

Joint Representation Learning and Novel Category Discovery on Single- and Multi-modal Data

Xuhui Jia¹ Kai Han^{1,2,3*} Yukun Zhu¹ Bradley Green¹

¹Google ²University of Bristol ³The University of Hong Kong

{xhjia, kaihanx, yukun, brg}@google.com

Abstract

This paper studies the problem of novel category discovery on single- and multi-modal data with labels from different but relevant categories. We present a generic, end-to-end framework to jointly learn a reliable representation and assign clusters to unlabelled data. To avoid over-fitting the learnt embedding to labelled data, we take inspiration from self-supervised representation learning by noise-contrastive estimation and extend it to jointly handle labelled and unlabelled data. In particular, we propose using category discrimination on labelled data and cross-modal discrimination on multi-modal data to augment instance discrimination used in conventional contrastive learning approaches. We further employ Winner-Take-All (WTA) hashing algorithm on the shared representation space to generate pairwise pseudo labels for unlabelled data to better predict cluster assignments. We thoroughly evaluate our framework on large-scale multi-modal video benchmarks Kinetics-400 and VGG-Sound, and image benchmarks CIFAR10, CIFAR100 and ImageNet, obtaining state-of-the-art results.

1. Introduction

With the tremendous advances in deep learning, recent machine learning models have shown superior performance on many tasks, such as image recognition [9, 30], object detection [54, 41], image segmentation [7], etc. While the state-of-the-art models might even outperform human in these tasks, the success of these models heavily relies on the huge amount of data with human annotations under the closed-world assumption. Applying deep learning in real (open) world brings many new challenges: it is cost-inhibitive to identify and annotate all categories, and new categories could keep emerging. Conventional methods struggle on handling unlabelled data from new categories [13]. On the flip side, real world provides rich unlabeled data, which are often multi-modal (e.g., video and audio), allowing more possi-

bilities for machine learning models to learn in a similar way as human. Indeed, humans learn from multi-modal data everyday with text, videos, audios, etc.

In this paper, we focus on automatically learning to discover new categories in the open world setting. Similar to recent works [18, 16] which transfer knowledge from labelled images of a few classes to other unlabelled image collections, we formulate the problem as partitioning unlabelled data from unknown categories into proper semantic groups, while some labelled data from other categories are available. This is a more realistic setting than pure unsupervised clustering which may produce equally valid data partitions following different unconstrained criteria (e.g., images can be clustered by texture, color, illumination, etc) and closed-world recognition which can not handle unlabelled data from new categories without any labels. Meanwhile, our setting is more similar to the human cognition process where humans can easily learn the concept of a new object by transferring knowledge from known objects.

Specifically, we introduce a flexible end-to-end framework to discover categories in unlabelled data, with the goal of utilizing both labelled and unlabelled data to build unbiased feature representation, while transferring more knowledge from labelled to unlabelled data. In particular, we extend the conventional contrastive learning [6, 19] to consider both instance discrimination and category discrimination to learn a reliable feature representation on labelled and unlabelled data. We also demonstrate that the cross-modal discrimination would further benefit representation learning on data with multi-modalities. To leverage more of unlabelled data, we employ the Winner-Take-All (WTA) hashing [48] on the shared representation space to generate pair-wise pseudo labels on-the-fly, which is the key for robust knowledge transfer from the labelled data to unlabelled data. With the weak pseudo labels, the model can be trained with a simple binary cross-entropy loss on the unlabelled data together with the standard cross-entropy loss on the labelled data. This way our model can simultaneously learn feature representation and perform the cluster assignment

* Corresponding author.

using an unified loss function.

The main contributions of the paper can be summarized as follows: (1) we propose a generic, end-to-end framework for novel category discovery that can be trained jointly on labelled and unlabelled data; (2) to the best of our knowledge, we are the first to extend contrastive learning in novel category discovery task by category discrimination on labelled data and cross-modal discrimination on multi-modal data; (3) we propose a strategy to employ WTA hashing on the shared representation space of both labelled and unlabelled data to generate additional (pseudo) supervision on unlabeled data; and (4) we thoroughly evaluate our end-to-end framework on challenging large scale multi-modal video benchmarks and single-modal image benchmarks, outperforming existing methods by a significant margin.

2. Related work

Our method is related to self-supervised learning, semi-supervised learning, and clustering, while different from each of them. We review the most relevant works below.

Self-supervised learning aims at learning reliable feature representations using the data itself to provide supervision signals during training. Many pretext tasks (e.g., relative position [11], colorization [49], rotation prediction [14]) have been proposed for self-supervised learning, showing promising results. Recently, the contrastive learning based methods, such as [19] and [6], have attracted lots of attention by its simplicity and effectiveness. The key idea of contrastive learning is *instance discrimination*, i.e., pulling similar pairs close and pushing dissimilar pairs away in the feature space. [27] studied the supervised contrastive learning on labelled data as an alternative of cross-entropy. With the labels, more positive pairs can be generated from the intra-class instances, enabling *category discrimination*. Noise-Contrastive Estimation (NCE) [15, 45] is an effective contrastive loss widely used in these methods. When handling multi-modal data like videos, different self-supervised learning methods have been proposed to exploit data of different modalities, such as [36, 1, 3, 2, 33]. Among them, [33] suggests that cross-modal discrimination can be adopted to improve the representation learning for downstream tasks like image recognition and object detection, which implies that good representations are shared between multiple views of the world. [43] shows that cross-view prediction outperforms conventional alternatives in contrastive learning on images, depth, video and flow, and more views can lead to better representation. In this paper, we consider the visual and audio modalities for cross-modal learning in videos, and present a new way to incorporate contrastive learning for both labelled and unlabelled data to bootstrap representation learning for novel category discovery.

Semi-supervised learning [4] considers the setting with labelled and unlabelled data. Specifically, the unlabelled data

are assumed to come from the same classes as the labelled data. The objective is to learn a robust model making use of both labelled and unlabelled data to avoid over-fitting to the labelled data. While this problem is well studied in the literature (e.g., [34, 42, 38]), existing methods can not handle unlabelled data from new classes. In contrast, our method is designed to discover new categories in the unlabelled data automatically.

Clustering, which aims at automatically partitioning the unlabelled data into different groups, has long been studied in the machine learning community. There are many classic methods (e.g., *k*-means [32], mean-shift [8]) and deep learning based methods (e.g., [47, 10, 37]) showing promising results. However, the definition of a cluster can be intrinsically ambiguous, because different criteria can be used to cluster the data. For example, objects can be clustered by color, shape or texture, and the clustering results will be different by taking different criteria, while these criteria cannot be predefined in the clustering methods. In this paper, we aim to learn these criteria implicitly from the labelled data and transfer them to the unlabelled data on-the-fly.

Until recently, the problem of discovering new categories in unlabelled data by exploiting the labelled data starts to draw attention. Particularly, the task of novel category discovery is formalized by [18, 16, 17]. [18] introduces a method to first pretrain the model on labelled data followed by fine-tuning with an unsupervised clustering loss. [16, 17] present a three-stage method, namely, self-supervised pre-training [14], supervised fine-tuning, and joint learning to transfer knowledge. [21] and [23], which are designed for general cross-task transfer learning, can also be applied to discover new categories in image datasets. These two methods require maintaining two models separately. One for binary label prediction and another for clustering. To our knowledge, the most relevant work to ours is [16], while our method significantly differs from it in several aspects. First, our method is an end-to-end trainable framework that can be easily adopted for practical use, while [16] requires three training stages, thus hampering their practical utilization; second, we conduct self-supervised learning jointly with new category discovery to bootstrap the semantic feature embedding that are more suitable for discovering new classes, rather than simply using self-supervision to provide only robust low-level features as in [16]. Third, we employ multiple random partial rankings for holistic comparison of two data points, rather than simply comparing the global ranking, which tends to be more vulnerable to noise. Last, our model can be flexibly applied on both single- and multi-modal data showing superior results, while [16] is only applicable to images. As will be shown in the experiments, our method substantially outperforms others in all benchmarks. Concurrent to our work, several other methods are also proposed for novel category discovery on images showing promising

performance by augmenting the training samples by mixing the labelled and unlabelled data [52], aggregating pseudo-positive pairs with contrastive learning [51], designing a unified objective [12], and leveraging local part information [50].

3. Method

Given an unlabelled collection of instances $x_i^u \in D^u$, our objective is to automatically partition these instances into C^u different semantic groups. We also assume that there is a labelled collection of instances $(x_i^l, y_i^l) \in D^l$ where $y_i^l \in \{1, \dots, C^l\}$, from which we want to transfer the knowledge to the unlabelled data so that the unlabelled data can be grouped into proper classes, while the classes in D^l and D^u are different but relevant. Each $x_i \in D^u \cup D^l$ can be either an image or a multi-modal video consisting of a visual stream x_i^v and the corresponding audio stream x_i^a .

Our approach is an end-to-end trainable framework that can jointly learn the representation and clustering assignment for the unlabelled data. Figure 1 shows the overview of our approach. Consider a video clip $x_i = (x_i^v, x_i^a)$, the visual encoder f_v and the audio encoder f_a first encode the visual and audio streams into two feature vectors z_i^v and z_i^a , which are then concatenated to form the global representation $z_i = [z_i^v, z_i^a]$ for the input video clip. A projection function η then project z_i to \bar{z}_i to fuse multi-modal feature in a compact representation space. Note that η is an identity mapping function for single-modal data. If x_i is from labelled data, \bar{z}_i will be sent to the linear head ϕ^l for supervised learning. Otherwise, it will be sent to the linear head ϕ^u for new category discovery. In addition, z_i^v and z_i^a are further encoded to \hat{z}_i^v and \hat{z}_i^a by h_v and h_a respectively for contrastive learning to bootstrap the representation learning on both labelled and unlabelled data. h_v and h_a are two MLP functions followed by ℓ_2 normalization, which is a common practice in contrastive learning.

The end-to-end training of the model consists of several important components including training on labelled data with full supervision, contrastive learning with instance and category discrimination on both labelled and unlabelled data, and training on unlabelled data with pair-wise pseudo labels transferred from the representation jointly learned by the previous two. To effectively transfer knowledge from the labelled to unlabelled data, we employ the winner-take-all (WTA) hashing algorithm on \bar{z}_i . Next, we will introduce our extended contrastive learning for novel category discovery by considering both instance and category discrimination, and knowledge transfer via the WTA hashing in more details.

3.1. Unified contrastive learning on labelled and unlabelled data

Given a batch of N data points that are randomly drawn from both D^u and D^l , i.e., $x_i \sim \text{Unif}(D^u \cup D^l)$ for all

the elements x_i in batch B , our objective is to extract as much information as possible from B to learn a representation that can be used to cluster unknown classes. Recent developments in contrastive learning focus on learning image representation under self-supervised scenario, where no labels are available, therefore it can be thought of as contrasting instance representations from a pair of data augmentation (positive) to those of other samples (negative). In our case, we have a mixed setting with both labelled and unlabelled data. We adopt the contrastive learning to jointly learn from the mixed dataset for novel category discovery. Next, we first demonstrate our unified contrastive learning on single-modal data, and then extend it to the more flexible multi-modal scenario.

3.1.1 Single-modal learning

Let x_i be a single-modal data point and $i \in \mathcal{N} = \{1, \dots, 2N\}$ be the index of pairs of augmented samples from the batch B . With the embedded representation \hat{z}_i of x_i , we adopt NCE [15] as our contrastive loss function for instance discrimination, which can be written as:

$$\mathcal{L}_i^{NCE-I} = -\log \frac{\exp(\hat{z}_i \cdot \hat{z}_{i'}/\tau)}{\sum_n \mathbb{1}_{[n \neq i]} \exp(\hat{z}_i \cdot \hat{z}_n/\tau)}, \quad (1)$$

where $\hat{z}_{i'}$ is the augmented counterpart of \hat{z}_i , $\mathbb{1}_{[n \neq i]}$ is an indicator function evaluating 1 iff $n \neq i$, and τ is a scalar temperature. This loss is widely used in conventional contrastive learning approaches.

In our case, some data points in the mini-batch B are accompanied with labels. If x_i is a labelled data point in the batch, besides considering only x_i and its transformed counterpart as a positive pair, other data points from the same class as x_i can also be paired with x_i to form more positive pairs to be pulled together, allowing category discrimination. The contrastive loss for category discrimination with these additional positive pairs can be written as:

$$\mathcal{L}_i^{NCE-C} = -\frac{1}{|Q(i)|} \sum_{q \in Q(i)} \log \frac{\exp(\hat{z}_i \cdot \hat{z}_q/\tau)}{\sum_n \mathbb{1}_{[n \neq i]} \exp(\hat{z}_i \cdot \hat{z}_n/\tau)}, \quad (2)$$

where $Q(i) = \{q \in \mathcal{N} \setminus i : y_q = y_i\}$ denotes the indices of other data points which have the same label as x_i in the batch B . Note that this is effective only for labelled data, as $Q(i) = \emptyset$ for unlabelled data. Therefore, the unified contrastive loss can be written as:

$$\mathcal{L}^{CL} = \frac{1}{2N} \sum_i^{2N} (\mathcal{L}_i^{NCE-I} + \mathcal{L}_i^{NCE-C}). \quad (3)$$

Note that there is only one positive pair in \mathcal{L}_i^{NCE-I} which aims at instance discrimination, whereas multiple other samples from same class are considered as positive in

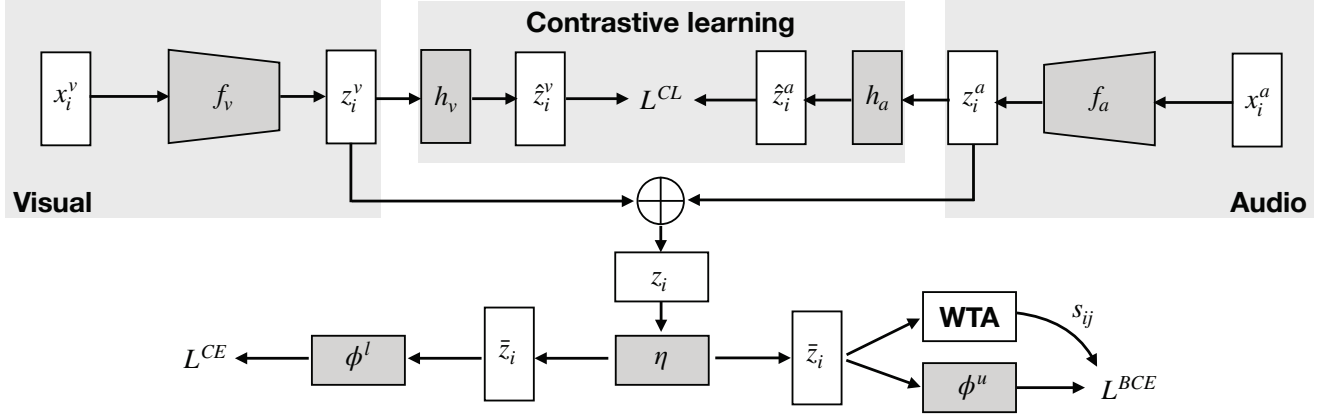


Figure 1: **Overview of our end-to-end framework.** For multi-modal videos, our framework consists of two feature encoders f_v and f_a , two MLPs h_v and h_a for contrastive learning, one fusion layer η , and two linear heads ϕ^l and ϕ^u for classification and clustering. The training signal for ϕ^u is obtained by WTA on-the-fly. For single-modal images, the audio encoder f_a is omitted and η turns to be an identity mapping function.

\mathcal{L}_i^{NCE-C} which aims at category discrimination. \mathcal{L}_i^{NCE-I} and \mathcal{L}_i^{NCE-C} have exactly the same denominator, which is the summation of $2N - 1$ scores for both positive and negative pairs. In this way, both labelled and unlabelled data are used for representation learning while making full use of the labels contained in the labelled data. \mathcal{L}_i^{NCE-C} is a critical complementary to \mathcal{L}_i^{NCE-I} , resulting in a more discriminative representation space for robust clustering as will be seen in the experiments.

3.1.2 Multi-modal learning

For multi-modal data like videos, besides the conventional within-modal contrastive learning for single-modal data like images, we can also have the additional cross-modal option. As noted by [2, 33], cross-modal agreement leads to better representation than within-modal agreement in self-supervised representation learning for supervised downstream tasks like object recognition and detection. However, our setting contains both labelled data and unlabelled data, resulting in a mixture of instance discrimination and category discrimination. It is not immediately obvious whether the within-modal or cross-modal choice is more effective under such a setting. For within-modal case, we can either discriminate visual or audio samples. Let the embedded representation for a multi-modal data point x_i be $\hat{z}_i = \{\hat{z}_i^v, \hat{z}_i^a\}$. We define modality selecting functions g_0 and g_1 to allow either within-modal discrimination (e.g., $g_0(\hat{z}_i) = \hat{z}_i^v$ and $g_1(\hat{z}_i) = \hat{z}_i^v$) or cross-modal discrimination (e.g., $g_0(\hat{z}_i) = \hat{z}_i^v$ and $g_1(\hat{z}_i) = \hat{z}_i^a$). The instance discrimination and category discrimination objectives in Eq. (1) and Eq. (2) can then be rewritten as

$$\mathcal{L}_i^{NCE-I} = -\log \frac{\exp(g_0(\hat{z}_i) \cdot g_1(\hat{z}_i)/\tau)}{\sum_n \mathbb{1}_{[n \neq i]} \exp(g_0(\hat{z}_i) \cdot g_1(\hat{z}_n)/\tau)}, \quad (4)$$

and

$$\mathcal{L}_i^{NCE-C} = -\frac{1}{|Q(i)|} \sum_{q \in Q(i)} \log \frac{\exp(g_0(\hat{z}_i) \cdot g_1(\hat{z}_q)/\tau)}{\sum_n \mathbb{1}_{[n \neq i]} \exp(g_0(\hat{z}_i) \cdot g_1(\hat{z}_n)/\tau)}. \quad (5)$$

As will be shown in the experiments, after investigating different strategies, we find that the cross-modal contrastive learning can produce better representation for novel category discovery. This might be due to the fact that cross-modal representations are intrinsically different, thus reducing the chance of obtaining trivial solutions such as discriminating objects by only verifying the most salient features in visual or audio modality alone. Therefore, unless stated otherwise, the modality selecting functions g_0 and g_1 choose visual and audio representations respectively in our formulation.

3.2. Knowledge transfer via Winner-take-all hash

To leverage the labelled data to help novel category discovery in unlabelled data, we transfer knowledge from labelled data to unlabelled data by adopting the Winner-Take-All (WTA) hash [48] during training. WTA is a sparse embedding method that maps the feature vectors into integer codes. In this work, we employ WTA to measure the similarity between each pair of unlabelled data points from new categories in the shared embedding space of both labelled and unlabelled data, so that we can transfer knowledge from the labelled categories to the unlabelled ones.

Winner-take-all hash [48] The idea of WTA is to measure similarity between high-dimensional feature vectors by comparing multiple partial ranking statistics. The WTA algorithm works as follows. First, we randomly generate a

set of H permutations $\mathcal{P} = \{\rho_1, \dots, \rho_H\}$. For an unlabelled sample $x_i \in D^u$, we extract its feature vector \bar{z}_i . We then apply each $\rho_h \in \mathcal{P}$ on \bar{z}_i to obtain a transformed feature vector $\rho_h(\bar{z}_i)$, i.e., a shuffled version of \bar{z}_i . Let c_i^h be the index of maximum value in the first k elements of $\rho_h(\bar{z}_i)$. We can then obtain the WTA hash code by $c_i = (c_i^1, \dots, c_i^H)$ for each x_i .

WTA for novel category discovery In our case, we employ WTA hash code to measure the similarity s_{ij} between x_i and x_j to generate pairwise pseudo labels for novel category discovery, by simply comparing their WTA hash codes c_i and c_j :

$$s_{ij} = \begin{cases} 1, & \mathbf{1}^T \cdot (c_i = c_j) \geq \mu \\ 0, & \text{otherwise} \end{cases}, \quad (6)$$

where μ is an empirical scalar threshold. Note that WTA is only applied during training to generate binary pseudo labels, and it is not needed at test time.

As discussed in [48], the precise values in the high-dimensional embedding is often not important, and the relative magnitude matters more. While requiring the total orderings to be identical is too strict for real application, as noise inevitably exists even using ranking statistics, therefore, WTA introduces multiple partial order statistics. Besides introducing more resilience to noise, we further emphasize that the partial orders for local rank correlation also captures the relative structural information of the objects. Intuitively, the global ranking statistics may only consider the most salient features in the embedding space, while the multiple partial orders spread more in the embedding space, thus capturing more structural information. As shown in [53], modern CNNs are very likely to make decisions by focusing on the salient patterns while overlooking the holistic structural composition. As will be seen in the experiments, WTA outperforms other alternatives for generating pairwise pseudo labels to discover new categories.

After applying the WTA algorithm, we can obtain pairwise similarity s_{ij} for each pair of unlabelled data points with Eq. (6). Assume we have M unlabelled samples in a mini-batch. By using these pairwise similarities as pseudo labels, we can then train the model with binary cross-entropy loss to simultaneously learn representation and cluster assignments on the unlabelled data from new classes:

$$\mathcal{L}^{BCE} = -\frac{1}{M^2} \sum_{i=1}^M \sum_{j=1}^M [s_{ij} \log \phi^u(\bar{z}_i)^\top \phi^u(\bar{z}_j) + \quad (7)$$

$$(1 - s_{ij}) \log(1 - \phi^u(\bar{z}_i)^\top \phi^u(\bar{z}_j))], \quad (8)$$

where $\phi^u : \mathbb{R}^d \rightarrow \mathbb{R}^{C^u}$ is a non-linear function mapping \bar{z}_i into an embedding space with the same dimension as the number of classes in the unlabelled data followed by

softmax normalization. In this way, we can obtain the cluster assignment for each unlabelled sample by indexing the location of the maximum value in $\phi^u(\bar{z}_i)$ after training, without requiring another offline clustering procedure for class assignment.

3.3. Joint learning objective

Inspired by the literature of semi-supervised learning [42], we include a consistency regularization loss on both labelled and unlabelled data. The purpose of such a consistency loss is to enforce the class predictions on a data point x_i and its transformed counterpart x'_i to be the same. This is especially important for unlabelled data samples. By enforcing the consistency, x_i and x'_i will be treated as a positive pair regardless of the WTA hash code, as the WTA hash code of x_i might be different from that of x'_i , thus smoothing the training. The consistency loss is commonly implemented as the mean squared error (MSE) between the class predictions. Let \mathcal{L}^{MSE} be the consistency loss and \mathcal{L}^{CE} be the cross-entropy loss on labelled data. The consistency loss between x_i and its transformed version x'_i is defined as

$$\mathcal{L}_i^{MSE} = (\phi(\bar{z}_i) - \phi(\bar{z}'_i))^2, \quad (9)$$

where ϕ is ϕ^l or ϕ^u depending on the input. The overall training loss of our end-to-end framework can then be written as

$$\mathcal{L} = \mathcal{L}^{CE} + \mathcal{L}^{BCE} + (1 - \omega(r))\mathcal{L}^{CL} + \omega(r)\mathcal{L}^{MSE}, \quad (10)$$

where $\omega(r)$ is a ramp-up function slowly increasing from 0 to 1 along with the training. In our experiment, we follow [31] to use $\omega(r) = \lambda e^{-5(1-\frac{r}{T})^2}$ where r , T and λ are current epoch number, total number of epochs and a positive scalar factor. We set $(1 - \omega(r))$ as the weight for contrastive learning. At the early stages, the cluster assignment predictions are noisy and we expect the model to focus more on representation learning, thus a higher weight is set to representation learning and a lower weight is set for consistency. In the late stages, the representation is good enough and we would like the model to focus on novel category discovery, therefore the weight is higher for consistency loss and lower for contrastive learning.

4. Experiments

Benchmarks and evaluation metric. We comprehensively evaluate our approach for novel category discovery on large-scale image benchmarks including ImageNet [9]/CIFAR-10 [28]/CIFAR-100 [28] and video benchmarks including Kinetics-400 [26]/ VGG-Sound [5]. We follow [16] to split the ImageNet/CIFAR-10/CIFAR-100 to have 30/5/20 classes in the unlabelled classes. For fair comparison with [18, 16, 22] three 30-class splits are used for ImageNet, and the results are averaged over the three splits. We

split Kinetics-400/VGG-Sound to have 50/39 classes in the unlabelled data, which are much more challenging than the image benchmarks. As only URLs are publicly available for video datasets, by the time of our experiments, we have 170k video clips with sound in Kinetics-400 and 183k videos with sound in VGG-Sound. To measure the novel category discovery accuracy, we adopt the widely used average clustering accuracy defined as: $\max_{e \in \mathcal{P}(C^u)} = \frac{1}{U} \sum_{i=1}^U 1\{y_i^u = e(\hat{y}_i^u)\}$, where y_i^u and \hat{y}_i^u denote the ground-truth label and predicted cluster assignment for each unlabelled data point, U is the total number of unlabelled instances in the whole dataset, $\mathcal{P}(C^u)$ denotes the set of all possible permutations of C^u elements, e is an arbitrary permutation in $\mathcal{P}(C^u)$. We obtain the optimal permutation e^* by Hungarian algorithm [29]. Note that as no supervision is used for unlabelled data, the same data are used for both training and evaluation following standard practice [25, 18].

Implementation details. We use R3D-18 [44, 24] as the video encoder and ResNet-18 [20] as the image and audio encoder. The feature vector dimension is 512 for both encoders. R3D-18 is an effective yet lightweight model for video recognition tasks, which has been shown effective for multi-modal self-supervised learning [36], allowing the use of a relatively larger batch size for contrastive learning. The choice of ResNet-18 for image datasets is to follow [18, 16] for fair comparison. The MLPs for contrastive learning consist of a hidden layer of size 512, a linear layer of size 128, and an ℓ_2 normalization layer [46]. The output dimension of the hidden layer η is 512. On image benchmarks, we follow SimCLR [6] to randomly apply cropping, resizing, horizontal flip, color distortion, and Gaussian blur for data augmentation. On video benchmarks, we follow [40] to use the input video and audio clips of 1 and 2 second duration respectively. Video frames are resized such that the shorter side has a size of 128, followed by random cropping with size 112. We preprocess the audio by randomly sampling within 0.5 second of the video and compute a log mel bank features with 257 filters and 199 time-frames, followed by SpecAugment [35]. During evaluation, we follow [36] to uniformly sample 10 clips from each video and take the mean score for prediction. For WTA, we set H to be equal to the feature dimension (i.e., 512), and follow [48] to set k as 4. For the threshold μ , we empirically set it to 240 in our experiments. We train our models with SGD [39] and use a batch size of 1024 for all benchmarks except CIFAR-10/CIFAR-100, which we use a batch size of 256. The models are trained with 8×8 TPUv2 Dragonfish devices.

4.1. Novel category discovery on image benchmarks

In table 1, we compare our approach with the k -means baseline and the previous state-of-the-art methods. By comparing row 10 with rows 1-6, we can clearly see that our method substantially outperforms the k -means baseline and

existing methods. For example, our method outperforms the previous state-of-the-art method [16], denoted as RS in row 5, by 3%/3.2%/4.2% on CIFAR-10/CIFAR-100/ImageNet. We also report the incremental learning result of RS [16] in row 6, where an additional step is introduced to create interactions between labelled and unlabelled heads. The k -means baseline in row 1 is evaluated on the features of unlabelled data extracted using the model pretrained on labelled data by cross-entropy loss. We further enhance the k -means baseline, KCL [21] and MCL [23] by introducing the same contrastive learning used in our framework (see rows 7-10). The contrastive learning successfully boosts the performance for all of them, but the results still largely lag behind our method. For example, after the enhancement, the ACC of KCL is increased to 73.9%/57.4%/74.3% on CIFAR-10/CIFAR-100/ImageNet, while the results of our method are 93.4%/76.4%/86.7%, demonstrating the effectiveness of our approach on image benchmarks.

Table 1: **Novel category discovery on image benchmarks.** ‘‘CL’’ denotes the contrastive learning we introduce in section 3.1, which performs both instance discrimination and category discrimination. Results on CIFAR-10/CIFAR-100 are averaged over 10 runs.

No	Method	CIFAR-10	CIFAR-100	ImageNet
(1)	k -means [32]	65.5 \pm 0.0%	56.6 \pm 1.6%	71.9%
(2)	KCL [21]	66.5 \pm 3.9%	14.3 \pm 1.3%	73.8%
(3)	MCL [23]	64.2 \pm 0.1%	21.3 \pm 3.4%	74.4%
(4)	DTC [18] w/ RotNet	88.7 \pm 0.3%	67.3 \pm 1.2%	-
(5)	RS [16]	90.4 \pm 0.5%	73.2 \pm 2.1%	82.5%
(6)	RS [16] w/ LL.	91.7 \pm 0.9%	75.2 \pm 4.2%	-
(7)	k -means [32] w/ CL	73.2 \pm 0.1%	58.8 \pm 1.5%	73.1%
(8)	KCL [21] w/ CL	73.9 \pm 0.1%	57.4 \pm 3.6%	74.3%
(9)	MCL [23] w/ CL	72.3 \pm 0.2%	50.8 \pm 3.1%	75.2%
(10)	RS [16] w/ CL	89.0 \pm 0.1%	54.6 \pm 1.8%	82.7%
(11)	Ours	93.4 \pm 0.6%	76.4 \pm 2.8%	86.7%

In fig. 2 we show the t-SNE projection for the features of data from the 5 unlabelled classes in CIFAR-10. The features are extracted using the model pretrained on the labelled data (fig. 2 left) and using our model (fig. 2 right) respectively. We can see that the embedding is rather cluttered using the model pretrained on the labelled data, while our model can successfully partition the unlabelled data into tight semantic groups.

4.2. Novel category discovery on video benchmarks

In table 2, we compare our framework with k -means baseline and [16] on Kinetics-400 and VGG-Sound benchmarks. In rows 1-3, we train the models on labelled data with cross-entropy loss, and extract the features on unlabelled data for k -means. We validate features from multi-modalities as well as each single-modality. The multi-modal features are more effective than the single-modal features. We also train the

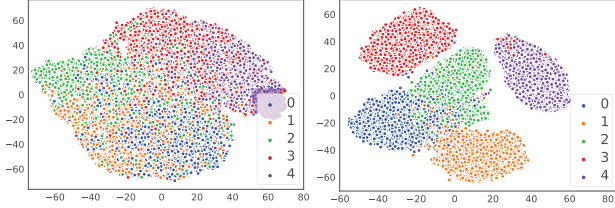


Figure 2: **Qualitative results on unknown classes of CIFAR-10.** Left: model trained on the labelled data only. Right: our model after end-to-end training on both labelled and unlabelled data. Color denotes the ground truth.

model with self-supervised contrastive learning and run k -means on the learned features (row 4). This is less effective than the baseline in row 3. To compare with the previous state-of-the-art method [16], which is designed for image category discovery, we implement its multi-modal counterpart (row 5) and also enhance it by replacing the original RotNet pretraining with our improved contrastive learning (row 6). It can be seen from row 7 that the contrastive learning introduced in our framework can further boost the performance of [16], and our approach consistently outperforms the k -means baseline and [16] by a large margin. For example, our method achieves 56.5%/50.0% ACC on Kinetics-400/VGG-Sound, while [16] only gives 31.2%/38.6% ACC. Interestingly, we found that [16] results in row 5-6 are less effective than the k -means baseline. We conjecture that ranking statistics in [16] only considers the most salient features when generating pseudo labels to train on unlabelled data. This might be sufficient for images but not for complex videos, e.g., different actions may share background or sub-actions, and ranking statistics may wrongly recognize them as positive pairs, leading to poor performance. In contrast, WTA considers holistic structure, and suffers much less from this.

Table 2: **Novel category discovery on video benchmarks.**

No	Method	audio	video	Kinetics-400	VGG-Sound
(1)	k -means [32]	✗	✓	40.9%	32.4%
(2)	k -means [32]	✓	✗	18.7%	41.7%
(3)	k -means [32]	✓	✓	41.1%	43.4%
(4)	CL [6] + k -means	✓	✓	34.7%	28.1%
(5)	RS [16] w/ RotNet	✓	✓	31.2%	38.6%
(6)	RS [16] w/ CL	✓	✓	33.5%	42.2%
(7)	Ours	✓	✓	56.5%	50.0%

We compare the features on videos in fig. 3. For visualisation purpose, we randomly choose unlabelled instances from 10 classes from Kinetics-400, we can see that our model can successfully separate novel classes into compact groups, while the novel categories are projected very close to each other for the baseline model trained with full supervision on the labelled data.

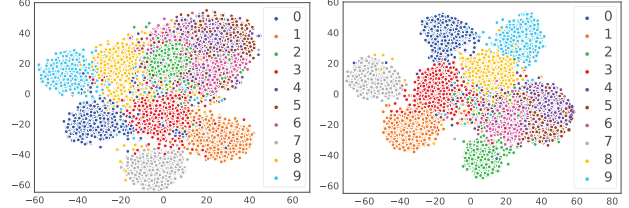


Figure 3: **Qualitative results on Kinetics-400.** Left: model trained on the labelled data only. Right: our model after end-to-end training on both labelled and unlabelled data. Color denotes the ground truth.

4.3. Ablation study

Component analysis. We validate the effectiveness of each component of our method during training on both image and video benchmarks. The results are reported in table 3. As can be seen, all components are important to our method. Removing the BCE loss causes the most performance drop. Namely, the ACC drops from 93.4% \rightarrow 32.2% on CIFAR-10, 76.4% \rightarrow 15.5% on CIFAR-100 and 56.5% \rightarrow 5.8% on Kinetics-400, suggesting that the WTA-hashing comparison can indeed generate reliable pairwise pseudo labels for the BCE loss. Removing the contrastive learning leads to significant performance drop (row 4 vs row 7). By comparing rows 5-7, we can see that category (NCE-C) or instance (NCE-I) discrimination alone is not good enough for the task of novel category discovery, and incorporating both is very effective.

Table 3: **Ablation study on image and video benchmarks.** MSE: consistency constraint; CE: cross entropy loss; BCE: binary cross entropy loss; NCE-I: NCE for instance discrimination; NCE-C: NCE for category discrimination.

No	MSE	CE	BCE	NCE-I	NCE-C	CIFAR-10	CIFAR-100	Kinetics-400
(1)	✗	✓	✓	✓	✓	91.3 \pm 1.4%	74.7 \pm 2.9%	50.9%
(2)	✗	✗	✓	✓	✓	77.4 \pm 5.4%	68.5 \pm 4.7%	24.5%
(3)	✓	✓	✗	✓	✓	32.2 \pm 1.8%	15.5 \pm 1.2%	5.8%
(4)	✓	✓	✓	✗	✗	89.8 \pm 0.9%	70.1 \pm 4.5%	47.6%
(5)	✓	✓	✓	✗	✗	91.3 \pm 0.8%	76.0 \pm 2.4%	52.8%
(6)	✓	✓	✓	✓	✗	91.6 \pm 1.2%	75.2 \pm 3.1%	51.7%
(7)	✓	✓	✓	✓	✓	93.4 \pm 0.6%	76.4 \pm 2.8%	56.5%

Table 4: **Multi-modal contrastive learning.** “(a, a)” denotes audio-modal contrastive; “(v, v)” denotes visual-modal contrastive; and “(a, v)” denotes cross-modal contrastive.

No	NCE-I	NCE-C	Kinetics-400
(1)	(a, a) + (v, v)	-	46.2%
(2)	(a, a) + (v, v)	(a, a) + (v, v)	51.4%
(3)	(a, a) + (v, v)	(a, v)	49.1%
(4)	(a, v)	-	51.7%
(5)	(a, v)	(a, a) + (v, v)	56.1%
(6)	(a, v)	(a, v)	56.5%

Multi-modal contrastive learning. As mentioned in section 3.1, for multi-modal data like videos, we can have multiple contrastive options. Namely, we can conduct within-

modal and cross-modal contrastive learning for each of NCE-I and NCE-C. We present the results of different options in table 4. By comparing rows 1 and 4, we can see that cross-modal is more effective than within-modal instance discrimination. It can be seen from rows 1-3 and rows 4-6 that NCE-C can effectively improve the representation for novel category discovery. Also, it is important to consistently use either within-modal or cross-modal discrimination for NCE-I and NCE-C (row 2 vs row 3; row 5 vs row 6). Overall, cross-modal contrastive learning for NCE-I and NCE-C (row 6) performs notably better than all other cases. We hypothesize this is because enforcing the cross-modal agreement can avoid the modal falling into the trivial solution, such as simply identifying the most salient visual or audio features. The modality discrepancy can effectively avoid such cases, thus improving the quality of learned representation.

4.4. Analysis on WTA

Hyperparameters of WTA. We obtain the two hyperparameters of WTA, i.e. threshold μ and window size k , by examining different values on the labelled data. To do so, we further split the labelled data into a smaller labelled subset and an unlabelled subset (i.e. pretending part of the labelled data to be unlabelled), and find μ and k that give the best results on the unlabelled subset. More specifically, for Kinetics-400, we split 350-class labelled set into a 300-class labelled subset and a 50-class unlabelled subset. For VGG-Sound, we split the 270-class labelled subset into 231-class labelled subset and 39-class unlabelled subset respectively. The number of unlabelled classes used here are the same as the unlabelled classes in the main experiments. Taking the empirical WTA window size $k = 4$ in [48], we experiment with different threshold μ and report the results in table 5. We find that the performance is generally stable for μ greater than 200. We set μ to 240 in our experiments. Note that the feature vectors have a dimension of 512. Given $\mu = 240$, we

Table 5: Performance of different WTA threshold.

Dataset	130	180	200	240	260	300
Kinetics-400	19.4	39.7	55.2	54.5	54.8	54.5
VGG-Sound	21.5	41.6	49.6	51.3	50.9	47.9

sweep different k and report the results in table 6. We find that the $k = 4$ and $k = 8$ perform comparably well, and they are both better than $k = 2$ and $k = 16$. Hence, we simply use the $k = 4$ in our experiments.

Table 6: Performance of different WTA window size.

Dataset	$k = 2$	$k = 4$	$k = 8$	$k = 16$
Kinetics-400	52.6	56.7	55.9	49.4
VGG-Sound	46.2	51.8	51.0	48.3

Comparing WTA with other alternatives. To better understand the role of WTA in our framework, we investigate other alternatives, namely, cosine similarity, ranking statistics, and nearest-neighbour, to transfer pseudo labels on-the-fly for both [16] and our framework. We report the results in table 7. It can be seen from the last row that simply replacing ranking statistics by WTA for [16] only shows marginal improvement, while our method significantly outperforms [16] in all cases (using WTA or any other alternative). WTA consistently performs better than ranking statistics. Intuitively, global ranking (ranking statistics) focuses on the most significant values in the feature space, while WTA considers multiple partial orders in random subsets of feature dimensions. Thus, WTA will consider values spread out the feature space when making the decision, avoiding the comparison to be dominated by high frequency noise or biased by small local regions, which is likely to be the case for global ranking. For example, as shown in [53], modern CNNs are likely to focus only on the (local) parts when recognizing objects (e.g., the beak of a bird), while sometimes the model could focus on the wrong parts, resulting completely wrong predictions. By considering multiple partial rankings, WTA can have a holistic view of the object, leading to more reliable comparison results in a unified framework learning.

Table 7: WTA vs other alternatives. We compare results of [16]/our method using different pairwise pseudo label generation methods.

Method	CIFAR-10	CIFAR-100	Kinetics-400
cosine	90.1%/92.1%	73.3%/75.4%	28.3%/35.4%
nearest-neighbour	90.2%/91.8%	69.7%/73.6%	22.7%/19.6%
ranking statistics [16]	90.4%/92.3%	73.2%/74.7%	31.2%/37.4%
WTA	90.5%/93.4%	73.4%/76.4%	38.5%/56.5%

5. Conclusion

We have presented a flexible end-to-end framework to tackle the challenging problem of novel category discovery. First, we extended the conventional contrastive learning to perform instance discrimination as well as category discrimination jointly by making full use of the labelled data and unlabelled data. Second, to successfully transfer knowledge from the labelled data to the unlabelled data, we employed the WTA hashing algorithm to generate pair-wise weak pseudo labels for training on unlabelled data, which is the key to automatically partition the unlabelled data into proper groups after training. Third, for multi-modal data, we investigated different ways of contrastive learning and we empirically found that cross-modal noise contrastive estimation performs consistently better than other options. Last, we thoroughly evaluated our approach on challenging image and video benchmarks and obtain superior results in all cases.

References

- [1] Jean-Baptiste Alayrac, Adrià Recasens, Rosalia Schneider, Relja Arandjelović, Jason Ramapuram, Jeffrey De Fauw, Lucas Smaira, Sander Dieleman, and Andrew Zisserman. Self-supervised multimodal versatile networks. *arXiv preprint arXiv:2006.16228*, 2020. 2
- [2] Humam Alwassel, Dhruv Mahajan, Lorenzo Torresani, Bernard Ghanem, and Du Tran. Self-supervised learning by cross-modal audio-video clustering. *arXiv preprint arXiv:1911.12667*, 2019. 2, 4
- [3] Yuki M. Asano, Mandela Patrick, Christian Rupprecht, and Andrea Vedaldi. Labelling unlabelled videos from scratch with multi-modal self-supervision. *arXiv preprint arXiv:2006.13662*, 2020. 2
- [4] Olivier Chapelle, Bernhard Scholkopf, and Alexander Zien. *Semi-Supervised Learning*. MIT Press, 2006. 2
- [5] Honglie Chen, Weidi Xie, Andrea Vedaldi, and Andrew Zisserman. Vgg5ound: A large-scale audio-visual dataset. In *ICASSP*, 2020. 5
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *ICML*, 2020. 1, 2, 6, 7
- [7] Bowen Cheng, Maxwell D Collins, Yukun Zhu, Ting Liu, Thomas S Huang, Hartwig Adam, and Liang-Chieh Chen. Panoptic-deeplab. *arXiv preprint arXiv:1910.04751*, 2019. 1
- [8] Dorin Comaniciu and Peter Meer. Mean shift: A robust approach toward feature space analysis. *IEEE TPAMI*, 1979. 2
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 1, 5
- [10] Kamran Ghasedi Dizaji, Amirhossein Herandi, Cheng Deng, Weidong Cai, and Heng Huang. Deep clustering via joint convolutional autoencoder embedding and relative entropy minimization. In *ICCV*, 2017. 2
- [11] Carl Doersch, Abhinav Gupta, and Alexei A. Efros. Unsupervised visual representation learning by context prediction. In *ICCV*, 2015. 2
- [12] Enrico Fini, Enver Sangineto, Stéphane Lathuilière, Zhun Zhong, Moin Nabi, and Elisa Ricci. A unified objective for novel class discovery. In *ICCV*, 2021. 3
- [13] Dario Fontanel, Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulò, Elisa Ricci, and Barbara Caputo. Boosting deep open world recognition by clustering. *arXiv preprint arXiv:2004.13849*, 2020. 1
- [14] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *ICLR*, 2018. 2
- [15] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 297–304, 2010. 2, 3
- [16] Kai Han, Sylvestre-Alvise Rebuffi, Sebastien Ehrhardt, Andrea Vedaldi, and Andrew Zisserman. Automatically discovering and learning new visual categories with ranking statistics. In *ICLR*, 2020. 1, 2, 5, 6, 7, 8
- [17] Kai Han, Sylvestre-Alvise Rebuffi, Sebastien Ehrhardt, Andrea Vedaldi, and Andrew Zisserman. Autonovel: Automatically discovering and learning novel visual categories. *IEEE TPAMI*, 2021. 2
- [18] Kai Han, Andrea Vedaldi, and Andrew Zisserman. Learning to discover novel visual categories via deep transfer clustering. In *ICCV*, 2019. 1, 2, 5, 6
- [19] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. *CVPR*, 2020. 1, 2
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 6
- [21] Yen-Chang Hsu, Zhaoyang Lv, and Zsolt Kira. Learning to cluster in order to transfer across domains and tasks. In *ICLR*, 2018. 2, 6
- [22] Yen-Chang Hsu, Zhaoyang Lv, Joel Schlosser, Phillip Odom, and Zsolt Kira. A probabilistic constrained clustering for transfer learning and image category discovery. In *CVPR Deep-Vision workshop*, 2018. 5
- [23] Yen-Chang Hsu, Zhaoyang Lv, Joel Schlosser, Phillip Odom, and Zsolt Kira. Multi-class classification without multi-class labels. In *ICLR*, 2019. 2, 6
- [24] Shuiwang Ji, Wei Xu, Ming Yang, and Kai Yu. 3d convolutional neural networks for human action recognition. *IEEE TPAMI*, 2012. 6
- [25] Xu Ji, João F Henriques, and Andrea Vedaldi. Invariant information clustering for unsupervised image classification and segmentation. In *ICCV*, pages 9865–9874, 2019. 6
- [26] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. 5
- [27] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *arXiv preprint arXiv:2004.11362*, 2020. 2
- [28] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. *Technical report*, 2009. 5
- [29] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 1955. 6
- [30] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, et al. The open images dataset v4. *International Journal of Computer Vision*, pages 1–26, 2020. 1
- [31] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. In *ICLR*, 2017. 5
- [32] James MacQueen. Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, 1967. 2, 6, 7
- [33] Pedro Morgado, Nuno Vasconcelos, and Ishan Misra. Audio-visual instance discrimination with cross-modal agreement. *arXiv preprint arXiv:2004.12943*, 2020. 2, 4

- [34] Avital Oliver, Augustus Odena, Colin Raffel, Ekin D. Cubuk, and Ian J. Goodfellow. Realistic evaluation of deep semi-supervised learning algorithms. In *NeurIPS*, 2018. 2
- [35] Daniel S Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D Cubuk, and Quoc V Le. Specaugment: A simple data augmentation method for automatic speech recognition. *arXiv preprint arXiv:1904.08779*, 2019. 6
- [36] Mandela Patrick, Yuki M Asano, Ruth Fong, João F Henriques, Geoffrey Zweig, and Andrea Vedaldi. Multi-modal self-supervision from generalized data transformations. *arXiv preprint arXiv:2003.04298*, 2020. 2, 6
- [37] Sylvestre-Alvise Rebuffi, Sebastien Ehrhardt, Kai Han, Andrea Vedaldi, and Andrew Zisserman. Lsd-c: Linearly separable deep clusters. *arXiv preprint arXiv:2006.10039*, 2020. 2
- [38] Sylvestre-Alvise Rebuffi, Sebastien Ehrhardt, Kai Han, Andrea Vedaldi, and Andrew Zisserman. Semi-supervised learning with scarce annotations. In *CVPR Deep Vision Workshop*, 2020. 2
- [39] Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the importance of initialization and momentum in deep learning. In *ICML*, 2013. 6
- [40] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *CVPR*, 2015. 6
- [41] Mingxing Tan, Ruoming Pang, and Quoc V Le. Efficientdet: Scalable and efficient object detection. In *CVPR*, 2020. 1
- [42] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NeurIPS*, 2017. 2, 5
- [43] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. *arXiv preprint arXiv:1906.05849*, 2019. 2
- [44] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*, 2018. 6
- [45] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 2
- [46] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, 2018. 6
- [47] Junyuan Xie, Ross Girshick, and Ali Farhadi. Unsupervised deep embedding for clustering analysis. In *ICML*, 2016. 2
- [48] Jay Yagnik, Dennis Strelow, David A. Ross, and Ruei sung Lin. The power of comparative reasoning. In *ICCV*, 2011. 1, 4, 5, 6, 8
- [49] Richard Zhang, Phillip Isola, and Alexei A Efros. Split-brain autoencoders: Unsupervised learning by cross-channel prediction. In *CVPR*, 2017. 2
- [50] Bingchen Zhao and Kai Han. Novel visual category discovery with dual ranking statistics and mutual knowledge distillation. *arXiv preprint arXiv:2107.03358*, 2021. 3
- [51] Zhun Zhong, Enrico Fini, Subhankar Roy, Zhiming Luo, Elisa Ricci, and Nicu Sebe. Neighborhood contrastive learning for novel class discovery. In *CVPR*, 2021. 3
- [52] Zhun Zhong, Linchao Zhu, Zhiming Luo, Shaozi Li, Yi Yang, and Nicu Sebe. Openmix: Reviving known knowledge for discovering novel visual categories in an open world. In *CVPR*, 2021. 3
- [53] Mohan Zhou, Yalong Bai, Wei Zhang, Tiejun Zhao, and Tao Mei. Look-into-object: Self-supervised structure modeling for object recognition. In *CVPR*, 2020. 5, 8
- [54] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019. 1