

# Contrasting quadratic assignments for set-based representation learning

Artem Moskalev<sup>1</sup>, Ivan Sosnovik<sup>1</sup>, Volker Fischer<sup>2</sup>, and Arnold Smeulders<sup>1</sup>

<sup>1</sup> UvA-Bosch Delta Lab, University of Amsterdam

{a.moskalev, i.sosnovik, a.w.m.smeulders}@uva.nl

<sup>2</sup> Bosch Center for AI (BCAI)

volker.fischer@de.bosch.com

**Abstract.** The standard approach to contrastive learning is to maximize the agreement between different views of the data. The views are ordered in pairs, such that they are either positive, encoding different views of the same object, or negative, corresponding to views of different objects. The supervisory signal comes from maximizing the total similarity over positive pairs, while the negative pairs are needed to avoid collapse. In this work, we note that the approach of considering individual pairs cannot account for both intra-set and inter-set similarities when the sets are formed from the views of the data. It thus limits the information content of the supervisory signal available to train representations. We propose to go beyond contrasting individual pairs of objects by focusing on contrasting objects as sets. For this, we use combinatorial quadratic assignment theory designed to evaluate set and graph similarities and derive set-contrastive objective as a regularizer for contrastive learning methods. We conduct experiments and demonstrate that our method improves learned representations for the tasks of metric learning and self-supervised classification.

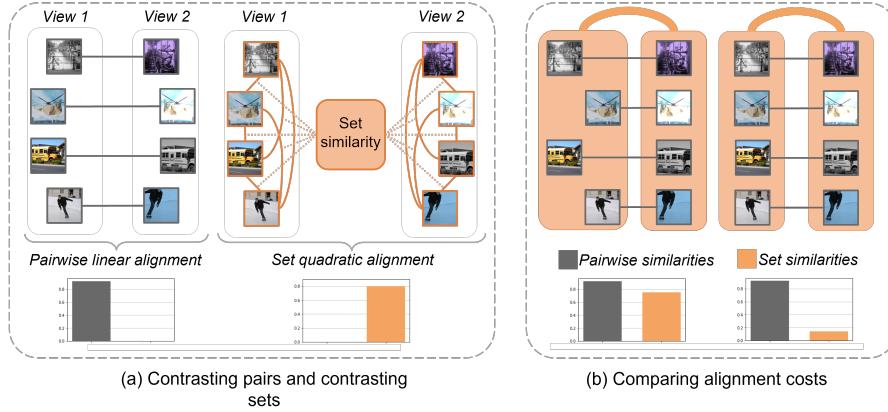
**Keywords:** contrastive learning, representation learning, self-supervised, optimal assignment

## 1 Introduction

The common approach to contrastive learning is to maximize the agreement between individual views of the data [33, 28]. The views are ordered in pairs, such that they either positive, encoding different views of the same object, or negative, corresponding to views of different objects. The pairs are compared against one another by a contrastive loss-objective [39, 44, 12]. Contrastive learning was successfully applied among others in metric learning [23], self-supervised classification [12, 21], and pre-training for dense prediction tasks [45]. However, when two views of an object are drastically different, the view of object A will resemble the same view of object B much more than it resembles the other view of object A. By comparing individual pairs, the common patterns in the differences

---

Source code: [github.com/amoskalev/contrasting-quadratic](https://github.com/amoskalev/contrasting-quadratic)



**Fig. 1.** (a) Contrasting views by similarity over pairs (pairwise linear alignment) versus similarity over sets with set-alignment (set-wise quadratic alignment). (b) Comparing total pairwise similarities versus set similarities for different configurations of representation graphs. In both configurations, the total similarity over pairs remains the same, being unable to discriminate between the internal structures of different views as quadratic alignment can.

between two views will not be exploited. In this paper, we propose to go beyond contrasting individual pairs of objects and focus on contrasting sets of objects.

In contrastive learning, the best alignment of objects follows from maximizing the total similarity over positive pairs, while the negative pairs are needed to encourage diversity. The diversity in representations is there to avoid collapse [42]. We note that one number, the total similarity from contrasting individual pairs, cannot both account for the inter-set similarities of the objects from the same view and the intra-set similarities with objects from another view. Therefore, considering the total similarity over pairs essentially limits the information content of the supervisory signal available to the model. We aim to exploit the information from contrasting sets of objects rather than contrasting objects pairwise only, further illustrated in Figure 1a. From the perspective of set assignment theory [8], two sets may expose the same linear pairwise alignment costs, yet have different internal structures and hence have different set alignment costs, see Figure 1b. Therefore, contrastive learning from sets provides a richer supervisory signal.

To contrast objects as sets, we turn to combinatorial quadratic assignment theory [8] designed to evaluate set and graph similarities. Quadratic assignment has advanced from linear assignment [9], which relies on pairwise similarities between objects. We pose contrastive learning by structural risk minimization [43] in assignment problems. In this, pairwise contrastive learning methods emerge from the linear assignment case. We aim to extend the contrastive objective to the set level by generalizing the underlying assignment problem from linear to quadratic, as illustrated in Figure 1a. Since directly computing set similarities from quadratic alignment is computationally expensive, we derive

an efficient approximation. It can be implemented as a regularizer for existing contrastive learning methods. We provide a theory for the proposed method from the perspective of assignment problems and dependence maximization. And, we experimentally demonstrate the advantages of contrasting objects by quadratic alignment.

We make the following contributions:

- Using combinatorial assignment theory, we propose to learn representations by contrasting objects as sets rather than as pairs of individual instances.
- We propose a computationally efficient implementation, where the set contrastive part is implemented as a regularizer for existing contrastive learning methods.
- We demonstrate that set-contrastive regularization improves recent contrastive learning methods for the tasks of metric learning and self-supervised classification.

As a byproduct of viewing representation learning through the lens of combinatorial assignment theory, we additionally propose SparseCLR contrastive loss, a modification of the popular InfoNCE [39] objective. Different from InfoNCE, SparseCLR enjoys sparse support over negative pairs, which permits better use of hard negative examples. Our experiments demonstrate that such an approach improves the quality of learned representations for self-supervised classification.

## 2 Related work

### 2.1 Contrastive learning

In contrastive learning, a model is trained to align representations of the data from different views (modalities), which dates back to the work of Becker *et al.* [4]. Contrastive learning is made up of a joint embedding architecture like Siamese networks [7] and a contrastive objective function. The network maps different views of the data into the embedding space, where the alignment between embeddings is measured by contrasting positive and negative pairs. The pairs either from a complete set of observations [41, 12, 40, 13] or from partially-complete views [26, 47, 32, 46]. And as a contrastive objective, contrastive loss functions are used [39, 12, 21, 44]. Van den Oord *et al.* [39] derived the InfoNCE contrastive loss as a mutual information lower bound between different views of the signal. In computer vision, Chen *et al.* apply the InfoNCE to contrast the images with their distorted version. In PIRL [37], the authors propose to maintain a memory bank of image representations to improve the generalization of the model by sampling negative samples better. Along the same lines, MoCO [21, 14] proposes a running queue of negative samples and a momentum-encoder to increase the contrasting effect of negative pairs. All of these contrastive methods are based on measuring the alignment between individual pairs of objects. This approach does not account for patterns in the views of objects beyond contrasting them pairwise. We aim to extend contrastive learning to include set similarities.

## 2.2 Information maximization methods

In contrastive learning information maximization methods aim to improve contrastive representations by maximizing the information content of the embeddings [2]. In W-MSE [16], batch representations are passed through a whitening, Karhunen-Loëve transformation before computing the contrastive loss. Such transformations help to avoid collapsing representations when the contrastive objective can be minimized without learning the discriminative representations. In [48], the authors follow the redundancy-reduction principle to design a loss function to bring the cross-correlation matrix of representations from different views close to identity. In the recent work of Zbontar *et al.*, the authors extend the above formulation of Barlow Twins with an explicit variance term, which helps to stabilize the training. In this paper, we also seek to maximize the information content of the embeddings but we aim to do so by computing the rich representation of set similarities between data views. In contrast to existing methods, our approach does not require any additional transformations like whitening and can be easily incorporated into other contrastive learning methods.

## 2.3 Distillation methods

Another branch of self-supervised methods is not based on contrastive learning but rather based on knowledge distillation [24]. Instead of contrasting positive and negative pairs of samples, these methods try to predict one positive view from another. We discuss them here for completeness. BYOL [19] uses a student network that learns to predict the output of a teacher network. SimSiam [15] simplifies the BYOL by demonstrating that the stop-gradient operation suffices to learn a generalized representation from distillation. In SWaV [11], the authors combine online clustering and distillation by predicting swapped cluster assignments. These self-supervised methods have demonstrated promising results in the task of learning representations. We prefer to follow a different path of contrastive learning, where we do not exclude the possibility that set-based methods may also be applicable to distillation.

## 2.4 Assignment theory

Generally, optimal assignment problems can be categorized into linear and higher-order assignments. Linear assignment problems can be viewed as a transportation problem over a bipartite graph, where the transportation distances are defined on individual pairs of nodes. Linear assignment problems have many real-world applications such as scheduling or vehicle routing, we refer to [9]. Higher-order or in particular Quadratic assignment problems [17, 8] extend the domain of the transportation distances from individual pairs to sets of objects. Where the linear assignment problem seeks to optimally match individual instances, its quadratic counterpart matches the inter-connected graphs of instances, bringing additional structure to the task. In this work, we aim to exploit linear and quadratic assignment theory to rethink contrastive learning.

## 2.5 Structured predictions

In a structured prediction problem, the goal is to learn a structured object such as a sequence, a tree, or a permutation matrix [1]. Training a structure is non-trivial as one has to compute the loss on the manifold defined by the structured output space. In the seminal work of Tschantaridis *et al.* [43], the authors derive a structured loss for support vector machines and apply it to a sequence labeling task. Later, structured prediction was applied to learn parameters of constraint optimization problems [20, 6]. In this work, we utilize structured prediction principles to implement set similarities into contrastive losses.

## 3 Background

We start by formally introducing contrastive learning and linear assignment problems. Then, we argue about the connection between these two problems and demonstrate how one leads to another. This connection is essential for the derivation of our contrastive method.

### 3.1 Contrastive representation learning

Consider a dataset  $\mathcal{D} = \{d_i\}_{i=1}^N$  and an encoder function  $f_\theta : \mathcal{D} \rightarrow \mathbb{R}^{N \times E}$ , which outputs an embedding vector of dimension  $E$  for each of the objects in  $\mathcal{D}$ . The task is to learn a  $f_\theta$  such that embeddings of objects belonging to the same category are pulled together, while the embeddings of objects from different categories are pushed apart.

As category labels are not available during training, the common strategy is to maximize the agreement between different views of the instances, i.e. to maximize the mutual information between different data modalities. The mutual information between two latent variables  $X, Y$ , can be expressed as:

$$MI(X, Y) = H(X) - H(X|Y) \quad (1)$$

where  $X$  and  $Y$  correspond to representations of two different views of a dataset. Minimizing the conditional entropy  $H(X|Y)$  aims at reducing uncertainty in the representations from one view given the representations of another, while  $H(X)$  enforces the diversity in representations and prevents trivial solutions. In practice, sample-based estimators of the mutual information such as InfoNCE [39] or NT-Xent [12] are used to maximize the objective in (1).

### 3.2 Linear assignment problem

Given two input sets,  $\mathcal{A} = \{a_i\}_{i=1}^N$  and  $\mathcal{B} = \{b_j\}_{j=1}^N$ , we define the inter-set similarity matrix  $\mathbf{S} \in \mathbb{R}^{N \times N}$  between each element in set  $\mathcal{A}$  and each element in  $\mathcal{B}$ . This similarity matrix encodes *pairwise* distances between the elements of the sets, *i.e.*  $[\mathbf{S}]_{i,j} = \phi(a_i, b_j)$ , where  $\phi$  is a distance metric. The goal of the linear

assignment problem is to find a one-to-one assignment  $\hat{y}(\mathbf{S})$ , such that the sum of distances between assigned elements is optimal. Formally:

$$\hat{y}(\mathbf{S}) = \underset{\mathbf{Y} \in \Pi}{\operatorname{argmin}} \operatorname{tr}(\mathbf{S}\mathbf{Y}^T) \quad (2)$$

where  $\Pi$  corresponds to a set of all  $N \times N$  permutation matrices.

The linear assignment problem in (2) is also known as the bipartite graph matching problem. It can be efficiently solved by linear programming algorithms [5].

### 3.3 Learning to produce correct assignments

Consider two sets,  $\mathcal{D}_{\mathcal{Z}_1}$  and  $\mathcal{D}_{\mathcal{Z}_2}$ , which encode two different views of the dataset  $\mathcal{D}$  under two different views  $\mathcal{Z}_1$  and  $\mathcal{Z}_2$ . By design, the two sets consist of the same instances, but the modalities and the order of the objects may differ. With the encoder, which maximizes the mutual information, objects in both sets can be uniquely associated with one another by comparing their representations as the uncertainty term in (1) should be minimized. In other words,  $\hat{y}(\mathbf{S}) = \mathbf{Y}_{gt}$ , where  $\mathbf{Y}_{gt} \in \Pi$  is a ground truth assignment between elements of  $\mathcal{D}_{\mathcal{Z}_1}$  and  $\mathcal{D}_{\mathcal{Z}_2}$ .

Thus, a natural objective to supervise an encoder function is to train it to produce the correct assignment between the different views of the data. As the assignment emerges as a result of the optimization problem, we employ a structured prediction framework [43] to define a structural loss from the linear assignment problem as follows:

$$L(\mathbf{S}, \mathbf{Y}_{gt}) = \operatorname{tr}(\mathbf{S}\mathbf{Y}_{gt}^T) - \min_{\mathbf{Y} \in \Pi} \operatorname{tr}(\mathbf{S}\mathbf{Y}^T) \quad (3)$$

where  $\mathbf{S} = \phi(f_\theta(\mathcal{D}_{\mathcal{Z}_1}), f_\theta(\mathcal{D}_{\mathcal{Z}_2}))$ . Note that  $L \geq 0$  and  $L = 0$  only when the similarities produced by an encoder lead to the correct optimal assignment. Intuitively, the structured linear<sup>3</sup> assignment loss in (3) encodes the discrepancy between the true cost of the assignment and the cost induced by an encoder.

By minimizing the objective in (3), we train the encoder  $f_\theta$  to correctly assign objects from one view to another. In practice, it is desirable that such assignment is robust under small perturbations in the input data. The straightforward way to do so is to enforce a separation margin  $m$  in the input similarities as  $\mathbf{S}_m = \mathbf{S} + m\mathbf{Y}_{gt}$ . Then, the structured linear assignment loss with separation margin reduces to a known margin Triplet loss [23].

**Proposition 1** *The structured linear assignment loss  $L(\mathbf{S}_m, \mathbf{Y}_{gt})$  with separation margin  $m \geq 0$  is equivalent to the margin triplet loss.*

---

<sup>3</sup> We emphasize that the term *linear* here is only used with regard to an underlying assignment problem.

*Mining strategies.* By default, the loss in (3) enforces a one-to-one negative pair mining strategy due to the structural domain constraint  $\mathbf{Y} \in \Pi$ . By relaxing this domain constraint to row-stochastic binary matrices, we arrive at the known batch-hard mining [23]. This is essential to have a computationally tractable implementation of structured assignment losses.

*Smoothness.* An immediate issue when directly optimizing the structured linear assignment loss is the non-smoothness of the minimum function in (3). It is known that optimizing smoothed functions can be more efficient than directly optimizing their non-smooth counterparts [3]. The common way to smooth a minimum is by *log-sum-exp* approximation. Thus, we can obtain a smoothed version of structured linear assignment loss:

$$L_\tau(\mathbf{S}, \mathbf{Y}_{gt}) = \text{tr}(\mathbf{S}\mathbf{Y}_{gt}^T) + \tau \log \sum_{\mathbf{Y} \in \Pi} \exp\left(-\frac{1}{\tau} \text{tr}(\mathbf{S}\mathbf{Y}^T)\right) \quad (4)$$

where  $\tau$  is a temperature parameter controlling the degree of smoothness. Practically, the formulation in (4) requires summing  $N!$  terms, which makes it computationally intractable under the default structural constraints  $\mathbf{Y} \in \Pi$ . Fortunately, similar to the non-smooth case, we can utilize batch-hard mining, which leads to  $O(N^2)$  computational complexity. The smoothed structured linear assignment loss with batch-hard mining reduces to known normalized-temperature cross entropy [12] also known as the InfoNCE [39] loss.

**Proposition 2** *The smoothed structured linear assignment loss  $L_\tau(\mathbf{S}, \mathbf{Y}_{gt})$  with batch-hard mining is equivalent to the normalized-temperature cross entropy loss.*

*Connection to mutual information* It is known that the InfoNCE objective is a lower bound on the mutual information between the representations of data modalities [39]. Thus, Propositions 1 and 2 reveal a connection between mutual information maximization and minimization of structured losses for assignment problems. In fact, the assignment cost  $\text{tr}(\mathbf{S}\mathbf{Y}_{gt}^T)$  in (2) is related to the conditional entropy  $H(X|Y)$ , while  $\min_{\mathbf{Y} \in \Pi} \text{tr}(\mathbf{S}\mathbf{Y}^T)$  aims to maximize the diversity in representations. This connection allows for consideration of contrastive representation learning methods based on InfoNCE as a special case of the structured linear assignment loss.

## 4 Extending contrastive losses

We next demonstrate how to exploit the connection between contrastive learning and assignment problems to extend contrastive losses to the set level. As a byproduct of this connection, we also derive the SparseCLR contrastive objective.

#### 4.1 Contrastive learning with quadratic assignments on sets

**Quadratic assignment problem** As in the linear case, we are given two input sets  $\mathcal{A}, \mathcal{B}$  and the inter-set similarity matrix  $\mathbf{S}$ . For the Quadratic Assignment Problem (QAP), we define intra-set similarity matrices  $\mathbf{S}_{\mathcal{A}}$  and  $\mathbf{S}_{\mathcal{B}}$  measuring similarities within the sets  $\mathcal{A}$  and  $\mathcal{B}$  respectively, *i.e.*  $[\mathbf{S}_{\mathcal{A}}]_{ij} = \phi(a_i, a_j)$ . A goal of the quadratic assignment problem is to find a *one-to-one* assignment  $\hat{y}_Q(\mathbf{S}, \mathbf{S}_{\mathcal{A}}, \mathbf{S}_{\mathcal{B}}) \in \Pi$  that maximizes the set similarity between  $\mathcal{A}, \mathcal{B}$ , where the set similarity is defined as follows:

$$\mathcal{Q}(\mathbf{S}, \mathbf{S}_{\mathcal{A}}, \mathbf{S}_{\mathcal{B}}) = \min_{\mathbf{Y} \in \Pi} \{ \text{tr}(\mathbf{S}\mathbf{Y}^T) + \text{tr}(\mathbf{S}_{\mathcal{A}}\mathbf{Y}\mathbf{S}_{\mathcal{B}}^T\mathbf{Y}^T) \} \quad (5)$$

Compared to the linear assignment problem, the quadratic term  $\mathcal{Q}(\mathbf{S}, \mathbf{S}_{\mathcal{A}}, \mathbf{S}_{\mathcal{B}})$  in (5) additionally measures the discrepancy in internal structures between sets  $\mathbf{S}_{\mathcal{A}}$  and  $\mathbf{S}_{\mathcal{B}}$ .

**Learning with quadratic assignments** Following the similar steps as for the linear assignment problem in Section 3.3, we next define the structured quadratic assignment loss by extending the linear assignment problem in (3) with the quadratic term:

$$L_{QAP} = \text{tr}(\mathbf{S}\mathbf{Y}_{gt}) - \mathcal{Q}(\mathbf{S}, \mathbf{S}_{\mathcal{A}}, \mathbf{S}_{\mathcal{B}}) \quad (6)$$

Minimization of the structured quadratic assignment loss in (6) encourages the encoder to learn representations resulting in the same solutions of the linear and quadratic assignment problems, which is only possible when the inter-set and intra-set similarities are sufficiently close [8]. Note that we do not use a ground truth quadratic assignment in  $L_{QAP}$ , but a ground truth linear assignment objective is used for the supervision. This is due to a quadratic nature of  $\mathcal{Q}$ , where minimizing the discrepancy between ground truth and predicted assignments is a subtle optimization objective, *e.g.* for an equidistant set of points the costs of all quadratic assignments are identical.

To compute  $L_{QAP}$ , we first need to evaluate the quadratic term  $\mathcal{Q}(\mathbf{S}, \mathbf{S}_{\mathcal{A}}, \mathbf{S}_{\mathcal{B}})$ , which requires solving the quadratic assignment. This problem is known to be notoriously hard to solve exactly even when an input dimensionality is moderate [8]. To alleviate this, we use a computationally tractable upper-bound:

$$\begin{aligned} L_{QAP} &\leq \text{tr}(\mathbf{S}\mathbf{Y}_{gt}) - \min_{\mathbf{Y} \in \Pi} \text{tr}(\mathbf{S}\mathbf{Y}^T) - \min_{\mathbf{Y} \in \Pi} \text{tr}(\mathbf{S}_{\mathcal{A}}\mathbf{Y}\mathbf{S}_{\mathcal{B}}^T\mathbf{Y}^T) \\ &\leq L(\mathbf{S}, \mathbf{Y}_{gt}) - \langle \lambda_{\mathcal{A}}, \lambda_{\mathcal{B}} \rangle_- \end{aligned} \quad (7)$$

where  $\langle \lambda_{\mathcal{A}}, \lambda_{\mathcal{B}} \rangle_-$  corresponds to a minimum dot product between eigenvalues of matrices  $S_{\mathcal{A}}$  and  $S_{\mathcal{B}}$ .

The first inequality is derived from Jensen's inequality for (6) and the second inequality holds due to the fact that  $\langle \lambda_F, \lambda_D \rangle_- \leq \text{tr}(\mathbf{F}\mathbf{X}\mathbf{D}^T\mathbf{X}^T) \leq \langle \lambda_F, \lambda_D \rangle_+$

---

**Algorithm 1** Pseudocode for set-based InfoNCE with Euclidean distance similarity and Quadratic Assignment Regularization (QARe).

---

```

# f: encoder network
# alpha: weighting for the pairwise contrastive part (linear assignment)
# beta: weighting for the set contrastive part
# N: batch size
# E: dimensionality of the embeddings

for x in loader: # load a batch with N samples
    # two randomly augmented views of x
    y_a, y_b = augment(x)

    # compute embeddings
    z_a = f(y_a) # NxE
    z_b = f(y_b) # NxE

    # compute inter-set and intra-set similarities
    S_AB = similarity(z_a, z_b) # NxN
    S_A = similarity(z_a, z_a) # NxN
    S_B = similarity(z_b, z_b) # NxN

    # compute pairwise contrastive InfoNCE loss
    pairwise_term = infonce(S_AB)

    # compute eigenvalues
    eigs_a = eigenvalues(S_A) #N
    eigs_b = eigenvalues(S_B) #N

    # compute QARe from minimum dot product of eigenvalues
    eigs_a_sorted = sort(eigs_a, descending = True) #N
    eigs_b_sorted = sort(eigs_b, descending = False) #N
    qare = -1*(eigs_a_sorted.T@eigs_b_sorted)

    # combine pairwise contrastive loss with QARe
    loss = alpha*pairwise_term + beta*qare/(N^2)

    # optimization step
    loss.backward()
    optimizer.step()

```

---

for symmetric matrices  $\mathbf{F}, \mathbf{D}$  and  $\mathbf{X} \in \Pi$  as shown in *Theorem 5* by Burkard [8]. The above derivations are for the case when the similarity metric is a valid distance function, i.e. minimizing a distances leads to maximizing a similarity. The derivation for the reverse case can be done analogously by replacing  $\min$  with  $\max$  in the optimal assignment problem formulation (we provide more details in supplementary material).

Optimizing the upper-bound in (7) is computationally tractable compared to optimizing the exact version of  $L_{QAP}$ . Another advantage is that the upper-bound in (7) breaks down towards minimizing the sum of the structured linear assignment loss  $L(\mathbf{S}, \mathbf{Y}_{gt})$  and the regularizing term, which accounts for the set similarity. This allows to easily modify existing contrastive learning approaches like those in [12, 23, 36, 39] that are based on pairwise similarities and thus stem from the linear assignment problem by simply plugging in the regularizing term. We provide a simple pseudocode example demonstrating InfoNCE with the quadratic assignment regularization (Algorithm 1).

**Computational complexity** Computing the upper-bound in (7) has a computational complexity of  $O(N^3)$  as one needs to compute eigenvalues of the intra-set

similarity matrices. This is opposed to (6) that requires directly computing quadratic assignments, for which it is known there exist no polynomial-time algorithms [8]. The computational complexity can be pushed further down to  $O(k^2N)$  by evaluating the only top-k eigenvalues instead of computing all eigenvalues. In supplementary materials, we also provide an empirical analysis of how the proposed approach influences a training time of a baseline contrastive method.

## 4.2 SparseCLR

In Section 3 we noted that the smoothness of a loss function is a desirable property for optimization and that the *log-sum-exp* smoothing of the structured linear assignment loss yields the normalized temperature cross-entropy objective as a special case. Such an approach, however, is known to have limitations. Specifically, the *log-sum-exp* smoothing yields dense support over samples, being unable to completely rule out irrelevant examples [38]. That can be statistically and computationally overwhelming and can distract the model from fully utilizing information from hard negative examples.

We propose to use *sparsemax* instead of *log-sum-exp* approximation to alleviate the non-smoothness of the structured linear assignment objective, yet keep the sparse support of the loss function. Here *sparsemax* is defined as the minimum distance Euclidean projection to a k-dimensional simplex as in [38, 35].

Let  $\tilde{x} \in \mathbb{R}^{N!}$  be a vector that consists of the realizations of all possible assignment costs for the similarity matrix  $\mathbf{S}$ . With this we can define SparseCLR as follows:

$$L_{\text{sparse}}(\mathbf{S}, \mathbf{Y}_{gt}) = \text{tr}(\mathbf{S}\mathbf{Y}_{gt}^T) - \frac{1}{2} \sum_{j \in \Omega(\tilde{x})} (\tilde{x}_j^2 - \mathcal{T}^2(-\tilde{x})) \quad (8)$$

where  $\Omega(X) = \{j \in X : \text{sparsemax}(X)_j > 0\}$  is the support of *sparsemax* function, and  $\mathcal{T}$  denotes the thresholding operator [35]. In practice, to avoid summing over a factorial number of terms, as well as for other methods, we use batch-hard mining strategy resulting in  $O(N^2)$  computational complexity for SparseCLR.

## 5 Experiments

In this section, we evaluate the quality of representations trained with and without our Quadratic Assignment Regularization (QARe). We also evaluate the performance of SparseCLR and compare it with other contrastive losses. As we consider the representation learning from the perspective of the assignment theory, we firstly conduct the instance matching experiment, where the goal is to learn to predict the correct assignment between different views of the dataset. Then, we test the proposed method on the task of self-supervised classification and compare it with other contrastive learning approaches. Next, we present ablation studies to visualize the role of the weighting term, when combining QARe with the baseline contrastive learning method.

Method	Triplet	SparseCLR	InfoNCE	NTLogistic
<i>pairwise</i>	54.85±0.77	53.03±0.43	57.87±1.19	40.30±0.42
+QARe	<b>58.48±1.67</b>	<b>54.84±1.40</b>	<b>61.96±1.38</b>	<b>43.48±1.41</b>

**Table 1.** Matching accuracy for instances from different views of CUHK-03 dataset. The *pairwise* corresponds to the contrastive losses acting on the level of pairs. +QARe denotes methods with quadratic assignment regularization. Best results are in bold.

### 5.1 Matching instances from different views

In this experiment, the goal is to train representations of objects from different views, such that the learned representations provide a correct matching between identities in the dataset. This problem is closely related to a metric learning and can be solved by contrasting views of the data [23].

**Data** We adopt the CUHK-03 dataset [31] designed for the ReID [23] task. CUHK-03 consist of 1467 different identities recorded each from front and back views. To train and test the models, we randomly divide the dataset into 70/15/15 train/test/validation splits.

**Evaluation** We evaluate the quality of representations learned with contrastive losses by computing the matching accuracy between front and back views. In practice, we first obtain the embeddings of the views of the instances from the encoder and then compute their inter-view similarity matrix using the Euclidean distance. Given this, the matching accuracy is defined as the mean Hamming distance between the optimal assignment from the inter-view similarity matrix and the ground truth assignment. We report the average matching accuracy and the standard deviation over 3 runs with the same fixed random seeds. As a baseline, we chose Triplet with batch-hard mining [23], InfoNCE [39] and NTLogistic [36] contrastive losses, which we extend with the proposed QARe.

**Implementation details** As the encoder, we use ResNet-18 [22] with the top classification layer replaced by the linear projection head that outputs feature vectors of dimension 64. To obtain representations, we normalize the feature vectors to be L2 unit-norm.

We train the models for 50 epochs using Adam optimizer [29] with the cosine annealing without restarts [34]. Initial learning rate is 0.01 and a batch size is set to 128. During the training, we apply random color jittering and horizontal flip augmentations. To compute the test matching accuracy, we select the best model over the epochs based on the validation split. We provide detailed hyperparameters of the losses, QARe weighting, and augmentations in the supplementary material.

**Results** The results are reported in Table 1. As can be seen, adding QARe regularization to a baseline contrastive loss leads to a better matching accuracy,

Method	Conv-4			ResNet-32		
	CIFAR-10	CIFAR-100	tiny-ImageNet	CIFAR-10	CIFAR-100	tiny-ImageNet
Supervised (upper bound)	80.46±0.39	49.29±0.85	36.47±0.36	90.87±0.41	65.32±0.22	50.09±0.32
Random Weights (lower bound)	32.92±1.88	10.79±0.59	6.19±0.13	27.47±0.83	7.65±0.44	3.24±0.43
DeepCluster [10]	42.88±0.21	21.03±1.56	12.60±1.23	43.31±0.62	20.44±0.80	11.64±0.21
RotationNet [18]	56.73±1.71	27.45±0.80	18.40±0.95	62.00±0.79	29.02±0.18	14.73±0.48
Deep InfoMax [25]	44.60±0.27	22.74±0.21	14.19±0.13	47.13±0.45	24.07±0.05	17.51±0.15
SimCLR* [12]	59.64±0.93	30.75±0.23	21.15±0.16	76.45±0.10	40.18±0.13	23.44±0.27
SimCLR+QARe (Ours)	<b>61.54±1.17</b>	<b>31.18±0.41</b>	<b>21.33±0.18</b>	<b>76.85±0.39</b>	<b>41.20±0.34</b>	<b>25.05±0.18</b>

**Table 2.** Self-supervised training and linear probing for a self-supervised classification. Average classification accuracy (percentage) and standard deviation over 3 runs with common fixed random seeds are reported. \* we report the results from our reimplementa-tion of SimCLR.

which indicates an improved quality of representations. Notably, QARe delivers 3.6% improvement when combined with Triplet loss and 4.1% increase in accuracy with the InfoNCE objective. This demonstrates that the quadratic assignment regularization on sets helps the model to learn generalized representation space.

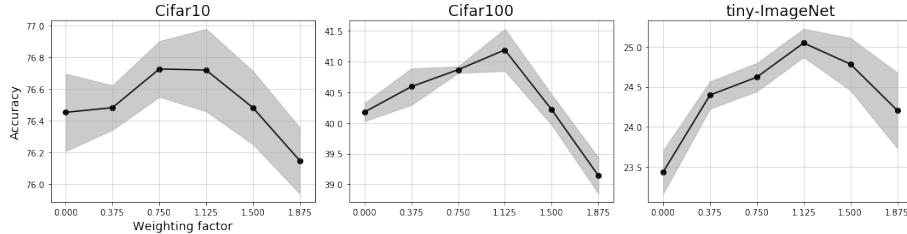
## 5.2 Self-supervised classification

Next, we evaluate the quadratic assignments on sets in the task of self-supervised classification. The goal in this experiment is to learn a representation space from an unlabeled data that provides a linear separation between classes in a dataset. This is a common problem to test contrastive learning methods.

**Evaluation** We follow the standard linear probing protocol [30]. The evaluation is performed by retrieving embeddings from a contrastively trained backbone and fitting a linear classifier on top of the embeddings. Since linear classifiers have a low discriminative power on their own, they heavily rely on input representations. Thus, a higher classification accuracy indicates a better quality of learned representations. As the baseline for pairwise contrastive learning methods, we select popular SimCLR and proposed SparseCLR. We extend the methods to the set level by adding the quadratic assignment regularization. We compare the set-level contrastive methods with other self-supervised [10, 18] and contrastive approaches [12, 11].

**Implementation details** We train shallow convolutional (Conv-4) and deep (ResNet-32) backbone encoders to obtain representations. We follow the standard approach and attach an MLP projection head that outputs feature vectors of dimension 64 to compute a contrastive loss. We use the cosine distance to get similarities between embeddings. We provide more details on the encoder architectures in the supplementary material.

We train the models for 200 epochs using Adam optimizer with a learning rate of 0.001 and a batch size is set to 128. For contrastive learning methods, we adopt



**Fig. 2.** Influence of the weighting of the QARe term in the self-supervised classification for ResNet-32 trained with SimCLR. The accuracy is recorded over 3 runs with common fixed random seeds.

Method	Conv-4			ResNet-32		
	CIFAR-10	CIFAR-100	tiny-ImageNet	CIFAR-10	CIFAR-100	tiny-ImageNet
SimCLR* [12]	59.64±0.93	30.75±0.23	21.15±0.16	<b>76.45±0.10</b>	40.18±0.13	23.44±0.27
SparseCLR (Ours)	59.86±0.56	32.04±0.26	<b>22.41±0.39</b>	70.37±0.12	41.05±0.20	25.87±0.59
SparseCLR+QARe (Ours)	<b>61.06±0.39</b>	<b>33.05±0.32</b>	<b>22.45±0.30</b>	71.39±0.30	<b>42.07±0.35</b>	<b>27.03±0.40</b>

**Table 3.** Self-supervised training on unlabeled data and linear evaluation on labeled data. Comparing SimCLR [12] with the proposed SparseCLR and SparseCLR with QARe. Average accuracy and standard deviation over three runs over 3 runs with common fixed random seeds.

a set of augmentations from [12] to form different views of the data. For linear probing, we freeze the encoder weights and train a logistic regression on top of the learned representation for 100 epochs with Adam optimizer, using a learning rate of 0.001 and a batch size of 128. We provide detailed hyperparameters of the contrastive losses and augmentations in the supplementary material.

We perform training and testing on standard CIFAR-10, CIFAR-100, and tiny-ImageNet datasets. For each dataset, we follow the same training procedure and fine-tune the optimal weighting of the quadratic assignment regularizer.

**QARe results** The results for the quadratic assignment regularization are reported in Table 2 for SimCLR and in Table 3 for proposed SparseCLR. For SimCLR, adding QARe delivers 1.9% accuracy improvement on CIFAR-10 for the shallow encoder network and up to 1.6 % improvement for ResNet-32 on tiny-ImageNet. We observed a consistent improvement in other dataset-architecture setups, except tiny-ImageNet with shallow Conv-4 encoder, where the performance gain from adding QARe is modest. For this case, we investigated the training behavior of the models and observed that both SimCLR and SimCLR+QARe for Conv-4 architecture very quickly saturate to a saddle point, where the quality of representations stops improving. Since this does not happen with ResNet-32 architecture, we attribute this phenomenon to the limited discriminative power of the shallow Conv-4 backbone, which can not be extended by regularizing the loss.

For SparseCLR, we observed the same overall pattern. As can be seen from Table 3, extending SparseCLR to set level with QARe delivers 1.6% accuracy improvement for ResNet-32 on tiny-ImageNet and also steadily improve the performance on other datasets. This indicates that QARe helps to provide a richer supervisory signal for the model to train representations.

**SparseCLR results** Next, we compare SimCLR against the proposed SparseCLR. As can be seen in Table 3, SparseCLR consistently improves over the baseline on CIFAR-100 and tiny-ImageNet datasets, where it delivers 2.4% improvement. Surprisingly, we noticed a significant drop in performance for ResNet-32 on CIFAR-10. For this case, we investigated the training behavior of the model and observed that in the case of CIFAR-10, the batch often includes many false negative examples, which results in uninformative negative pairs. Since SparseCLR assigns the probability mass only to a few hardest negative pairs, a lot of false-negative examples in the batch impede obtaining a clear supervisory signal. We assume the problem can be alleviated by using false-negative cancellation techniques [27].

### 5.3 Ablation study

Here we illustrate how the weighting of the QARe term influences the quality of representations learned with SimCLR under a linear probing evaluation. We search for the optimal weighting in the range from 0 to 1.875 with a step of 0.125. The results are depicted in Figure 2. In practice, we observed that QARe is not too sensitive to a weighting and the values in the range 0.75-1.25 provide consistent improvement.

## 6 Discussion

In this work, we present set contrastive learning method based on quadratic assignments. Different from other contrastive learning approaches, our method works on the level of set similarities as opposed to only pairwise similarities, which allows improving the information content of the supervisory signal. For derivation, we view contrastive learning through the lens of combinatorial assignment theory. We show how pairwise contrastive methods emerge from learning to produce correct optimal assignments and then extend them to a set level by generalizing an underlying assignment problem, implemented as a regularization for existing methods. As a byproduct of viewing representation learning through the lens of assignment theory, we additionally propose SparceCLR contrastive loss.

Our experiments in instance matching and self-supervised classification suggest that adding quadratic assignment regularization improves the quality of representations learned by backbone methods. We suppose, that our approach would be most useful in the problems where the joint analysis of objects and their groups appears naturally and labeling is not readily available.

## References

1. BakIr, G., Hofmann, T., Schölkopf, B., Smola, A.J., Taskar, B., Vishwanathan, S.: Predicting Structured Data. The MIT Press (2007)
2. Bardes, A., Ponce, J., LeCun, Y.: Vicreg: Variance-invariance-covariance regularization for self-supervised learning. In: ICLR (2022)
3. Beck, A., Teboulle, M.: Smoothing and first order methods: A unified framework. SIAM Journal on Optimization **22**(2), 557–580 (2012)
4. Becker, S., Hinton, G.E.: Self-organizing neural network that discovers surfaces in random-dot stereograms. Nature **355**, 161–163 (1992)
5. Bertsimas, D., Tsitsiklis, J.: Introduction to Linear Optimization. Athena Scientific (1998)
6. Blondel, M., Martins, A., Niculae, V.: Learning classifiers with fenchel-young losses: Generalized entropies, margins, and algorithms. In: Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics. Proceedings of Machine Learning Research, vol. 89, pp. 606–615. PMLR (2019)
7. Bromley, J., Guyon, I., LeCun, Y., Säckinger, E., Shah, R.: Signature verification using a "siamese" time delay neural network. In: Advances in Neural Information Processing Systems. vol. 6. Morgan-Kaufmann (1993)
8. Burkard, R.E.: Quadratic Assignment Problems, pp. 2741–2814. Springer New York (2013)
9. Burkard, R.E., Çela, E.: Linear assignment problems and extensions. In: Handbook of Combinatorial Optimization (1999)
10. Caron, M., Bojanowski, P., Joulin, A., Douze, M.: Deep clustering for unsupervised learning of visual features. In: European Conference on Computer Vision (2018)
11. 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. vol. 33, pp. 9912–9924 (2020)
12. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning. pp. 1597–1607 (2020)
13. Chen, T., Kornblith, S., Swersky, K., Norouzi, M., Hinton, G.E.: Big self-supervised models are strong semi-supervised learners. Advances in neural information processing systems **33**, 22243–22255 (2020)
14. Chen, X., Fan, H., Girshick, R., He, K.: Improved baselines with momentum contrastive learning. arXiv preprint arXiv:2003.04297 (2020)
15. Chen, X., He, K.: Exploring simple siamese representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15750–15758 (2021)
16. Ermolov, A., Siarohin, A., Sangineto, E., Sebe, N.: Whitening for self-supervised representation learning. In: International Conference on Machine Learning. pp. 3015–3024 (2021)
17. Finke, G., Burkard, R.E., Rendl, F.: Quadratic assignment problems. In: Surveys in Combinatorial Optimization, North-Holland Mathematics Studies, vol. 132, pp. 61–82. North-Holland (1987)
18. Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. arXiv preprint arXiv:1803.07728 (2018)
19. Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P., Buchatskaya, E., Doersch, C., Avila Pires, B., Guo, Z., Gheshlaghi Azar, M., et al.: Bootstrap your own latent-a new approach to self-supervised learning. Advances in neural information processing systems **33**, 21271–21284 (2020)

20. Hazan, T., Keshet, J., McAllester, D.: Direct loss minimization for structured prediction. In: Advances in Neural Information Processing Systems. vol. 23 (2010)
21. 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. pp. 9729–9738 (2020)
22. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
23. Hermans, A., Beyer, L., Leibe, B.: In defense of the triplet loss for person re-identification. arXiv preprint arXiv:1703.07737 (2017)
24. Hinton, G., Vinyals, O., Dean, J., et al.: Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 (2015)
25. Hjelm, D., Fedorov, A., Lavoie-Marchildon, S., Grewal, K., Bachman, P., Trischler, A., Bengio, Y.: Learning deep representations by mutual information estimation and maximization. ICLR (2019)
26. Huang, Z., Hu, P., Zhou, J.T., Lv, J., Peng, X.: Partially view-aligned clustering. Advances in Neural Information Processing Systems **33**, 2892–2902 (2020)
27. Huynh, T., Kornblith, S., Walter, M.R., Maire, M., Khademi, M.: Boosting contrastive self-supervised learning with false negative cancellation. Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV) (2022)
28. Jaiswal, A., Babu, A.R., Zadeh, M.Z., Banerjee, D., Makedon, F.: A survey on contrastive self-supervised learning. Technologies **9**(1) (2020)
29. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
30. Kolesnikov, A., Zhai, X., Beyer, L.: Revisiting self-supervised visual representation learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (2019)
31. Li, W., Zhao, R., Xiao, T., Wang, X.: Deepreid: Deep filter pairing neural network for person re-identification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 152–159 (2014)
32. Lin, Y., Gou, Y., Liu, Z., Li, B., Lv, J., Peng, X.: Completer: Incomplete multi-view clustering via contrastive prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
33. Linsker, R.: Self-organization in a perceptual network. IEEE Computer **21**, 105–117 (03 1988). <https://doi.org/10.1109/2.36>
34. Loshchilov, I., Hutter, F.: Sgdr: Stochastic gradient descent with warm restarts. arXiv preprint arXiv:1608.03983 (2016)
35. Martins, A.F.T., Astudillo, R.F.: From softmax to sparsemax: A sparse model of attention and multi-label classification. In: Proceedings of the 33rd International Conference on International Conference on Machine Learning (ICML) (2016)
36. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. arXiv preprint arXiv:1301.3781 (2013)
37. Misra, I., Maaten, L.v.d.: Self-supervised learning of pretext-invariant representations. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6707–6717 (2020)
38. Niculae, V., Martins, A., Blondel, M., Cardie, C.: Sparsemap: Differentiable sparse structured inference. In: International Conference on Machine Learning (ICML). pp. 3799–3808 (2018)
39. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018)

40. Patacchiola, M., Storkey, A.: Self-supervised relational reasoning for representation learning. In: Advances in Neural Information Processing Systems (2020)
41. Peng, X., Huang, Z., Lv, J., Zhu, H., Zhou, J.T.: Comic: Multi-view clustering without parameter selection. In: International Conference on Machine Learning (ICML) (2019)
42. Tschannen, M., Djolonga, J., Rubenstein, P.K., Gelly, S., Lucic, M.: On mutual information maximization for representation learning. In: International Conference on Learning Representations (ICLR) (2020)
43. Tschantzidis, I., Joachims, T., Hofmann, T., Altun, Y.: Large margin methods for structured and interdependent output variables. Journal of Machine Learning Research **6**(50), 1453–1484 (2005)
44. Wang, T., Isola, P.: Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In: International Conference on Machine Learning (ICML) (2020)
45. Wang, X., Zhang, R., Shen, C., Kong, T., Li, L.: Dense contrastive learning for self-supervised visual pre-training. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
46. Yang, M., Li, Y., Hu, P., Bai, J., Lv, J.C., Peng, X.: Robust multi-view clustering with incomplete information. IEEE Transactions on Pattern Analysis and Machine Intelligence (2022)
47. Yang, M., Li, Y., Huang, Z., Liu, Z., Hu, P., Peng, X.: Partially view-aligned representation learning with noise-robust contrastive loss. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
48. Zbontar, J., Jing, L., Misra, I., LeCun, Y., Deny, S.: Barlow twins: Self-supervised learning via redundancy reduction. In: International Conference on Machine Learning (ICML) (2021)