

How does Contrastive Learning Organize Images?

Yunzhe Zhang

Zhejiang University of Technology
288 Liuhe Rd

YunzheZh@outlook.com

Yao Lu

yaolu.zjut@gmail.com

Qi Xuan

xuanqi@zjut.edu.cn

Abstract

Contrastive learning, a dominant self-supervised technique, emphasizes similarity in representations between augmentations of the same input and dissimilarity for different ones. Although low contrastive loss often correlates with high classification accuracy, recent studies challenge this direct relationship, spotlighting the crucial role of inductive biases. We delve into these biases from a clustering viewpoint, noting that contrastive learning creates locally dense clusters, contrasting the globally dense clusters from supervised learning. To capture this discrepancy, we introduce the "RLD (Relative Local Density)" metric. While this cluster property can hinder linear classification accuracy, leveraging a Graph Convolutional Network (GCN) based classifier mitigates this, boosting accuracy and reducing parameter requirements. The code is available [here](#).

1. Introduction

In recent years, contrastive learning has revolutionized machine learning by effectively learning representation functions from unlabeled data. This method generates simple augmentations for each data point and enforces, through an appropriate loss function, that (1) augmentations of a single data point are clustered together and (2) augmentations of different data points remain distant. These representations yield competitive classification performance, even when using a linear classifier for various downstream tasks. This progress brings us closer to the long-standing aspiration of developing machine learning models capable of generalizing across diverse data distributions and tasks.

Researchers currently lack a conceptual framework for understanding such phenomena, which is not surprising considering the limited quantitative understanding of generalization for a single task and data distribution. However, even a partial conceptual understanding could advance the field, and researchers have started to address this challenge [2, 18, 41, 44].

Recent papers aim to guide the further development of

this emerging theory by (1) analyzing the properties of augmented views' distributions and making assumptions, and (2) mathematically proving that, given these properties, an optimal contrastive loss can guarantee linear classification accuracy to some extent. Initial works in this area, such as [37] and [27], assume that two positive samples, as augmented views of the same image, are (nearly) conditionally independent given the class y . In other words, augmented views from the same class share the same distribution, even if they are derived from different images. More recent works, like [44], adopt a more practical and weaker assumption: for images from the same class, their distributions of augmented views overlap, while those from different classes do not.

Papers following this paradigm exhibit a limitation. Although the mathematical proofs are rigorous, there is no guarantee that the assumptions made are correct. These assumptions are often based on intuition, which is understandable given our limited knowledge of the underlying data distributions. However, this means that such assumptions may not hold true in real-world scenarios. For example, [36] points out that the assumptions in [44] may not be valid, as there is little overlap between the distributions of augmented views in practice, regardless of the class they belong to.

The paper [36] also demonstrates that when augmentation distributions for inputs do not overlap (which almost holds in practice), any discussion that does not consider inductive bias will lead to vacuous guarantees on classification accuracy. This raises a natural question: what does optimal contrastive loss mean in real-world scenarios when inductive bias is taken into account?

To address this question, we investigate the role of inductive bias in optimal contrastive loss from a clustering perspective and argue that a Graph Convolutional Network (GCN)-based classifier [24] is more suitable for contrastive learning than a linear one. In summary, we make the following observations:

- Visually similar images have similar vector representations in contrastive learning, regardless of whether

they belong to the same class. Conversely, images from the same class that are visually different have very distinct representations.

- Due to this, all images from the same class form locally dense clusters in contrastive learning, as most images visually similar to a specific image usually belong to the same class. However, visually different images with the same class label remain distinct in vector space, resulting in clusters that are not globally dense. To quantitatively measure this property, we develop a metric called "RLD (Relative Local Density)" that assesses density in graph space.
- Considering that contrastive learning creates dense clusters (communities) in graph space, we employ a GCN classifier, which not only achieves higher accuracy compared to the linear classifier but also requires fewer parameters.

2. Related Works

In this section, we discuss the underlying principles of contrastive learning and review relevant literature on this topic.

Contrastive learning has achieved remarkable success in solving downstream tasks by learning representations from similar pairs of data obtained using temporal information [28, 43] or different views or augmentations of inputs [3, 6, 8, 13, 15, 19, 22, 40, 45]. By minimizing a contrastive loss [17, 31], the similarity of a representation across various 'views' of the same image is maximized, while minimizing their similarity with distracting negative samples.

Multiple views of a single data point can be naturally extracted from multimodal or multisensory data [1, 25, 30, 32, 33, 39], while for an image-only modality, they are typically constructed via local and global cropping [3, 21, 22, 31] or data augmentation [6, 11, 13, 19, 45]. Positive pairs correspond to views of the same data point, while negatives are sampled views of different data points (typically from the same mini-batch), although the necessity for negative samples has recently been questioned [8, 16, 46].

With the rapid development of contrastive learning, a number of studies have attempted to understand it. In line with the framework discussed in the introduction, [37] and [27] initially assumed that positive samples are (nearly) conditionally independent, and then established a bound connecting contrastive loss and classification performance. In [44], the author posited an overlap between intra-class augmented views and provided a bound for downstream performance in relation to optimal contrastive loss.

However, [36] contends that neglecting inductive biases of the function class and training algorithm fails to adequately explain the success of contrastive learning, and can

even provably result in vacuous guarantees in certain situations. Furthermore, the study also asserts that there is minimal overlap between the distributions of augmented images, which contradicts basic assumptions made in numerous previous works.

With this understanding of contrastive learning and the existing literature, we will now proceed to explore the role of inductive biases and their impact on clustering and classification performance.

3. Preliminaries

3.1. Cluster and Community Evaluation

In this paper, we define clusters as distinct subsets within a data partition, where each subset represents a group of similar data points. Cluster evaluation has been a long-standing question in the literature [23, 38], focusing on determining whether clusters are well-structured [35] or consistent with ground truth labels [34]. In this study, we primarily concentrate on the former aspect, as ground truth labels themselves are treated as a partition, instead of those generated by the data-based cluster method.

The conventional notion of well-structured clusters, adopted in numerous works, is described as "dense and well-separated". A quantitative description of this concept is the Calinski-Harabasz score, also known as the Variance Ratio Criterion, introduced by [5]. Given a set of data points $\{x_1, x_2, \dots, x_n\}$ and the cluster indices they belong to $\{y_1, y_2, \dots, y_n\}$, the Calinski-Harabasz score can be defined as:

$$\text{CH score} = \frac{\sum_{i=1}^k n_i \|\mu_i - \mu\|^2 / (k - 1)}{\sum_{i=1}^k \sum_{x \in C_i} \|x - \mu_i\|^2 / (n - k)} \quad (1)$$

where k is the number of clusters, n is the total number of data points, n_i is the number of data points in cluster C_i , μ_i is the centroid of cluster C_i , and μ is the centroid of all data points. In essence, the Calinski-Harabasz score is the ratio of the sum of between-cluster dispersion to the sum of within-cluster dispersion for all clusters, where dispersion is defined as the sum of squared distances.

However, as we will illustrate later, the Calinski-Harabasz score, based on the conventional concept of clusters, does not correspond to the clusters formed by contrastive learning, necessitating the consideration of an alternative notion of clusters.

In this context, communities can be viewed as analogs to clusters but are based on graph space rather than Euclidean space, offering increased flexibility. The most commonly used metric for assessing community quality is modularity [9], which quantifies the density of connections within communities relative to the expected density of connections between randomly placed nodes. For graph G with adja-

gency matrix $A \in \mathbb{R}^{N \times N}$, modularity Q is:

$$Q = \frac{1}{2m} \sum_{i,j} \left[A_{ij} - \frac{k_i k_j}{2m} \right] \delta(y_i, y_j) \quad (2)$$

where m is the sum of edge weights in the graph, k_i and k_j are the degrees of nodes i and j , respectively, y_i and y_j are the community assignments of nodes i and j , and $\delta(y_i, y_j)$ is the Kronecker delta function, which equals 1 if $y_i = y_j$ and 0 otherwise. Our experiments demonstrate that, with the proper construction of G , modularity can better reveal underlying cluster structures in contrastive learning.

3.2. Graph Convolutional Network

Graph Convolutional Networks (GCNs) are a class of neural networks specifically designed for handling graph-structured data. They have gained considerable attention in recent years due to their ability to learn and generalize information from graph data [4, 10, 24].

The term GCN can refer to all graph convolutional networks or a specific architecture of a graph convolutional network proposed by [24]. In this paper, we focus on the latter for simplicity, as it serves as the origin of many other GCN models and captures the key characteristics of the broader class of GCN models.

Typical components of graph-structured data include node features $X \in \mathbb{R}^{N \times C}$ and an adjacency matrix $A \in \mathbb{R}^{N \times N}$, where N refers to the number of nodes and C refers to the dimensions of node features. The adjacency matrix captures the relationships between nodes in the graph.

A GCN aims to learn a function that maps the input node features X to a new set of node features $Z \in \mathbb{R}^{N \times F}$, where F is the output feature dimension. The GCN proposed by [24] performs graph convolutions in the following manner:

$$Z = \sigma \left(\tilde{A} X W \right), \quad (3)$$

where $\tilde{A} = \hat{D}^{-\frac{1}{2}} \hat{A} \hat{D}^{-\frac{1}{2}}$, with $\hat{A} = A + I$ being the adjacency matrix with added self-connections and \hat{D} is the diagonal degree matrix of \hat{A} , having elements $\hat{D}_{ii} = \sum_j \hat{A}_{ij}$. The matrix $W \in \mathbb{R}^{C \times F}$ is a learnable weight matrix, and $\sigma(\cdot)$ is the activation function applied element-wise.

For deeper GCN architectures with L layers, the output feature matrix $Z^{(L)}$ can be computed through a series of graph convolutions:

$$Z^{(l)} = \sigma \left(\tilde{A} Z^{(l-1)} W^{(l)} \right), \quad (4)$$

where $l = 1, \dots, L$, $Z^{(0)} = X$, and $W^{(l)} \in \mathbb{R}^{F_{l-1} \times F_l}$ are the learnable weight matrices for each layer. The dimensions of the input and output features change through layers, with $F_0 = C$ and $F_L = F$.

In this work, we introduce a minor modification to the GCN by incorporating a learnable scale parameter α in each layer, which transforms the adjacency matrix as $\hat{A} = \alpha A + I_N$. This adaptation enables the model to dynamically balance the significance of self-connections and connections to neighboring nodes, potentially enhancing the network’s capability to extract valuable information from the graph structure. Consequently, the updated adjacency matrix \tilde{A} becomes:

$$\tilde{A} = \hat{D}^{-\frac{1}{2}} (\alpha A + I_N) \hat{D}^{-\frac{1}{2}}, \quad (5)$$

With this modified \tilde{A} , the GCN can better adjust to the underlying structure of the graph during training, as it learns the optimal value of α alongside the weight matrices W .

We employ a GCN as the classifier due to its compatibility with graph space rather than Euclidean space, making it a better fit for the properties of clusters formed by contrastive learning. Experimental results demonstrate that the GCN consistently outperforms a linear classifier across various scenarios, further justifying its use in this context.

4. Experimental Settings

In our analysis of contrastive learning models, several models were meticulously trained as subjects using standard procedures. This section delineates the experimental settings employed, leveraging the CIFAR-10 dataset [26]. We adopt the framework and augmentation proposed in [7] and incorporate the prediction head improvements suggested by [14], using a representation space dimension of 512. The model architectures tested include ResNet18, ResNet101, and a modified Vision Transformer with a parameter count similar to ResNet101 [12, 20]. For comparison, we also train three models with identical architectures using supervised learning on CIFAR-10, maintaining a representation dimension of 512. All models presented in this paper utilize the AdamW optimizer [29]. Further information regarding model architectures and hyperparameters can be found in the appendix. All experiments are done on one A100 GPU.

In subsequent sections, image vector representations are derived from raw images in the dataset using both contrastive and supervised models, consistent with standard evaluation methods. For contrastive models, all features are before the prediction head and normalized to unit vectors due to the undefined lengths during the training process.

We emphasize that our experimental settings are not exhaustively tuned for optimal performance. As such, achieving state-of-the-art results in **not our primary objective**. Instead, our main focus lies in uncovering novel insights into the inner workings of contrastive learning, rather than proposing practical improvements or techniques.

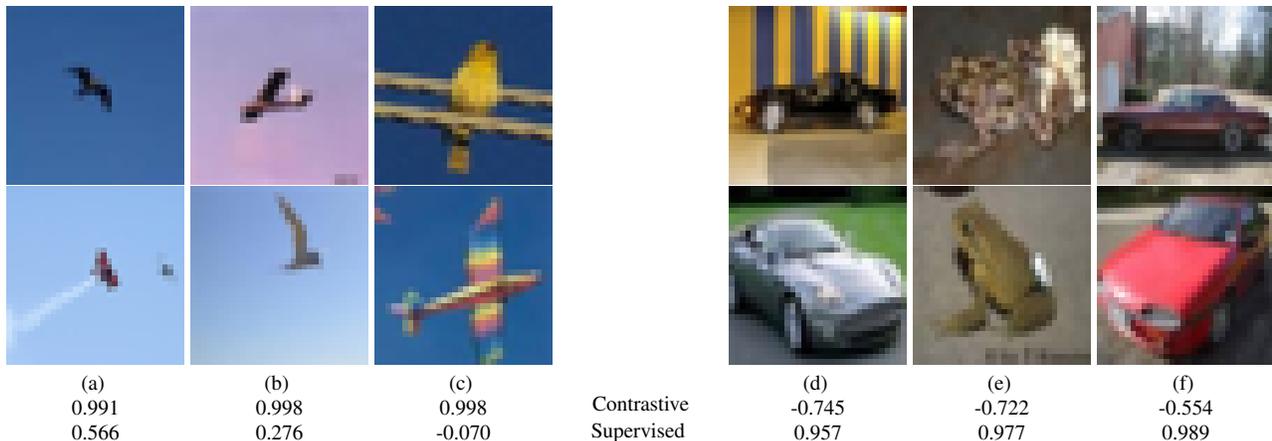


Figure 1. **Exploring Image Pairs and Their Cosine Similarities in Representations.** (a) displays the image pair with the highest similarity across different labels in contrastive ResNet18, while (b) demonstrates the same for ResNet101, and (c) for ViT. In contrast, (d) illustrates the image pair with the lowest similarity within the same class in contrastive ResNet18, (e) for ResNet101, and (f) for ViT. The first line of the numbers beneath each image pair indicates their similarity in contrastive space, and the second line denotes the similarity in supervised space (using a supervised model with the same architecture). In (a), the top image portrays a bird, and the bottom one depicts an airplane. The blue sky backgrounds and dark main bodies contribute to their visual similarity. Contrastive ResNet18 has difficulty distinguishing them, whereas supervised ResNet18 can. This observation also applies to (b) and (c). Conversely, (d), (e), and (f) present pairs of visually distinct images (including hue, shape, perspective, etc.) within the same class, which are recognized by supervised models but not by contrastive models.

5. Micro View: Image Pairs Under Contrastive Learning

In this section, we delve into the distinct behaviors of contrastive and supervised learning methods in organizing images within the representation space. Our analysis encompasses a comprehensive examination of how these methods group images based on visual similarities and class labels. We begin by demonstrating how contrastive learning effectively groups visually similar images together, irrespective of their class labels. We then proceed to discuss the organization of images in supervised learning, which primarily focuses on class labels while often overlooking visual similarities within and across classes. Furthermore, we investigate the properties of k-nearest neighbors in the contrastive representation space and introduce the Class Homogeneity Index (CHI) to measure the average proportion of k-nearest neighbors that share the same class label. By comparing CHI across diverse models under contrastive and supervised learning, we provide insights into the local and global class-wise homogeneity of neighborhoods in the representation space and how architectural choices influence the behavior of contrastive learning.

5.1. Contrastive Representation Space and Visual Similarity

we analyze specific image pairs from the CIFAR-10 training dataset to illustrate the properties of contrastive representation space. Fig. 1a, Fig. 1b, and Fig. 1c display im-

age pairs with the highest cosine similarity in contrastive representation space, even though they belong to different classes. Conversely, Fig. 1d, Fig. 1e, and Fig. 1f present image pairs with the lowest cosine similarity in contrastive representation space, despite having the same class label.

It is important to note that image pairs with high similarity in contrastive space exhibit visual resemblance, regardless of their class affiliation. On the other hand, image pairs that share the same label but display divergent visual appearances have relatively low similarity in contrastive space. This pattern is not observed in supervised space, where class labels rather than visual likenesses primarily influence similarity.

We hypothesize that this is a result of the unsupervised nature of contrastive learning, which relies solely on the information inherent in the samples. In contrast, supervised learning depends on the provided labels to organize images.

5.2. K-Nearest Neighbors in Contrastive Representation Space

As we have seen in the previous subsections, contrastive learning tends to group visually similar images together, regardless of their class labels. However, empirically, it is reasonable to assume that for most images, their visually similar counterparts belong to the same class. Building on our former observations and the insights from the analysis of visually similar image pairs, we can then infer that most representations have neighbors with the same label.

| k | Contrastive | | | Supervised | | |
|--------|-------------|-----------|--------|------------|-----------|--------|
| | ResNet18 | ResNet101 | ViT | ResNet18 | ResNet101 | ViT |
| k=1 | 0.8938 | 0.8818 | 0.8730 | 0.9680 | 0.9611 | 0.9903 |
| k=10 | 0.8762 | 0.8662 | 0.8473 | 0.9674 | 0.9611 | 0.9893 |
| k=100 | 0.8457 | 0.8438 | 0.8047 | 0.9669 | 0.9605 | 0.9886 |
| k=500 | 0.7785 | 0.7945 | 0.6675 | 0.9666 | 0.9605 | 0.9881 |
| k=1000 | 0.7094 | 0.7267 | 0.5273 | 0.9663 | 0.9606 | 0.9879 |

Table 1. **CHI across diverse models under contrastive and supervised learning.** The table presents the CHI for diverse models—ResNet18, ResNet101, and ViT—under both contrastive and supervised learning conditions, across different k values. CHI measures the average proportion of k -nearest neighbors that share the same class label.

To show that, in this subsection, we delve into the properties of the k -nearest neighbors for each data point within the contrastive representation space. We use cosine similarity as our distance metric and compute the average proportion of k -nearest neighbors that share the same class, a metric we term the Class Homogeneity Index (CHI). We explore the effects of varying k values ($k = 1, 10, 100, 500, 1000$) on CHI and summarize our findings in Tab. 1.

In light of the findings presented in Tab. 1, several implications emerge. First, for smaller k values, the high CHI observed in contrastive learning underscores that nearest neighbors typically belong to the same class, which exactly matches our earlier inference.

Second, the steep decrease in CHI as k increases in contrastive learning compared to supervised learning reveals an intriguing attribute. The neighborhood’s class-wise homogeneity in the contrastive representation space seems to diminish more rapidly as we expand our neighborhood size. This implies that while contrastive features exhibit a local class-wise homogeneity, they do not extend this effect at a more global scale, unlike the supervised features.

Third, the most noticeable decline in CHI is observed for the ViT. This could indicate that its weaker inductive bias is a contributing factor. Therefore, the behavior of contrastive learning might be shaped not only by the data but also by architectural choices and their inherent inductive biases.

6. Macro View: Contrastive Learning Forms Locally Dense Clusters

Expanding from our previous analysis of individual images and their nearest neighbors within the representation space, we now widen our scope to a macroscopic perspective. The focus of this paper now shifts towards understanding how contrastive learning arranges data on a larger scale, specifically at the cluster level. Here, a “cluster” signifies the grouping formed by all data points belonging to the same class.

To characterize these clusters, we present a novel descriptor: “Locally Dense”. A cluster is deemed “Locally Dense” if data points within the cluster are closely packed,

creating a dense local neighborhood. Consequently, for a specific data point, its nearest neighbors are predominantly from its own cluster, leading to high class homogeneity within local regions of the cluster. To quantitatively measure the local density of a specific cluster, we propose a new graph-based metric called Relative Local Density (RLD).

Following this, we provide examples to illustrate what locally dense clusters appear like, and how they contrast from the typical “dense and well-separated” clusters. Importantly, “dense and well-separated” refers to clusters that are not just densely packed internally (locally dense), but are also clearly separated from other clusters, implying a global density and separation.

Lastly, we will substantiate that clusters formed by contrastive learning are locally dense but not globally, while clusters in supervised learning exhibit both local and global density and separation.

6.1. Relative Local Density: A Quantitative Measure for Locally Dense Clusters

While the qualitative notion of locally dense clusters provides a conceptual foundation, it leaves room for a concrete, quantifiable measure. Thus, we introduce Relative Local Density (RLD), a novel metric aimed at capturing the notion of local density. This subsection details the construction of RLD, illuminating its correspondence to local density and its utility in comparing learning methods.

The computation of RLD involves several stages, each contributing to the encapsulation of local density. Initially, we construct a similarity matrix, S , for a given set of data points X . Each entry, S_{ij} , captures the pairwise distance between data points X_i and X_j , normalized by the mean distance and the square root of the dimension of X to ensure scale-invariance:

$$S_{ij} = \frac{-dist(X_i, X_j)}{\sqrt{dim(X)} \cdot mean(dist(X_p, X_q) : p \neq q)} \quad (6)$$

This process converts distances into similarity scores, with higher scores indicating closer data points. The ma-

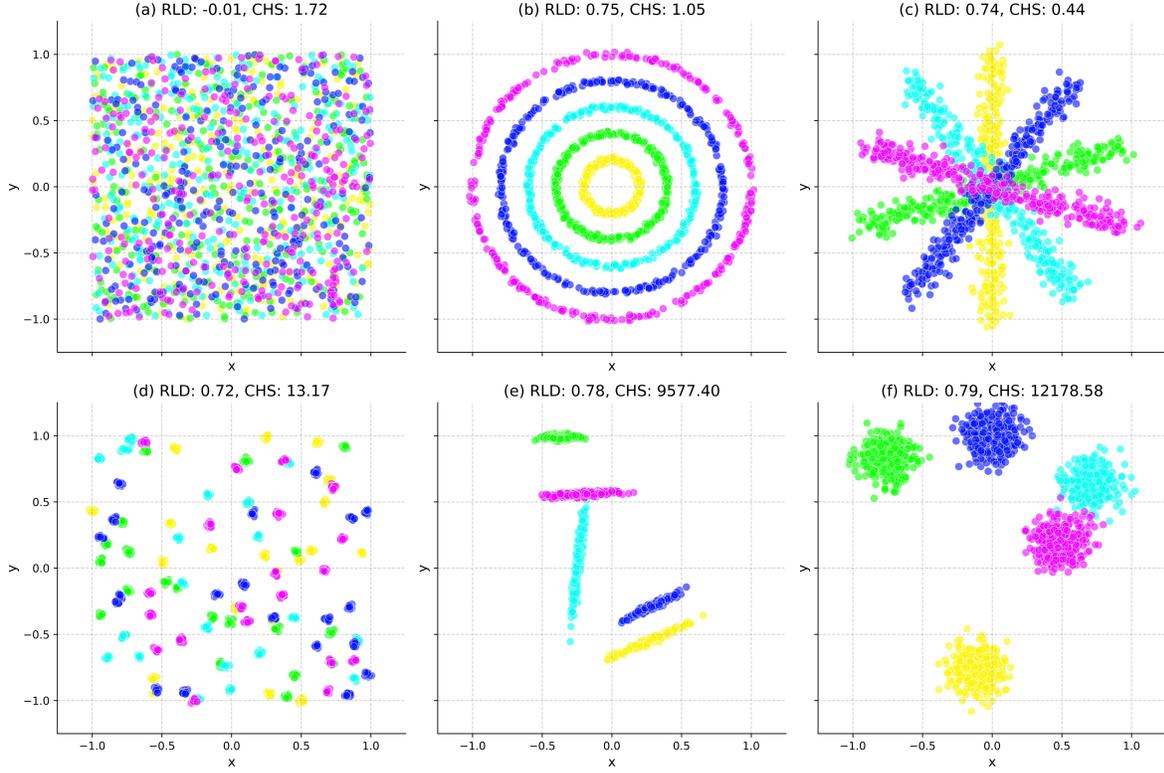


Figure 2. **Varied Cluster Configurations and Their Corresponding RLDs and CH scores.** Each subfigure, from (a) to (f), symbolizes a unique cluster configuration. Subfigure (a) presents data points uniformly scattered without discernible clustering, whereas (b) demonstrates a circular pattern with identifiable local clusters. Subfigure (c) showcases clusters arranged along distinct lines, and (d) displays small, tightly grouped local clusters dispersed throughout the plot. Subfigure (e) illustrates clusters formed along random lines, while (f) portrays multiple Gaussian distributions, each constituting a separate cluster. These examples underscore the varied characteristics of local and global densities and highlight the limitations of relying solely on global metrics, like the CH score, for evaluating cluster quality.

trix diagonal elements are set to negative infinity, ensuring a data point does not regard itself as its neighbor.

The similarity matrix is then transformed into an adjacency matrix, A , which encapsulates the relationships between data points in the feature space. A temperature parameter, T , modulates this transformation, balancing the emphasis on local and global structures. A higher T yields a matrix with a more global structure, while a lower T retains more local information. The matrix is normalized, each entry divided by T and exponentially transformed to accentuate differences between data points:

$$A_{ij} = \frac{n^2 \cdot \exp(S_{ij}/T)^2}{\sum_{k=1}^n \exp(S_{ik}/T)^2 \cdot \sum_{k=1}^n \exp(S_{kj}/T)^2} \quad (7)$$

The final RLD computation step measures the modularity of the adjacency matrix concerning class labels, y , with Eq. (2). Modularity quantifies the strength of a graph’s division into clusters. A high modularity score signifies dense intra-cluster connections and sparse inter-cluster connections, aligning with our local density intuition.

As a cluster evaluation metric, RLD offers several advantages:

1. **Local structure emphasis:** Unlike global metrics like the CH score, RLD captures local structure details, offering a more granular data organization understanding.
2. **Differentiability:** RLD’s full differentiability enables integration into gradient-based optimization processes.
3. **Scale invariance:** RLD’s insensitivity to the scale and number of data points enhances its versatility across diverse datasets.
4. **Graph techniques compatibility:** RLD’s graph-based nature permits integration with other graph techniques, such as community detection.
5. **Comparability:** The modularity scaling to $(-0.5, 1)$ standardizes RLD, facilitating comparisons. Its application enables direct cluster comparison, as scores are relative to the same scale.

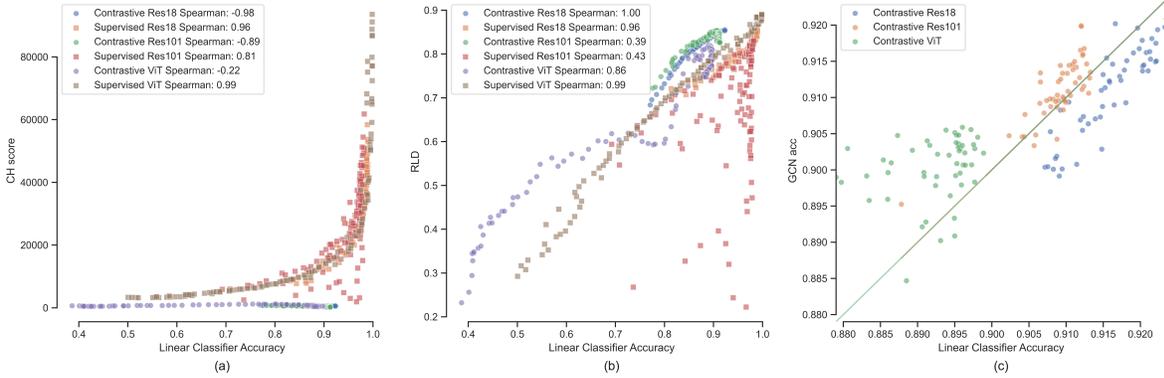


Figure 3. **Comparative Analysis of Cluster Evaluation Metrics and Classifier Accuracy across Models.** The figure presents three scatter plots, each point represents a model, comparing different metrics with linear classifier accuracy for contrastive and supervised learning methods across three models: ResNet18, ResNet101, and ViT. (a) shows the correlation between CH scores and linear classifier accuracy. (b) illustrates the relationship between Relative Local Density (RLD) and linear classifier accuracy. (c) provides a comparison between GCN classifier accuracy and linear classifier accuracy for contrastive learning models.

6.2. Visualizing Locally Dense Clusters: Examples and Comparisons

To gain a more intuitive understanding of local and global density concepts, it is important to visualize them before delving into the cluster analysis formed by contrastive and supervised learning methods. To facilitate this, we introduce illustrative examples that distinguish between locally dense clusters, as determined by RLD, and globally dense ones, as gauged by the CH score.

Fig. 2 presents six distinct cluster examples, each illustrating the divergent nature of local and global densities. Let’s examine some key takeaways:

1. A high RLD does not guarantee a high CH score. This fact is exemplified by clusters (b), (c), and (d), which, despite their high RLDs, register significantly lower CH scores compared to clusters (e) and (f).
2. The CH score fails to acknowledge the well-structured nature of clusters (b) and (c). This observation underscores the insensitivity of the CH score to certain types of cluster formations.
3. Clusters (b), (c), and (d) pose challenges for linear classifiers, which often struggle to define boundaries effectively. On the other hand, these classifiers are likely to perform well with clusters (e) and (f). Conversely, classifiers that operate in a neighbor-centric space, such as KNN, may deliver better separation results for clusters (b), (c), and (d).

6.3. Comparing Clusters Formed by Contrastive Learning and Supervised Learning

With a solid understanding of local and global density concepts, we can now turn our attention to examining the clusters produced by contrastive and supervised learning methods.

To provide a comprehensive and dynamic picture of cluster formation, we consider not only clusters created by fully trained models, but also those formed during the training process. This approach enables us to capture the evolution of clusters throughout the training phase and observe the unique ways in which different learning methods shape them over time.

In Fig. 3 (a) and (b), we present a real-case analysis. The RLDs ($T = 0.1$) and CH scores for all contrastive and supervised features are displayed in relation to linear classifier accuracy, along with their respective Spearman correlation coefficients. This collective view provides a comparative perspective on how these learning methods impact data organization.

As demonstrated in Fig. 3 (a), there is a visible positive correlation between the CH score of supervised learning and linear classifier accuracy. In contrast, contrastive learning shows a negative correlation and maintains a nearly constant CH score throughout the entire training process. This observation suggests that contrastive learning does not create globally dense clusters.

However, as depicted in Fig. 3 (b), the Spearman correlation between linear classifier accuracy and RLD remains positive for both learning methods, given the same architecture. This result implies that both contrastive learning and supervised learning form locally dense clusters.

Applying a GCN Classifier Given the challenges faced by linear classifiers in distinguishing locally dense clusters, as observed in Fig. 2 (b), (c), and (d), we decided to explore the use of a Graph Convolutional Network (GCN) classifier as an alternative. To facilitate this, we first constructed a graph identical to the one used in the computation of the

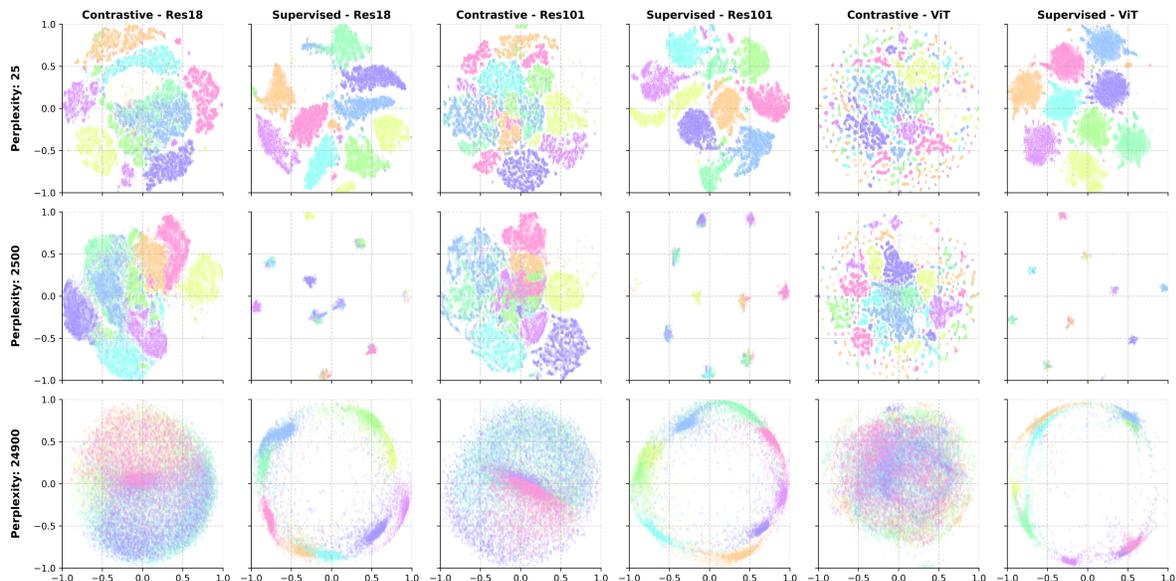


Figure 4. **t-SNE Visualization of Normalized CIFAR-10 Features Generated by Contrastive and Supervised Learning.** This figure presents t-SNE visualizations of 25,000 data points for contrastive (left column) and supervised (right column) learning methods, across three different model architectures: ResNet18, ResNet101, and ViT. Each row corresponds to a different perplexity value (25, 2500, and 24900), showcasing the influence of this parameter on the visualization. As perplexity increases, contrastive clusters begin to dissipate, while supervised clusters remain relatively stable, highlighting the distinct clusters formed by the two learning methods.

RLDs.

When assigning node features, we utilized a one-hot encoding scheme. Each node was assigned a one-hot encoded vector of length equal to the number of classes, denoted as $num_classes$. In this scheme, each vector represented a specific class label, with the element corresponding to the class label set to 1, and all other elements set to 0.

Following the training method described in [24], we applied a mask rate of 0.2 and used a four-layer GCN. We then plotted the model accuracy in Fig. 3 (c), comparing it with linear accuracy. The results indicated that for the ResNet101 and ViT models, GCN accuracy was slightly higher than linear accuracy for most of the models, although the opposite was observed for the ResNet18 model. This suggests that a GCN classifier can sometimes outperform linear classifiers when dealing with clusters formed by contrastive learning methods.

Visual Evidence via t-SNE Lastly, we employ t-SNE [42] as a visualization tool to further illustrate the differences between the features generated by contrastive and supervised learning methods. t-SNE is known for its ability to manage both 'local' and 'global' views via a tunable parameter known as perplexity. When we adjust this parameter in Fig. 4, it becomes evident that contrastive clusters start to fade as the perplexity increases. This effect, however, is not observed in the clusters formed by supervised

learning, thereby reinforcing our earlier observations about the distinct clusters formed by these two types of learning methods.

7. Conclusion and Future Directions

This study illuminates distinct differences in how contrastive and supervised learning algorithms structure data in representational space. We discover that contrastive learning primarily fosters 'locally dense' clusters while supervised learning generates globally dense clusters that align with class labels. The novel Relative Local Density (RLD) metric introduced in this study quantifies local density, offering a contrast to the traditional Calinski-Harabasz score. Implementing a Graph Convolutional Network (GCN) classifier demonstrates potential in tackling locally dense clusters, and the differences between learning methods are further highlighted through t-SNE visualizations.

Looking ahead, we propose two potential research directions. First, the development of more efficient classifiers tailored to clusters created by contrastive learning should be a priority. While GCNs show promise, their computational and memory requirements present challenges due to the need for an $N \times N$ adjacency matrix and a N node feature matrix. Second, creating innovative augmentation algorithms could help prevent models from misclassifying visually similar images from different classes by distinguishing these images in augmented views.

References

- [1] Relja Arandjelovic and Andrew Zisserman. Look, listen and learn. In *Proceedings of the IEEE international conference on computer vision*, pages 609–617, 2017. [2](#)
- [2] Sanjeev Arora, Hrishikesh Khandeparkar, Mikhail Khodak, Orestis Plevrakis, and Nikunj Saunshi. A theoretical analysis of contrastive unsupervised representation learning. In *36th International Conference on Machine Learning, ICML 2019*, pages 9904–9923. International Machine Learning Society (IMLS), 2019. [1](#)
- [3] Philip Bachman, R Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. *Advances in neural information processing systems*, 32, 2019. [2](#)
- [4] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. Spectral networks and deep locally connected networks on graphs. In *2nd International Conference on Learning Representations, ICLR 2014*, 2014. [3](#)
- [5] Tadeusz Caliński and Jerzy Harabasz. A dendrite method for cluster analysis. *Communications in Statistics-theory and Methods*, 3(1):1–27, 1974. [2](#)
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. [2](#)
- [7] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. [3](#)
- [8] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15750–15758, 2021. [2](#)
- [9] Aaron Clauset, Mark EJ Newman, and Cristopher Moore. Finding community structure in very large networks. *Physical review E*, 70(6):066111, 2004. [2](#)
- [10] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pages 3844–3852, 2016. [3](#)
- [11] Carl Doersch and Andrew Zisserman. Multi-task self-supervised visual learning. In *Proceedings of the IEEE international conference on computer vision*, pages 2051–2060, 2017. [2](#)
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2020. [3](#)
- [13] Alexey Dosovitskiy, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with convolutional neural networks. *Advances in neural information processing systems*, 27, 2014. [2](#)
- [14] Yann Dubois, Stefano Ermon, Tatsunori B Hashimoto, and Percy S Liang. Improving self-supervised learning by characterizing idealized representations. *Advances in Neural Information Processing Systems*, 35:11279–11296, 2022. [3](#)
- [15] Tianyu Gao, Xingcheng Yao, and Danqi Chen. Simcse: Simple contrastive learning of sentence embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, 2021. [2](#)
- [16] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020. [2](#)
- [17] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 1735–1742. IEEE, 2006. [2](#)
- [18] Jeff Z HaoChen, Colin Wei, Adrien Gaidon, and Tengyu Ma. Provable guarantees for self-supervised deep learning with spectral contrastive loss. *Advances in Neural Information Processing Systems*, 34:5000–5011, 2021. [1](#)
- [19] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. [2](#)
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [3](#)
- [21] Olivier Henaff. Data-efficient image recognition with contrastive predictive coding. In *International conference on machine learning*, pages 4182–4192. PMLR, 2020. [2](#)
- [22] R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. In *International Conference on Learning Representations*, 2018. [2](#)
- [23] Lawrence Hubert and Phipps Arabie. Comparing partitions. *Journal of classification*, 2:193–218, 1985. [2](#)
- [24] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2016. [1](#), [3](#), [8](#)
- [25] Bruno Korbar, Du Tran, and Lorenzo Torresani. Cooperative learning of audio and video models from self-supervised synchronization. *Advances in Neural Information Processing Systems*, 31, 2018. [2](#)
- [26] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical Report 0, University of Toronto, Toronto, Ontario, 2009. [3](#)
- [27] Jason D Lee, Qi Lei, Nikunj Saunshi, and Jiacheng Zhuo. Predicting what you already know helps: Provable self-supervised learning. *Advances in Neural Information Processing Systems*, 34:309–323, 2021. [1](#), [2](#)
- [28] Lajanugen Logeswaran and Honglak Lee. An efficient framework for learning sentence representations. In *International Conference on Learning Representations*, 2018. [2](#)

- [29] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2017. 3
- [30] Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. Howto100m: Learning a text-video embedding by watching hundred million narrated video clips. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2630–2640, 2019. 2
- [31] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 2
- [32] Andrew Owens and Alexei A Efros. Audio-visual scene analysis with self-supervised multisensory features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 631–648, 2018. 2
- [33] Adria Recasens, Pauline Luc, Jean-Baptiste Alayrac, Luyu Wang, Florian Strub, Corentin Tallec, Mateusz Malinowski, Viorica Pătrăucean, Florent Althé, Michal Valko, et al. Broaden your views for self-supervised video learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1255–1265, 2021. 2
- [34] Andrew Rosenberg and Julia Hirschberg. V-measure: A conditional entropy-based external cluster evaluation measure. In *Proceedings of the 2007 joint conference on empirical methods in natural language processing and computational natural language learning (EMNLP-CoNLL)*, pages 410–420, 2007. 2
- [35] Peter J Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65, 1987. 2
- [36] Nikunj Saunshi, Jordan Ash, Surbhi Goel, Dipendra Misra, Cyril Zhang, Sanjeev Arora, Sham Kakade, and Akshay Krishnamurthy. Understanding contrastive learning requires incorporating inductive biases. In *International Conference on Machine Learning*, pages 19250–19286. PMLR, 2022. 1, 2
- [37] Nikunj Saunshi, Orestis Plevrakis, Sanjeev Arora, Mikhail Khodak, and Hrishikesh Khandeparkar. A theoretical analysis of contrastive unsupervised representation learning. In *International Conference on Machine Learning*, pages 5628–5637. PMLR, 2019. 1, 2
- [38] Alexander Strehl and Joydeep Ghosh. Cluster ensembles—a knowledge reuse framework for combining multiple partitions. *Journal of machine learning research*, 3(Dec):583–617, 2002. 2
- [39] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. Videobert: A joint model for video and language representation learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7464–7473, 2019. 2
- [40] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*, pages 776–794. Springer, 2020. 2
- [41] Christopher Tosh, Akshay Krishnamurthy, and Daniel Hsu. Contrastive learning, multi-view redundancy, and linear models. In *32nd International Conference on Algorithmic Learning Theory*, 2021. 1
- [42] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(86):2579–2605, 2008. 8
- [43] Xiaolong Wang and Abhinav Gupta. Unsupervised learning of visual representations using videos. In *Proceedings of the IEEE international conference on computer vision*, pages 2794–2802, 2015. 2
- [44] Yifei Wang, Qi Zhang, Yisen Wang, Jiansheng Yang, and Zhouchen Lin. Chaos is a ladder: A new theoretical understanding of contrastive learning via augmentation overlap. In *International Conference on Learning Representations*, 2022. 1, 2
- [45] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3733–3742, 2018. 2
- [46] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning*, pages 12310–12320. PMLR, 2021. 2