathbf{K}(\mathbf{A})) + R_{\text{var}}(\mathbf{K}(\mathbf{B}))) + \frac{\nu}{d} (R(\mathbf{K}(\mathbf{A})) + R(\mathbf{K}(\mathbf{B}))). \quad (15)$$

Setting  $R = R_{\text{off}}$  in these formulas gives the original Barlow Twins and VICReg, when hyperparameter  $q = 2$ .

## 5. Experiments

We empirically evaluate the effect of the proposed regularizers. To be precise, we train SSL models using the Barlow Twins-style loss function of Eq. (14) and the VICReg-style loss function of Eq. (15) and compare their performance with Barlow Twins and VICReg in downstream tasks. The training time and memory consumption are also evaluated. For our models, feature permutation is performed in every batch iteration except for the ablation study.

In the following, we briefly present the tasks and datasets used in the experiments. See Supplementary Material for the complete experimental setup, including the hyperparameter values and the results of additional experiments.

Table 1. Linear evaluation accuracy (%) on ImageNet-100 with  $d = 2048$ . Bold numbers indicate the best performance within each family (cross-correlation regularization, covariance regularization, or other SSL models). †: quoted from the solo-learn [26] GitHub repository as of December 28, 2022; ‡: quoted from [30].

Model	Top-1	Top-5
Barlow Twins†	80.16	95.14
Barlow Twins	80.12	<b>95.24</b>
Proposed (Barlow Twins-style; no grouping)	79.94	94.76
Proposed (Barlow Twins-style; $b = 128$ )	<b>81.02</b>	<b>95.24</b>
VICReg†	79.40	<b>95.02</b>
VICReg	79.30	94.30
Proposed (VICReg-style; no grouping)	79.20	94.96
Proposed (VICReg-style; $b = 128$ )	<b>80.04</b>	94.98
W-MSE† [10]	69.06	91.22
Zero-FCL‡ [30]	79.32	94.94
Zero-CL‡ [30]	79.26	94.98
NNCLR† [9]	80.16	<b>95.30</b>
BYOL† [13]	80.32	94.94
MoCo V3† [6]	<b>80.36</b>	94.96

## 5.1. Tasks and Datasets

Models are pretrained with images in the ImageNet dataset [8] or its subset, ImageNet-100 [24], depending on the experiment. ResNet-50 [16] is used as the backbone for ImageNet, while ResNet-18 is used for ImageNet-100.

To evaluate downstream SSL performance, we follow the standard linear evaluation protocol: After the backbone network is pretrained by a SSL method, we train a linear classifier on top of the frozen backbone using labeled data from the ImageNet or ImageNet-100 training set. The resulting classifier is then evaluated by the top-1 and top-5 accuracy on the respective validation sets.

For transfer learning evaluation, we apply the pretrained models to an object detection task on Pascal VOC07+12 [11]. Following previous studies [1, 15, 29], we use the trainval set of VOC2007 and VOC2012 for training and the test split of VOC2007 for testing. We fine-tune Faster R-CNN [21] with R50-C4. Models are evaluated by three types of average precision (AP): AP, AP<sub>50</sub>, and AP<sub>75</sub>, where AP<sub>x</sub> signifies the intersection-over-union (IoU) threshold of  $x$  %. We report the average scores over five trials.

## 5.2. Results and Discussion

**Linear evaluation on ImageNet-100.** Tab. 1 presents the results. We can see that the accuracy of the proposed models is comparable to that of all existing models in the table, including Barlow Twins and VICReg, with or without feature grouping.

**Linear evaluation on ImageNet.** Tab. 2 presents the results. Due to the high computational cost of this large-scale

Table 2. Linear evaluation accuracy (%) on ImageNet; highest accuracy over 100 epochs of linear head training.  $d = 8192$  for the proposed model, Barlow Twins, and VICReg. †: quoted from the original papers of the respective methods; ‡: quoted from the MoCo V3 GitHub repository.

Model	Epochs	Top-1	Top-5
Barlow Twins†	1000	<b>73.2</b>	91.0
Barlow Twins	1000	72.4	90.6
Proposed (Barlow Twins-style; no grouping)	1000	73.0	91.2
Proposed (Barlow Twins-style; $b = 128$ )	1000	<b>73.2</b>	<b>91.3</b>
VICReg†	1000	<b>73.2</b>	<b>91.1</b>
VICReg	1000	72.6	90.9
Proposed (VICReg-style; no grouping)	1000	72.8	<b>91.1</b>
W-MSE 4† [10]	400	72.6	—
Zero-CL† [30]	400	72.6	90.5
SimCLR† [3]	1000	69.3	89.0
NNCLR† [9]	1000	<b>75.4</b>	<b>92.3</b>
BYOL† [13]	1000	74.3	91.6
MoCo V3‡ [6]	1000	74.6	—

Table 3. Results of transfer learning for object detection on VOC07+12. †: quoted from original papers; ‡: quoted from [15].

Model	AP <sub>50</sub>	AP	AP <sub>75</sub>
Supervised‡	81.3	53.5	58.8
Barlow Twins†	<b>82.6</b>	<b>56.8</b>	<b>63.4</b>
Proposed (Barlow Twins-style; no grouping)	82.5	55.0	61.1
VICReg†	<b>82.4</b>	—	—
Proposed (VICReg-style; no grouping)	82.3	56.1	62.1

experiment, we do not evaluate feature grouping with the VICReg-style regularization. The proposed models perform slightly worse than NNCLR, BYOL, and MoCo V3, but have comparable performance to that of Barlow Twins and VICReg.

**Transfer learning evaluation on Pascal VOC object detection.** Tab. 3 presents the results of transfer learning. Again, the proposed models demonstrate performance comparable to that of Barlow Twins and VICReg.

**Training time on ImageNet.** Tab. 4 presents the training time over 1000 epochs on ImageNet. We tested two situations: training using 8 and 4 GPUs. In both situations, we set the per GPU batch size to 128, which leads to the effective batch size of 1024 for 8 GPUs and 512 for 4 GPUs. As indicated in Tab. 4, the accuracy of the proposed model is comparable to that of Barlow Twins, with a noticeable reduction in training time.<sup>5</sup> More precise evaluation of train-

<sup>5</sup>This experiment was conducted on a commercial cloud platform that limits a session to a maximum of three days. To finish training Barlow Twins with 8 GPUs for 1000 epochs, three sessions were required, whereas

Table 4. Linear evaluation accuracy (%) and the total training time on ImageNet with ResNet-50 backbone ( $d = 8192$ ). The per GPU batch size is 128. Proposed refers to the proposed method using the Barlow Twins-style regularizer without grouping. Barlow Twins\* refers to the results quoted from [29, Figure 2].

#GPUs (Batch size)	Model	Top-1	Top-5	Time
8 (1024)	Barlow Twins*	<b>73.2</b>	91.0	—
	Barlow Twins	72.4	90.7	6d 14h 0m
	Proposed	73.0	<b>91.3</b>	5d 14h 58m
4 (512)	Barlow Twins*	—	—	—
	Barlow Twins	72.1	90.2	12d 19h 30m
	Proposed	<b>72.8</b>	<b>91.2</b>	10d 21h 6m

Table 5. The effect of feature permutation: accuracy (%) and training time per 10 epochs (second) on ImageNet-100.

(a) Barlow Twins-style cross-correlation regularization				
Grouping	Permutation	Top-1	Top-5	Time
no	no	59.64	85.20	<b>1646.2</b>
	yes	<b>79.94</b>	<b>94.76</b>	1668.7
$b = 128$	no	73.58	93.36	<b>1697.5</b>
	yes	<b>81.02</b>	<b>95.24</b>	1709.6

(b) VICReg-style covariance regularization				
Grouping	Permutation	Top-1	Top-5	Time
no	no	57.42	84.26	<b>1692.2</b>
	yes	<b>79.20</b>	<b>94.96</b>	1718.0
$b = 128$	no	66.26	89.68	<b>1802.1</b>
	yes	<b>80.04</b>	<b>94.98</b>	1813.3

ing speed is provided in the next experiment, and additional results are provided in Supplementary Material.

### Dimensionality of embeddings and computational cost.

Fig. 2 presents the elapsed time and the peak GPU memory allocation over 10 epochs on ImageNet-100 on a single GPU, with varying dimensionality of the projected embeddings:  $d \in \{2048, 4096, 8192, 16384\}$ . The improvement over Barlow Twins and VICReg becomes noticeable as  $d$  is increased; at  $d = 8192$ , the proposed models (without grouping) are 2.8 times as fast as VICReg, and 2.2 times as fast as Barlow Twins; while at  $d = 16384$ , they are 5.7 times as fast as VICReg, and 4.0 times as fast as Barlow Twins. For both  $d = 8192$  and 16384, memory consumption is reduced by more than half. See Supplementary Material for the results with the ResNet-50 backbone on multiple GPUs and the detail of forward/backward/loss computation time.

the proposed model required two sessions. The time reported in Tab. 4 is the total run time of these sessions, which includes time (3 to 5 seconds) for reinitialization at the beginning of each session.

Table 6. The regularizers of Barlow Twins and VICReg applied to the embeddings produced by the proposed models. Perm.: permutation; Diff: difference to the baselines.

(a) Barlow Twins-style cross-correlation regularization				
Model	Grouping	Perm.	Normalized Barlow Twins loss (Eq. (16))	Diff
Barlow Twins	—	—	0.005	0
Proposed	no	no	0.564	0.559
		yes	<b>0.010</b>	<b>0.005</b>
	$b = 128$	no	0.049	0.044
		yes	<b>0.009</b>	<b>0.004</b>

(b) VICReg-style covariance regularization				
Model	Grouping	Perm.	Normalized VICReg loss (Eq. (17))	Diff
VICReg	—	—	0.002	0
Proposed	no	no	1.999	1.997
		yes	<b>0.011</b>	<b>0.009</b>
	$b = 128$	no	0.379	0.377
		yes	<b>0.007</b>	<b>0.005</b>

**Effectiveness of feature permutation.** Tab. 5 presents the effect of feature permutation on ImageNet-100 at  $d = 2048$ . Regardless of whether grouping is used, the accuracy decreases significantly without permutation, suggesting that feature permutation is essential for the effectiveness of our proposed regularizer. As illustrated in the column “Time” in Tab. 5, the cost of permutation is negligible, even though it was performed as frequently as every batch iteration.

We also quantitatively evaluate the degree of decorrelation obtained by feature permutation. After the proposed models are trained, we compute the normalized values of Barlow Twins’ and VICReg’s regularizers applied to the embeddings  $A, B$  output by the proposed models, given as follows:

$$\frac{R_{\text{off}}(\mathbf{C}(A, B))}{d(d-1)} \quad (16)$$

$$\frac{R_{\text{off}}(\mathbf{K}(A)) + R_{\text{off}}(\mathbf{K}(B))}{2d(d-1)} \quad (17)$$

The normalization factor  $d(d-1)$  is to make the resulting values the means over  $d(d-1)$  off-diagonal elements of  $\mathbf{K}(A)$ ,  $\mathbf{K}(B)$ , and  $\mathbf{C}(A, B)$ . Since VICReg has regularization terms for  $\mathbf{K}(A)$  and  $\mathbf{K}(B)$ , its loss is further divided by 2. Tab. 6 presents the results on ImageNet-100 ( $d=2048$ ). We observe that feature permutation promotes decorrelation from the point of view of the baseline loss functions.

**Impact of block size in feature grouping.** To evaluate the effect of feature grouping on ImageNet-100, we fix the embedding dimension at  $d = 2048$  and change the block

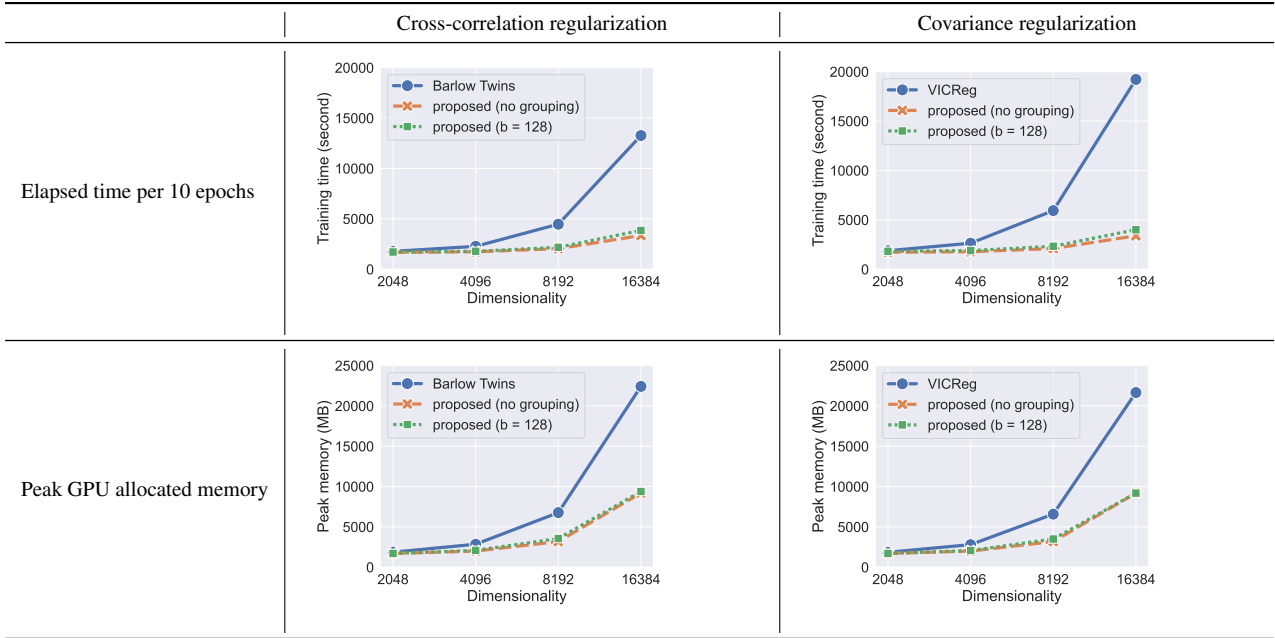


Figure 2. Training time and memory usage on ImageNet-100 with ResNet-18 on a single GPU.




Figure 3. The impact of the block size on ImageNet-100. The  $x$ -axis indicates the block size  $b$ .

size  $b \in \{2, 4, 8, \dots, 2048\}$ . The block size  $b = d = 2048$  corresponds to no feature grouping. Figure 3 presents the result, which indicates that unless  $b$  is extremely small (i.e., 8 or less), there is no significant increase in training time or GPU memory usage. Setting  $b$  to a moderate size, such as  $b = 128$ , improves performance.

## 6. Conclusion

We have proposed a new decorrelating regularizer for non-contrastive SSL. Given  $d$ -dimensional embeddings of  $n$  samples, the regularizer can be computed in  $O(nd \log d)$  time, which is an improvement over the  $O(nd^2)$  time required by the existing models (Barlow Twins and VICReg). The reduced memory consumption allows for a larger batch size, which in general enables more effective models to be learned. We also proposed a feature permutation technique to alleviate the weakness of our regularizer.

With our regularizer, the speed-up is most notable with networks with a lightweight backbone, wherein computation at the loss node takes a large part of training time. This makes our method an attractive option to use with knowledge-distilled backbones [12, 17].

The proposed method is not limited to SSL but potentially useful to a wider range of applications that involve decorrelation. We plan to pursue this avenue in future work.

## Acknowledgments

We thank anonymous reviewers for their helpful comments. This work was partially supported by the New Energy and Industrial Technology Development Organization (NEDO) and JSPS Kakenhi Grant 19H04173.



## References

- [1] Adrien Bardes, Jean Ponce, and Yann LeCun. VI-CReg: Variance-invariance-covariance regularization for self-supervised learning. In *ICLR*, 2022. 1, 2, 6
- [2] A. Borsellino and T. Poggio. Convolution and correlation algebras. *Kybernetik*, 13(2):113–122, 1973. 2
- [3] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, pages 1597–1607, 2020. 1, 2, 6
- [4] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E. Hinton. Big self-supervised models are strong semi-supervised learners. In *NeurIPS*, pages 22243–22255, 2020. 1, 2
- [5] Xinlei Chen and Kaiming He. Exploring simple Siamese representation learning. In *CVPR*, pages 15750–15758, 2021. 1, 2
- [6] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *ICCV*, pages 9640–9649, 2021. 2, 6
- [7] James W. Cooley and John W. Tukey. An algorithm for the machine calculation of complex Fourier series. *Mathematics of Computation*, 19(90):297–301, 1965. 1
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. 6
- [9] Debidatta Dwibedi, Yusuf Aytar, Jonathan Tompson, Pierre Sermanet, and Andrew Zisserman. With a little help from my friends: Nearest-neighbor contrastive learning of visual representations. In *ICCV*, pages 9588–9597, 2021. 2, 6
- [10] Aleksandr Ermolov, Aliaksandr Siarohin, Enver Sangineto, and Nicu Sebe. Whitening for self-supervised representation learning. In *ICML*, pages 3015–3024, 2021. 1, 2, 6
- [11] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The PASCAL visual object classes (VOC) challenge. *International Journal of Computer Vision*, 88(2):303–338, 2010. 6
- [12] Yuting Gao, Jia Xin Zhuang, Shaohui Lin, Hao Cheng, Xing Sun, Ke Li, and Chunhua Shen. DisCo: Remediating self-supervised learning on lightweight models with distilled contrastive learning. In *ECCV*, pages 237–253, 2022. 8
- [13] Jean-bastien Grill, Florian Strub, Florent Althé, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap your own latent: A new approach to self-supervised learning. In *NeurIPS*, pages 21271–21284, 2020. 1, 2, 6
- [14] Katsuhiko Hayashi and Masashi Shimbo. On the equivalence of holographic and complex embeddings for link prediction. In *ACL '17*, pages 554–559, 2017. 2
- [15] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9729–9738, 2020. 1, 2, 6
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 6
- [17] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 8
- [18] Tianyu Hua, Wenxiao Wang, Zihui Xue, Yue Wang, Sucheng Ren, and Hang Zhao. On feature decorrelation in self-supervised learning. In *ICCV*, pages 9598–9608, 2021. 2
- [19] Maximilian Nickel, Lorenzo Rosasco, and Tomaso Poggio. Holographic embeddings of knowledge graphs. In *AAAI*, pages 1955–1961, 2016. 2
- [20] Tony Plate. *Holographic Reduced Representation: Distributed Representation for Cognitive Structures*. CSLI Lecture Notes No. 150. CSLI Publications, 2003. 2, 4
- [21] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *NIPS*, pages 91–99, 2015. 6
- [22] P. H. Schönemann. Some algebraic relations between involutions, convolutions, and correlations, with applications to holographic memories. *Biological Cybernetics*, 56:367–374, 1987. 2, 4
- [23] Julius O. Smith, III. *Mathematics of the Discrete Fourier Transform (DFT) with Audio Applications*. W3K Publishing, 2nd edition, 2008. 4
- [24] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *ECCV*, pages 776–794, 2020. 2, 6
- [25] Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. Complex embeddings for simple link prediction. In *ICML*, pages 2071–2080, 2016. 2
- [26] Victor G. Turrissi da Costa, Enrico Fini, Moin Nabi, Nicu Sebe, and Elisa Ricci. solo-learn: A library of self-supervised methods for visual representation learning. *Journal of Machine Learning Research*, 23(56):1–6, 2022. 6
- [27] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv.cs preprint*, 1807.03748, 2018. 1, 2
- [28] Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *ICML*, pages 9929–9939, 2020. 1, 2
- [29] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphanie Deny. Barlow Twins: Self-supervised learning via redundancy reduction. In *ICML*, pages 12310–12320, 2021. 1, 2, 6, 7
- [30] Shaofeng Zhang, Feng Zhu, Junchi Yan, Rui Zhao, and Xi-aokang Yang. Zero-CL: Instance and feature decorrelation for negative-free symmetric contrastive learning. In *ICLR*, 2022. 1, 2, 6