

Representation Learning via Consistent Assignment of Views to Clusters

Thalles Silva
University of Campinas
Campinas, SP

thalles.silva@students.ic.unicamp.br

Adín Ramírez Rivera*[†]
Reykjavik University
Reykjavik
adinr@uio.no

ABSTRACT

We introduce Consistent Assignment for Representation Learning (CARL), an unsupervised learning method to learn visual representations by combining ideas from self-supervised contrastive learning and deep clustering. By viewing contrastive learning from a clustering perspective, CARL learns unsupervised representations by learning a set of general prototypes that serve as energy anchors to enforce different views of a given image to be assigned to the same prototype. Unlike contemporary work on contrastive learning with deep clustering, CARL proposes to learn the set of general prototypes in an online fashion, using gradient descent without the necessity of using non-differentiable algorithms or K-Means to solve the cluster assignment problem. CARL surpasses its competitors in many representation learning benchmarks, including linear evaluation, semi-supervised learning, and transfer learning. Code at: <https://gitlab.com/mipl/carl/>.

CCS CONCEPTS

• **Computing methodologies** → **Semi-supervised learning settings; Learning latent representations;**

KEYWORDS

Contrastive learning, learning through clustering, representation learning

ACM Reference Format:

Thalles Silva and Adín Ramírez Rivera. 2022. Representation Learning via Consistent Assignment of Views to Clusters. In *The 37th ACM/SIGAPP Symposium on Applied Computing (SAC'22), April 25–29, 2022, Virtual Event*. ACM, New York, NY, USA, Article 4, 9 pages. <https://doi.org/10.1145/3477314.3507267>

1 INTRODUCTION

Unsupervised visual representation learning focuses on creating meaningful representations from data and inductive biases. Lately, methods based on Siamese neural networks [4] and contrastive

loss functions [8, 22] have significantly reduced the gap between supervised and unsupervised based representations. Indeed, for some downstream tasks, unsupervised-based representations already surpass their supervised counterparts [7]. In computer vision, approaches to unsupervised representation learning can be categorized into three groups: (1) contrastive learning methods using instance discrimination, (2) clustering-based methods, and (3) a mixture of the two.

Recent state-of-the-art unsupervised representation learning rely on contrastive learning [8, 9, 11, 22, 31, 34]. These methods optimize an instance discrimination pretext task where each image and its transformations are treated as individual classes. They compare feature vectors of individual images intending to organize the feature space such that similar concepts are placed closer while moving different ones farther.

On the other hand, traditional clustering methods aim to learn the data manifold by comparing groups of features that share semantic structure based on a distance metric. When combined with deep learning, clustering methods are often designed as two-step algorithms. First, the complete dataset is clustered, and then the meta clustering information, e.g., prototypes and pseudo-labels, are used as supervisory signals in a posterior optimization task [2, 5, 6, 37].

Recent work has attempted to combine the benefits of contrastive learning and clustering [7, 27]. In particular, Expectation-Maximization (EM) approaches alternate between finding the clusters and maximizing the mutual information (MI) between the embeddings and the cluster centroids [27]. Inspired by them, our work merges the benefits of both approaches by bridging the gap between clustering and contrastive learning. On the one hand, we use unsupervised clustering dynamics to generate robust prototypes that organize the feature space. On the other, we use contrastive learning to compare the distributions of the views' assignments w.r.t. the clusters. Our experiments show that we can learn unsupervised visual representations that outperform existing methods by mixing both approaches.

We can think of contrastive learning as learning clustered representations at the image level. However, given the nature of the task, these clusters fail to capture semantic information from the heterogeneous unknown classes since the learned clusters only comprise representations from synthetic views of an image. Moreover, since contrastive learning methods handle different images as negatives in the training process, even if two distinct images share the same class information, their representations will be pushed farther apart from each other. In the end, each image will have its own cluster structure.

We propose an alternative method to learn high-level features

*Also with University of Oslo.

[†]Also with University of Campinas.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SAC '22, April 25–29, 2022, Virtual Event,

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8713-2/22/04...\$15.00

<https://doi.org/10.1145/3477314.3507267>

by clustering views based on consistent assignments. Unlike previous work that uses K -Nearest Neighbors [1] or K -Means [28] as priors to enforce (learn) a cluster mapping, our method learns the prototypes online. While contrastive learning with instance discrimination [8, 22] poses a classification pretext task at the image embedding level, we propose a pretext task that operates on the assignment of views to a set of clusters. Rather than directly maximizing similarities between image embeddings, we force the distribution of positive views' assignments to be consistent among a set of finite learnable prototypes. If the number of prototypes equals the number of observations in the dataset, we would be forcing each cluster only to contain synthetic views of a given observation. This is equivalent to contrastive learning with instance discrimination. However, if we set the number of prototypes to be smaller than the number of observations in the dataset, by the pigeonhole principle, the learned prototypes will not only cluster different views of an image together, but they will also aggregate representations of different images that are similar enough to be assigned to the same cluster.

Regarding our contributions, (i) we propose a learning framework that leverages current contrastive learning methods with clustering-based algorithms to improve the learned representations. Unlike contemporary work, our method proposes to learn the clusters' assignments in an online fashion using gradient descent with no need for pre-clustering steps e.g., K -Means or non-differentiable methods to solve the clustering assignment problem, such as the Sinkhorn-Knopp algorithm [14]. (ii) We contrast high-level structures (the distributions of the views over the cluster assignments) instead of low-level ones (such as the representations). Reasoning in this high-level space gives the representations more robustness that translates to better performance in downstream tasks. And, (iii) unlike contrastive learning methods that can be viewed as learning clusters containing only synthetic transformations of a given image, our learned prototypes do not need to hold the semantics of the data but rather become energy anchors that self-organize the space to learn better representations. Moreover, the proposed loss function does not require a more extensive set of negative representations, which avoids the common problem of treating representations of the same class as negatives.

2 RELATED WORK

Self-supervised learning relates to the idea of extracting supervisory signals from the data. Methods including relative patch prediction [16], jigsaw puzzle [30], rotation prediction [18], and colorization [39] propose manually crafted pretext tasks that, when optimized, can make a deep neural network learn useful representations without the need of human-annotated datasets. Most of these methods work with the same principle. They corrupt the input with stochastic random transformations and challenge the network to predict some property of the corrupted input.

One such pretext task is instance discrimination [17, 36]. It describes a classification task in which each image is treated as a unique class and, therefore, stochastic transformations of the same image, called views, should belong to the same class. Dosovitskiy et al. [17] proposed to optimize this task by learning a linear classifier where the number of output classes matches the number of

observations in the dataset. Following, Wu et al. [36] proposed to use a Noise-Contrastive Estimation (NCE) approximation of the non-parametric softmax classifier that could scale to large datasets.

Currently, contrastive learning [21] methods rely on an NCE based loss function called InfoNCE [31, 33]. Recent work describes optimizing the InfoNCE loss through the lens of maximizing the MI between representations of the same image [3, 24, 25]. InfoNCE characterizes an $M + 1$ classification task where a pair of positive examples needs to be identified among the set of M negatives. In practice, the success of InfoNCE requires a large number of negative examples. Nonetheless, since negatives are usually randomly sampled from the dataset, it often leads to a false-negative problem where representations from images of the same class are treated as negatives [32].

He et al. [22] and Chen et al. [10] presented MoCo, a contrastive learning framework that employs an additional momentum encoder to provide consistent instance representations for the InfoNCE loss. Chen et al. [8] presented SimCLR, a Siamese-based [4] contrastive learning method trained with InfoNCE that relies on large batch sizes to draw a high number of negative samples. BYOL [20] proposes a framework that does not require negative samples and learns visual representations by approximating augmented views of the same data point using an ℓ_2 loss in the latent space. Unlike contrastive learning with instance discrimination, we seek to learn prototype vectors that act as anchors in the embedding space. We use these anchors as energy beacons for the images. Our goal is to use the energy distributions induced by the similarity of the images w.r.t. the prototypes to find representations that share similarities in the feature space.

Recent work proposed clustering-based methods for deep unsupervised representation learning [2, 6, 7, 27, 37]. DeepCluster [5] learns representations by predicting cluster assignments. Once per epoch, the algorithm clusters the entire dataset using K -Means and uses the pseudo-labels as supervised signals in an additional classification task. One of the limitations of this approach is that the classification layer needs to be reinitialized once per clusterization. DeeperCluster [6] builds on top of Caron et al.'s [5] work and presents an algorithm to combine hierarchical clustering with unsupervised feature learning using the rotation prediction pretext task [18]. Similarly, Prototypical Contrastive Learning (PCL) [27], formulates a self-supervised visual representation learning framework as an Expectation-Maximization algorithm.

Our method utilizes a conceptually distinct methodology. Unlike Caron et al.'s [5] and Li et al.'s [27] works, our method does not require a pre-clusterization step of the entire corpus, which vastly reduces memory and computing power requirements. Moreover, since we do not use K -Means clustering as a proxy to learn an additional task, we do not need to reinitialize any layers during optimization, nor is our method susceptible to limitations and assumptions implied by the K -Means algorithm. Instead, we propose to learn the prototypes end-to-end by consistently enforcing different views of the same image to be assigned to the same prototypes. Lastly, our method does not rely on handcrafted pretext tasks.

Similar to our work, Caron et al. [7] proposes an online clustering method to learn visual representations by contrasting cluster assignments on a second set of latent variables. To avoid collapsing modes where all observations are assigned to a few classes,

they use the Sinkhorn-Knopp algorithm [14] to solve the cluster assignment problem over the latent variables to guide the encoder during learning. Unlike Caron et al.'s [7] work, our method learns cluster assignments end-to-end via gradient descent and avoids trivial solutions by enforcing a regularization term over the cluster assignments of views in a given batch.

Van Gansbeke et al. [35] presented a two-step algorithm for unsupervised classification and proposed the SCAN-Loss (Semantic Clustering by Adopting Nearest neighbors) as part of the learning pipeline. The idea is to decouple feature learning and clustering by learning two encoders in distinct optimization stages. The algorithm extracts a set of nearest neighbors from each observation and uses them as priors to learn a second network for semantic clustering. The nearest neighbors are mined from representations of an encoder trained using a self-supervised pretext task. Our implementation builds on top of the SCAN-Loss, but unlike Van Gansbeke et al. [35], we employ a Siamese network architecture to learn representations via cluster assignments, end-to-end, without the necessity of optimizing for a second self-supervised task or mining of nearest neighbors. Our objective is to compare consistent assignments to a set of clusters for each view instead of comparing the underlying representations nor their true clusters. By contrasting high-level representations of the data, we obtain reliable data representations that perform well in downstream tasks.

3 PROPOSED METHOD

Our proposal, Consistent Assignment for Representation Learning (CARL), re-frames Contrastive Learning through a clustering perspective to learn robust representations (Section 3.1). CARL builds distributions of the similarities between the prototypes and the image's views (Section 3.2). To avoid collapsing the representation to a single prototype, we impose an uninformative prior to CARL's prototypes (Section 3.3). CARL jointly optimizes similarity to the learned prototypes and the uninformative prior (with a decay schedule for learning). Fig. 1 illustrates the overall idea.

3.1 Contrastive Learning from a Clustering Perspective

Let $\mathcal{X} = \{x_1, x_2, \dots, x_N\}$ be a dataset containing N unlabeled images. And, let a view $v_i = T(x_i)$ of an observation x_i be the application of a stochastic function T that is designed to change the content of x_i subjected to preserving the task-relevant information encoded in it. In practice, we can create as many views as needed by applying the stochastic function T . Recent contrastive learning methods learn visual embeddings by solving an instance discrimination pretext task that is usually optimized using the InfoNCE loss [31] defined as

$$\mathcal{L}_{\text{InfoNCE}} = -\log \frac{\exp\left(\text{sim}\left(z_i^a, z_i^+\right) / \tau\right)}{\sum_j^M \exp\left(\text{sim}\left(z_i^a, z_j\right) / \tau\right)}, \quad (1)$$

where z_i^a and z_i^+ are anchor and positive representations taken from an encoder function f such that $z_i^{(\cdot)} = f(v_i^{(\cdot)})$, τ is the temperature parameter, and $\text{sim}(\cdot, \cdot)$ is a similarity function, e.g., the cosine similarity. From a clustering perspective, the InfoNCE loss (1) is minimized when all possible variants $\{v_i^j\}_j$ of an image x_i are

clustered into the same prototype while representations from within a cluster are far apart from representations of other clusters.

We propose an approach where, instead of comparing against other instances [22] or prototypes of the classes [27], we learn a set of K general prototypes $C = \{c_1, c_2, \dots, c_K\}$, $K \ll N$, against which we compare the views to determine their similarity. Instead of maximizing a similarity function between positive embeddings of different views, as contrastive learning methods do, we maximize the similarity between assignment vectors of positive views to our general prototypes to promote consistency and confidence when assigning views to clusters. I.e., views must agree with high confidence in their cluster assignments. Our method learns the prototype matrix C online using gradient descent, and it does not require any pre-clusterization step as is typically the case for clustering-based representation learning algorithms [5, 27]. Moreover, since the training process is unsupervised, to avoid trivial solutions, where all images are assigned to only a handful of prototypes, we enforce a non-informative prior over the class distribution of the views to guide the learning process.

3.2 Consistent Assignment for Representation Learning

As with previous methods, we treat augmented versions of a given image as views and use them as positive examples for optimization. Our objective is to transform two positive samples into a distribution of their likelihood to belong to a set of K clusters. To do so, we encode each view through an encoder function $z_i = f(v_i) \in \mathbb{R}^d$. The encoder f comprises a backbone convolutional neural network, such as a ResNet [23], followed by a non-linear multilayer perceptron (MLP) head.

Our objective is not to cluster the data in an unsupervised manner but rather to learn a set of prototypes that will serve as anchors to differentiate views. We hypothesize that similar views should have similar assignments w.r.t. the prototypes. Hence, to convert the representations into these assignments, we first compare the representation z_i against all the prototypes to obtain an energy distribution

$$q_i[j] = \langle z_i, c_j \rangle, \quad (2)$$

where $q_i[j]$ is the j -th element of the energy distribution q_i for the i -th view. We learn the set of prototypes through a linear layer. To get the distribution of a view over all prototypes, we normalize the energy using the softmax function and obtain the posterior probability distribution, i.e., the probability of assigning the view v_i to the cluster k . Hence, our normalized probability for the k -th class given our view v_i is

$$p_i[k] = P(i \text{ assigned to } k | v_i) = \frac{\exp(q_i[k])}{\sum_{j=1}^K \exp(q_i[j])}, \quad (3)$$

where $q_i[j]$ denotes the j -th element of the i -th un-normalized vector output of the classifier for the respective view.

As mentioned before, our objective is to contrast the distributions of two views' likelihood w.r.t. the clusters. In CARL, the encoder and assigner operate in a Siamese setup where a pair of views from an image is independently transformed in its corresponding distribution. To ensure the similarity between the views, we optimize the views' distributions p_i^a and p_i^+ over the clusters in C , so that the

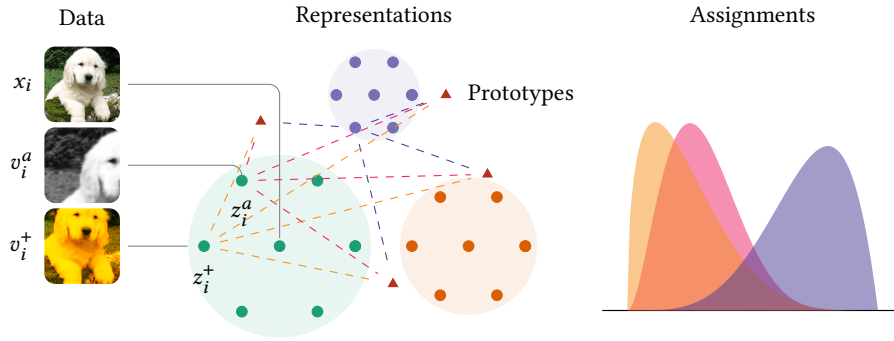


Figure 1: The data is first augmented and then encoded into their representations. We learn a set of **prototypes that serve as anchors to compute a distribution of assignments. For a **positive (augmented) pair of points** and for the rest of **negatives** we compute their agreement to the prototypes. We collect these comparisons as distributions that we later compare. We expect that the positive examples will have similar distributions over the assignments (cf. orange and pink distributions), while the negatives will not (cf. purple distribution).**

two distributions are consistent with one another. In other words, by learning a consistent assignment of views over the clusters, a given prototype will be invariant to augmented versions of an input image. Moreover, because the number of prototypes is smaller than the number of observations in the dataset, the clusters will contain different observations that share similarities in the embedding space.

We compute the similarity between the views' distributions as their dot product

$$\mathcal{L}_c = -\frac{1}{B} \sum_i \log \langle p_i^a, p_i^+ \rangle, \quad (4)$$

where B is the size of a minibatch over which we are aggregating the samples. In the ideal case of two one-hot vectors signaling the same perfect assignment, the dot product above yields its maximum value of one, and the negative log is minimized.

3.3 Preventing Trivial Solutions

Only forcing different views, v_i^a and v_i^+ , to have the same cluster assignment using our consistency loss (4) leads to finding a trivial solution where all representations z_i end up assigned to the same cluster (or just to a handful of them). To prevent such triviality, we force the distribution over the classes, P , to be uninformative by minimizing the Kullback-Leibler divergence w.r.t. a uniform distribution, U . Our regularization is

$$\mathcal{L}_{\text{KL}} = \text{KL}(P \parallel U) = \log(K) + \sum_{c \in C} \hat{p}^c \log(\hat{p}^c), \quad (5)$$

where \hat{p}^c is the expected distribution over a minibatch of size B ,

$$\hat{p}^c = \frac{1}{B} \sum_i p_i^c. \quad (6)$$

In other words, we maximize the Shannon entropy of the average distribution of the predictions. We can interpret the KL-divergence (5) as regularizing the encoder f to encourage the approximate posterior (3) to be closer to the uniform distribution.

Minimizing the KL-divergence (5) will force the predictions for a view v_i to be spread across all clusters. Since we do not know the underlying class distribution in advance, the KL-divergence (5) acts as a prior where we assume that the observations \mathcal{X} are uniformly assigned among all K prototypes.

By combining the consistency assignment loss (4) with the KL-divergence regularization (5) we obtain our final learning objective

$$\mathcal{L} = \mathcal{L}_c + \lambda_e \mathcal{L}_{\text{KL}}, \quad (7)$$

where λ_e is an epoch-dependent function that returns a scalar that prevents mode collapse at the beginning of training. We observed that training is very susceptible to such collapsing to a single assignment if not regularized. However, in practice, we noticed that keeping a large fixed value of λ_e during training also prevents the encoder from learning more complex representations. Thus, we recommend a function λ_e that decreases as training progresses. In theory, any decay schedule, such as an exponential or cosine, could be used. We propose a linear decay schedule

$$\lambda_e(a, b) = \begin{cases} b - \frac{b-a}{E}e & \text{if } e \leq E, \\ a & \text{otherwise,} \end{cases} \quad (8)$$

where b and a denote the start and ending values of the decay, E represents the number of epochs in which the decay will happen, and e is the epoch counter.

We noticed that mode collapse can happen in two scenarios: (1) if the regularizer is not added to the final loss, which is equivalent to setting the KL weight penalty $\lambda_e(0, 0)$, or (2) if the KL weight penalty λ_e is set to a value below a certain threshold, cf. Fig. 2. In both situations, CARL finds suboptimal solutions and learns bad representations. In practice, we found that slowly decreasing the KL weight penalty throughout training can work twofold since it increases the quality of the representations and prevents such trivialities during optimization. Refer to Section 4.1 and Fig. 2 for an in-depth analysis.

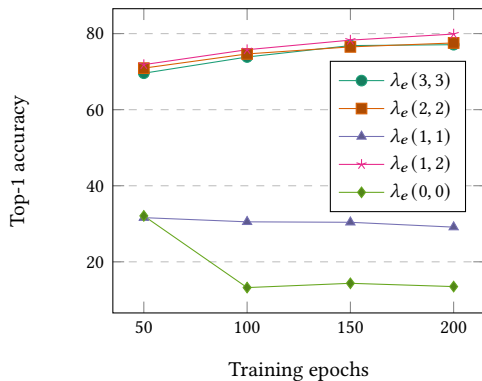


Figure 2: Effect of the uninformative prior’s scheduling λ_e on the overall performance in STL-10. Notice that a linear decay scheduling outperforms its constant counterparts. Two situations may result in non-optimal solutions: (1) a lower value of λ_e , and (2) not having the KL regularizer in the final loss.

4 HYPER-PARAMETER EXPLORATION

In this section, we evaluate the effects of the primary hyperparameters of our method. For exploring hyper-parameters, we learn representations using a ResNet-18 backbone trained for 150 epochs, and the KL weight penalty λ_e is linearly decayed over the first $E = 100$ epochs. To evaluate the multiple experimental setups, we train linear classifiers on top of the encoder’s frozen features following the linear evaluation protocol proposed by He et al. [22] and report average Top-1 accuracy over three independent runs.

4.1 Does Decreasing the KL Weight Penalty Improves Representation Learning?

The hyperparameter λ_e controls the contribution of the KL regularization (5) to the consistency loss (4). Especially at the beginning of training, a higher contribution for the KL term avoids mode collapsing, where the network optimizes the consistency loss (4) by assigning all observations to the same prototype. Van Gansbeke et al. [35] make similar claims for the entropy regularization in their SCAN-loss [35], and suggest a high (constant) value for the scalar hyperparameter λ_e to avoid such trivialities.

We hypothesize that keeping a high value of λ_e throughout training prevents the network from learning complex features. To verify this hypothesis, we trained CARL on the STL-10 [13] unsupervised dataset for 200 epochs. We measure the performance by training a linear classifier on top of the frozen features of the ResNet-18 backbone. We linearly decay the magnitude of the λ_e hyperparameter, following (8), from $b = 2$ to $a = 1$ over the first $E = 100$ epochs instead of keeping it constant for one of the experiments. As shown in Fig. 2, we observe that the quality of the representations learned by CARL benefits from decreasing the contribution of the KL regularization. However, if the λ_e is decreased below a certain threshold, mode collapse may happen.




Figure 3: Effect of learning a different number of prototypes on the quality of the representations. Empirical tests suggest an inverse U-shape curve where the optimal number of prototypes lies near one order of magnitude w.r.t. the actual number of classes of the dataset.

4.2 Does the Number of General Prototypes Influence the Quality of the Representations?

To evaluate the effect of learning a different number of prototypes, we trained CARL with ResNet-18 backbones on the CIFAR-10/100 datasets [26] for 150 epochs. For all experiments, the KL weight term λ_e starts as $b = 4$ and is linearly decreased to $a = 1.5$ over the first $E = 100$ epochs. Fig. 3 shows that over-clustering benefits the quality of the learned representations and that the optimal number of general prototypes depends on the number of actual classes of the dataset. The experiments suggest an inverse U-shape for both datasets, and the optimal number of general prototypes lies near an order of magnitude w.r.t. the actual number of classes.

4.3 Does the Batch Size Improves the Representations Learned by CARL?

We regularize CARL using a non-informative prior (5) applied over the input batch. This uniform prior prevents the views’ assignment from collapsing to a single prototype. We work over the expected distribution of assignments instead of each individual one. This expectation raises the question of how large a sample (batch size) should be to model different numbers of general prototypes.

Table 1 shows the effect of training CARL with different batch sizes and their relationship with the number of general prototypes. For this experiment, we trained CARL on CIFAR-100 for 150 epochs. Note that to compensate for the reduced number of iterations (as we increase the batch size), we scale the learning rate following the recipe proposed by Goyal et al. [19].

With smaller batches, we can see that the performance of the representations learned by CARL degrades as the number of prototypes increases. On the other hand, large batch sizes tend to retain high performance even when we increase the number of prototypes to very large numbers. Indeed, the gap in performance due to the batch size can reach nearly 6 points of average accuracy when optimizing a downstream task. We hypothesize that this is a

Table 1: The effect of training CARL (on CIFAR-100) with different batch sizes and general prototypes.

Batch	Number of prototypes						
	10	100	500	1000	3000	5000	10000
64	35.34	38.72	38.60	37.64	35.98	35.75	34.12
128	35.25	39.43	39.50	38.97	38.38	38.59	36.33
256	34.46	40.31	40.51	40.88	39.98	38.83	39.18
512	35.01	40.91	42.559	41.23	41.06	40.61	38.65
1024	34.67	41.259	41.44	42.049	41.479	41.548	40.689

direct effect of the sample size used to estimate the entropy in (5). In other words, as the distribution increases (with a higher number of general prototypes), the KL regularizer requires a larger batch size to better estimate the entropy and force it to be high (5).

5 UNSUPERVISED FEATURE EVALUATION

We evaluate CARL’s representations extracted from a ResNet-18 encoder and compare its performance with different state-of-the-art methods using the linear evaluation protocol [22], cross-domain transfer, semi-supervised learning, and fine-tuning.

5.0.1 Linear Evaluation. Models were trained for 200 epochs, closely following their original hyper-parameters. Refer to Appendix A for more details. Table 2 compare results with previous approaches on CIFAR-10/100 [26], STL-10 [13], and on a downsampled version (DS) of the ImageNet dataset [12]. We compare the representations learned by CARL against two methods based on instance discrimination with contrastive learning, i.e., SimCLR and MoCo v2, against two clustering-based methods, i.e., PCL and SwAV, and for completeness, we also include BYOL, which does not rely on negatives during optimization.

CARL outperforms all major competitors in most setups and only stays behind on early training cases (100 epochs) for the STL-10 and DS ImageNet datasets. Unlike SimCLR and MoCo v2, our method does not use negative examples, nor does it require a momentum-based target encoder to prevent collapsing [20].

We adapted PCL and SwAV to use a ResNet-18 encoder and kept official hyperparameters for a fair comparison. Most of these methods were only tested on large-scale datasets, such as the ImageNet [15], usually with very deep convolutional encoders. Thus, it naturally raises the question of how these methods would perform on smaller datasets, using faster and not-so-deep encoders, without enormous batch sizes and trained on modest GPUs. Our experiments highlight some of the limitations of training these data-hungry methods on limited resources.

5.0.2 Cross-domain transfer. To evaluate how well the representations learned by CARL transfer between different distributions, the first two columns of Table 3 show experimental results on cross-domain transfer. We took a self-supervised encoder trained on the CIFAR-10 dataset and used it as a feature extractor to train a linear classifier (for 50 epochs) to solve a 100-class classification task using the CIFAR-100 dataset (and vice versa). Representations learned by CARL outperform (on average) the main competitors in both setups.

5.0.3 Semi-supervised learning. The third and fourth columns of Table 3 display results on semi-supervised learning. First, we trained self-supervised models (for 200 epochs) on the unsupervised portion of the STL-10 dataset. Afterward, we fine-tuned the encoders (for 50 epochs) using different portions (1 % and 10 %) of labeled data from the supervised set of STL-10. Here, representations learned by CARL outperform competitors on the 10 % setup and perform on par with SimCLR on the more extreme case of only 1 % of labeled data.

5.0.4 Fine-tuning. Lastly, similar to the semi-supervised evaluation, the fifth column of table 3 shows results on fine-tuning the self-supervised encoder (trained on the unsupervised STL-10 for 200 epochs) using the entire supervised STL-10 dataset (for 50 epochs). CARL’s representations outperform competitors by at most 2.43 % average accuracy.

6 CONCLUSIONS

In this work, we presented Consistent Assignment for Representations Learning (CARL). An unsupervised method that learns visual representations by forcing augmented versions of an image to be consistently assigned over a finite set of learnable prototypes. Unlike contrastive learning methods, we propose a higher-level pretext task that operates over the distributions of views instead of directly optimizing the view’s embeddings. Our method differs from recent work that merges clustering with contrastive learning since it does not require pre-clusterization steps or non-differentiable algorithms to solve the cluster assignment problem. Instead, we learn a set of general prototypes that act as energy anchors for the views’ representations, entirely online using gradient descent. We studied some of the main components of CARL and the effects of different configurations of hyper-parameters. Lastly, our results show that representations learned by CARL surpass many of the recent contrastive learning and clustering methods without resorting to a large number of negative samples or extra encoders. The representations are qualitatively measured using many datasets and benchmarks, including linear evaluation, semi-supervised and finetuning.

A IMPLEMENTATION DETAILS

For evaluating CARL’s performance against its competitors, all methods were trained for 200 epochs using the same cosine learning rate scheduler and the same batch size of 256 observations. If not specified otherwise, we kept all the hyperparameters fixed as the original implementations.

A.1 Datasets

To properly validate our model’s performance, we opted to use many standard computer vision datasets. For small to medium classes, we used the CIFAR-10/100 dataset family. Each dataset contains 60 000, 32×32 RGB images, with 50 000 images for training and 10 000 for testing. CIFAR-10 has 10 classes with 6000 images per class, and CIFAR-100 has 600 images for each of the 100 classes.

We used the STL-10 unsupervised and supervised sets for fine-tuning and semi-supervised experiments. The unsupervised portion contains 100 000, 96×96 RGB images, and the supervised set has 5000 images from 10 classes.

Table 2: Results are reported using mean accuracy (percentage) and standard deviation over three runs (except for DS ImageNet where we report a single run) on the linear evaluation protocol.

Method	CIFAR-10		CIFAR-100		STL-10		DS ImageNet	
	100	200	100	200	100	200	100	200
Supervised	87.76 ± 0.07	89.40 ± 0.22	59.23 ± 0.28	60.59 ± 0.09	73.11 ± 0.24	74.76 ± 0.25	35.06	43.50
BYOL [20]	68.90 ± 0.24	76.46 ± 0.37	37.26 ± 0.47	45.68 ± 0.16	76.58 ± 0.78	79.64 ± 0.20	28.12	30.96
SimCLR [8]	72.64 ± 0.34	75.41 ± 0.47	41.64 ± 0.14	44.94 ± 0.11	77.36 ± 0.39	80.57 ± 0.66	28.20	30.32
MoCo [22]	65.57 ± 0.65	71.62 ± 1.04	39.12 ± 0.43	44.35 ± 0.41	71.50 ± 0.96	75.46 ± 1.28	24.82	28.67
PCL [27]	66.37 ± 0.23	71.29 ± 0.54	35.97 ± 0.53	40.1 ± 07.6	69.61 ± 0.93	70.56 ± 0.41	22.60	26.18
SwAV [7]	72.75 ± 0.65	75.93 ± 0.48	39.98 ± 0.83	43.20 ± 0.35	70.89 ± 0.55	74.35 ± 0.47	18.39	25.31
CARL	73.39 ± 0.31	78.94 ± 0.52	42.91 ± 0.26	48.85 ± 0.79	76.9 ± 01.2	81.95 ± 0.10	24.62	31.86

Table 3: Results are reported using mean accuracy (percentage) and standard deviation over three runs. Best results in bold. Cross-Domain Transfer: We trained CARL on the unlabeled CIFAR-10 dataset and performed linear evaluation on the labeled CIFAR-100 (10 → 100), and vice versa (100 → 10). Fine-tune: We trained CARL on the unlabeled STL-10 dataset and fine-tune it on the labeled STL-10.

Method	Cross Domain Transfer		Semi-supervised		Fine-tuning
	C-10 → C-100	C-100 → C-10	1 %	10 %	100 %
Supervised	38.11 ± 0.41	68.54 ± 0.42	52.97 ± 0.07	69.87 ± 0.03	76.11 ± 0.01
BYOL [20]	43.47 ± 0.20	67.23 ± 0.45	36.64 ± 0.43	65.7 ± 09.8	85.72 ± 0.10
SimCLR [8]	40.39 ± 0.22	67.4 ± 03.0	53.66 ± 0.11	73.92 ± 0.26	86.97 ± 0.03
CARL	43.47 ± 0.02	67.89 ± 0.23	52.26 ± 1.14	74.85 ± 0.33	88.15 ± 0.37

For datasets with a significant number of classes, we used the Tiny-ImageNet and a downsampled version of the original ImageNet dataset, referred to as DS ImageNet. Tiny-ImageNet contains 100 000 64×64 RGB images balanced across 200 different categories. The downsampled version of ImageNet has 1000 classes with 1 281 167, 64×64 RGB images.

A.2 Backbones

For all experiments, the encoder f comprises a ResNet-18 backbone followed by a non-linear 2-hidden layer fully-connected network, as projection head. For CARL, SwAV, and BYOL, the projection head also has a batch normalization layer. The function f encodes an image x_i into a 128-dim representation. Specifically, the projection head receives a 512-dim vector from the final average pooling layer of the ResNet-18 encoder. The hidden layer of the projection head has 512 neurons. For all methods, the dimensionality of the embedding vector z_i and the complexity of both the encoder and projection head are equivalent.

A.3 Augmentations

To create views for optimizing CARL, we used the same pipeline of data augmentations proposed by Chen et al. [8]. First, we apply a random crop resize operation, which randomly extracts a portion of the image from 0.08 to 1.0 of the original size. Second, we apply a horizontal flipping operation with a 50 % probability. There is an 80 % probability of jittering the pixels of the image, altering its brightness, contrast, saturation, and hue. Then, a 20 % chance of a grayscale conversion, a 50 % chance of gaussian blurring, followed by normalization using the dataset’s mean and STD.

A.4 CARL

CARL is trained using SGD with momentum and a cosine learning rate decay scheduler [29] (without restarts) starting at 0.6 and decaying to 0.0006. We use a weight decay penalty of 5×10^{-4} and the LARS optimizer [38], but using pure SGD produces similar results. To choose the number of general-prototypes, we follow the findings from our hyper-parameter exploration, see Section 4. For CIFAR-10/100, and STL-10, we trained CARL using 100 prototypes for the first dataset and 300 for the last two. Following SwAV, for the downsampled ImageNet dataset, we used 3000 prototypes.

A.5 Other methods

To compare the performance of CARL with other implementations, except for MoCo, we developed our versions of SimCLR, BYOL, SwAV, and PCL by adapting their original code repositories and performing the minimal changes necessary to meet our requirements. The core changes to the official implementations regard (1) adaptation to other datasets, (2) support for single GPU training, (3) change of backbone encoder (from ResNet-50 to ResNet-18), (4) when necessary, change hyperparameters to achieve better results or to comply with our computational budget.

A.5.1 PCL. [27]. For training on CIFAR-10/100 and STL-10, the number of prototypes is linearly scaled based on the number of true classes. For CIFAR-10 and STL-10, we set the number of clusters to {250, 500, 1000}, and for CIFAR-100, we used {2500, 5000, 10000}.

A.5.2 SwAV. Caron et al. [7]. For CIFAR-10/100 and STL-10, we used the same number of prototypes used to train CARL. The number of prototypes was set to 100 for CIFAR-10, and 300 for both CIFAR-100 and STL-10. We used the standard configurations from their original repositories to train both PCL and SwAV on the downsampled version of the ImageNet dataset.

A.5.3 BYOL and SimCLR. We empirically found that removing the LARS optimizer (from their official implementations) produced better results in our datasets.

We performed a simple grid search to find an optimal learning rate of 0.03, which improved over original parameters.

Originally, BYOL [20] uses a denser representation vector of 256-dim and a more complex projection head with a hidden layer containing 4096 neurons. BYOL employs batch sizes of 4096 observations. All methods use the same projection head architecture, and the extra BYOL’s predictor head follows the same architecture as the projection head.

For SimCLR [8], we empirically found that a temperature parameter of 0.2, instead of 0.5 described in the paper [8], is best for our datasets. We could not reproduce the original training setups for SimCLR and BYOL due to limited hardware. To ensure a fair comparison, we used equal batch sizes of 256 for all implementations. This allows SimCLR to produce 130 560 pairs against 256 pairs used in CARL.

A.5.4 Supervised. To establish an upper bound for performance comparisons, we trained a supervised ResNet-18 using SGD and a learning rate of 0.03 with a cosine learning rate decay (without restarts) that decreases to 0. For data augmentation, we used a lighter version of the MoCo’s [22] augmentation pipeline composed of (1) random crop resize, (2) random horizontal flip (50 % chance), (3) random color jitter, (4) random gray scaling (20 % chance), and (5) random Gaussian blurring (20 % chance).

A.6 Computing infrastructure

Experiments have been conducted on a workstation with 4 NVIDIA GeForce RTX 2080 Ti GPUs (11 GB of RAM). We adapted all methods to fit in a single GPU.

ACKNOWLEDGMENTS

This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior—Brasil (CAPES)—Finance Code 001.

REFERENCES

- [1] Naomi S Altman. 1992. An introduction to kernel and nearest-neighbor nonparametric regression. *Am. Stat.* 46, 3 (1992), 175–185.
- [2] Yuki Markus Asano, Christian Rupprecht, and Andrea Vedaldi. 2019. Self-labelling via simultaneous clustering and representation learning. In *Inter. Conf. Learn. Represent. (ICLR)*.
- [3] Philip Bachman, R Devon Hjelm, and William Buchwalter. 2019. Learning representations by maximizing mutual information across views. In *Wksp. Adv. Neural Inf. Process. Sys. (NeurIPS)*.
- [4] Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah. 1994. Signature verification using a "siamese" time delay neural network. *Adv. Neural Inf. Process. Sys. (NeurIPS)* (1994), 737–737.
- [5] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. 2018. Deep clustering for unsupervised learning of visual features. In *European Conf. Comput. Vis. (ECCV)*. 132–149.
- [6] Mathilde Caron, Piotr Bojanowski, Julien Mairal, and Armand Joulin. 2019. Unsupervised pre-training of image features on non-curated data. In *IEEE/CVF Inter. Conf. Comput. Vis. Pattern Recog. (CVPR)*. 2959–2968.
- [7] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. 2020. Unsupervised learning of visual features by contrasting cluster assignments. In *Adv. Neural Inf. Process. Sys. (NeurIPS)*.
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *Inter. Conf. Mach. Learn. (ICML)*. PMLR, 1597–1607.
- [9] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. 2020. Big self-supervised models are strong semi-supervised learners. In *Wksp. Adv. Neural Inf. Process. Sys. (NeurIPS)*.
- [10] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. 2020. Improved Baselines with Momentum Contrastive Learning. *arXiv preprint 2003.04297* (2020).
- [11] Xinlei Chen and Kaiming He. 2021. Exploring Simple Siamese Representation Learning. In *IEEE/CVF Inter. Conf. Comput. Vis. Pattern Recog. (CVPR)*.
- [12] Patryk Chrabaszcz, Ilya Loshchilov, and Frank Hutter. 2017. A Downsampled Variant of ImageNet as an Alternative to the CIFAR datasets. *arXiv preprint 1707.08819* (2017).
- [13] Adam Coates, Andrew Ng, and Honglak Lee. 2011. An analysis of single-layer networks in unsupervised feature learning. In *Inter. Conf. Artif. Intell. Stat. (AISTATS)*. JMLR Workshop and Conference Proceedings, 215–223.
- [14] Marco Cuturi. 2013. Sinkhorn distances: Lightspeed computation of optimal transport. In *Adv. Neural Inf. Process. Sys. (NeurIPS)*, Vol. 26. 2292–2300.
- [15] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *IEEE/CVF Inter. Conf. Comput. Vis. Pattern Recog. (CVPR)*.
- [16] Carl Doersch, Abhinav Gupta, and Alexei A Efros. 2015. Unsupervised visual representation learning by context prediction. In *IEEE/CVF Inter. Conf. Comput. Vis. Pattern Recog. (CVPR)*. 1422–1430.
- [17] Alexey Dosovitskiy, Philipp Fischer, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. 2015. Discriminative unsupervised feature learning with exemplar convolutional neural networks. *IEEE Trans. Pattern Anal. Mach. Intell.* 38, 9 (2015), 1734–1747.
- [18] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. 2018. Unsupervised Representation Learning by Predicting Image Rotations. In *Inter. Conf. Learn. Represent. (ICLR)*. <https://openreview.net/forum?id=S1v4N2l0>
- [19] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. 2018. Accurate, Large Minibatch SGD: Training ImageNet in 1 Hour. In *IEEE Inter. Conf. Comput. Vis. Wksp. (ICCVW)*.
- [20] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, et al. 2020. Bootstrap your own latent: A new approach to self-supervised learning. In *Wksp. Adv. Neural Inf. Process. Sys. (NeurIPS)*.
- [21] Raia Hadsell, Sumit Chopra, and Yann LeCun. 2006. Dimensionality reduction by learning an invariant mapping. In *IEEE/CVF Inter. Conf. Comput. Vis. Pattern Recog. (CVPR)*, Vol. 2. IEEE, 1735–1742.
- [22] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *IEEE/CVF Inter. Conf. Comput. Vis. Pattern Recog. (CVPR)*. 9729–9738.
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *IEEE/CVF Inter. Conf. Comput. Vis. Pattern Recog. (CVPR)*. 770–778.
- [24] Olivier Henaff. 2020. Data-efficient image recognition with contrastive predictive coding. In *Inter. Conf. Mach. Learn. (ICML)*. PMLR, 4182–4192.
- [25] R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. 2019. Learning deep representations by mutual information estimation and maximization. In *Inter. Conf. Learn. Represent. (ICLR)*.
- [26] Alex Krizhevsky. 2009. *Learning multiple layers of features from tiny images*. Ph.D. Dissertation, University of Toronto.
- [27] Junnan Li, Pan Zhou, Caiming Xiong, Richard Socher, and Steven CH Hoi. 2021. Prototypical Contrastive Learning of Unsupervised Representations. In *Inter. Conf. Learn. Represent. (ICLR)*.
- [28] Stuart Lloyd. 1982. Least squares quantization in PCM. *IEEE Trans. Inf. Theory* 28, 2 (1982), 129–137.
- [29] Ilya Loshchilov and Frank Hutter. 2016. SGDR: Stochastic gradient descent with warm restarts. In *Inter. Conf. Learn. Represent. Wksp. (ICLRW)*.
- [30] Mehdi Noroozi and Paolo Favaro. 2016. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European Conf. Comput. Vis. (ECCV)*. Springer, 69–84.
- [31] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. In *Wksp. Adv. Neural Inf. Process. Sys. (NeurIPS)*.
- [32] Nikunj Saunshi, Orestis Plevrakis, Sanjeev Arora, Mikhail Khodak, and Hrishikesh Khandeparkar. 2019. A theoretical analysis of contrastive unsupervised representation learning. In *Inter. Conf. Mach. Learn. (ICML)*. PMLR, 5628–5637.
- [33] Yonglong Tian, Dilip Krishnan, and Phillip Isola. 2020. Contrastive multiview coding. In *European Conf. Comput. Vis. Wksp. (ECCVW)*.
- [34] Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. 2020. What makes for good views for contrastive learning. In *Wksp. Adv. Neural Inf. Process. Sys. (NeurIPS)*.
- [35] Wouter Van Gansbeke, Simon Vandenhende, Stamatios Georgoulis, Marc Proesmans, and Luc Van Gool. 2020. Learning to classify images without labels. In

- European Conf. Comput. Vis. Wksp. (ECCVW)*.
- [36] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. 2018. Unsupervised feature learning via non-parametric instance discrimination. In *IEEE/CVF Inter. Conf. Comput. Vis. Pattern Recog. (CVPR)*. 3733–3742.
 - [37] Xueting Yan, Ishan Misra, Abhinav Gupta, Deepti Ghadiyaram, and Dhruv Mahajan. 2020. ClusterFit: Improving Generalization of Visual Representations. In *IEEE/CVF Inter. Conf. Comput. Vis. Pattern Recog. (CVPR)*. 6509–6518.
 - [38] Yang You, Igor Gitman, and Boris Ginsburg. 2017. Large Batch Training of Convolutional Networks. *arXiv preprint 1708.03888* (2017).
 - [39] Richard Zhang, Phillip Isola, and Alexei A Efros. 2016. Colorful image colorization. In *European Conf. Comput. Vis. (ECCV)*. Springer, 649–666.