

This work was written as part of one of the author's official duties as an Employee of the United States Government and is therefore a work of the United States Government. In accordance with 17 U.S.C. 105, no copyright protection is available for such works under U.S. Law.

Public Domain Mark 1.0

<https://creativecommons.org/publicdomain/mark/1.0/>

Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback

Please support the ScholarWorks@UMBC repository by emailing scholarworks-group@umbc.edu and telling us what having access to this work means to you and why it's important to you. Thank you.

NEV-NCD: NEGATIVE LEARNING, ENTROPY, AND VARIANCE REGULARIZATION BASED NOVEL ACTION CATEGORIES DISCOVERY

Zahid Hasan^{*}, Masud Ahmed^{*}, Abu Zaher Md Faridee⁺, Sanjay Purushotham^{*}
Heesung Kwon[§], Hyungtae Lee[§], Nirmalya Roy^{*}

^{*}Department of Information Systems, University of Maryland, Baltimore County, USA

⁺Amazon, USA; [§]DEVCOM Army Research Laboratory (ARL), USA

ABSTRACT

Novel Categories Discovery (NCD) facilitates learning from a partially annotated label space and enables deep learning (DL) models to operate in an open-world setting by identifying and differentiating instances of novel classes based on the labeled data notions. One of the primary assumptions of NCD is that the novel label space is perfectly disjoint and can be equipartitioned, but it is rarely realized by most NCD approaches in practice. To better align with this assumption, we propose a novel single-stage joint optimization-based NCD method, **N**egative learning, **E**ntropy, and **V**ariance regularization NCD (**NEV-NCD**). We demonstrate the efficacy of NEV-NCD in previously unexplored NCD applications of video action recognition (VAR) with the public UCF101 dataset and a curated in-house partial action-space annotated multi-view video dataset. Further, we perform a thorough ablation study by varying the composition of final joint loss and associated hyper-parameters. During our experiments with UCF101 and multi-view action dataset, NEV-NCD achieves $\approx 83\%$ classification accuracy in test instances of labeled data. NEV-NCD achieves $\approx 70\%$ clustering accuracy over unlabeled data outperforming both naive baselines and state-of-the-art pseudo-labeling-based approaches by $\approx 40\%$ and $\approx 3.5\%$ over both datasets.

Index Terms— NCD, Negative Learning, Variance Loss.

1. INTRODUCTION

Deep learning (DL) based approaches form the backbone of modern video understanding, video action recognition (VAR), surveillance, and human-computer interactions. They learn to extract relevant spatio-temporal action features from videos and map them to predefined action classes. The traditional VAR models often assume the existence of exact labels (often laboriously annotated) for all action classes and fail to comprehend novel unlabeled action categories. *Novel Categories Discovery* (NCD) [1–3], an emerging field of DL, aims at more generalized model development under the partial class-space labeled data and alleviates the need for class-specific exact labels for all classes. NCD methods utilize the semantics knowledge of class notions from the labeled data to detect and structure unlabeled




Fig. 1: NCD objective aims to detect and structure the unlabeled novel categories based on the semantics of disjoint labeled data.

novel classes while retaining the classification performance (Figure 1). In this research, we study a novel NCD method in VAR to extract generalized spatial-temporal action features from videos based on labeled class notions to identify and cluster unlabeled actions.

The current NCD approaches rely on either a two-stage or a single-stage joint optimization [4] with various data-specific assumptions (simultaneous/sequential data availability, disjoint classes (mutually exclusive sets) between labeled and unlabeled data, knowledge about novel class quantity and their uniform distribution). These approaches optimize the classification loss of labeled data and cluster unlabeled data based on the class notions and semantics of the labeled data. The clustering constraints for unlabeled data are the primary key to NCD and can be enforced by pairwise similarity losses [5], pseudo-labeling [6], prototype learning [7], Sinkhorn-Knopp clustering [8], and self-training [9], using the uniformity property.

We observe that these NCD clustering objectives aim at achieving low entropy and equal variance among the unlabeled classes. However, these approaches fail to explicitly capitalize on the NCD’s assumption of unique, uniform, and disjoint labels for the unlabeled instances. Hence, we formulate the NCD’s disjoint properties as learning from complementary labels for unlabeled data and find the similarities of fair k -face dice statistics with the unique and uniform conditions of the unlabeled data. We propose an intuitive solution to achieve low-entropy, equal class variance aligning equipartition constraint for the novel classes in separate embedding spaces from the labeled classes taking inspiration from negative learning [10] and VICreg [11]. We hypothesize that negative learning (disjoint labels), explicit entropy (unique label), and variance (equipartition) regularization would cluster the unlabeled instances aligning with the semantics of the partially labeled data. Our contribution of

This research is supported by the NSF CAREER grant #1750936, REU Site grant #2050999 and U.S. Army grant #W911NF2120076. An extended version of this work is available at <https://arxiv.org/abs/2304.07354>.

incorporating this complementary information and k -face dice statistics into a novel NCD methodology is summarized below:

(1) We propose a novel joint optimization-based single-stage NCD method NEV-NCD (*Negative learning, Entropy, and Variance regularization NCD*) by integrating negative learning, entropy, and variance constraints. Further, we adapt supervised contrastive learning [12] to enhance the cluster discrimination.

(2) We validate NEV-NCD in the VAR by experimenting with public UCF101 [13] and our in-house multi-view video datasets. We thoroughly investigate the impacts of defined class notions, individual loss components, backbone networks, and viewpoints. NEV-NCD achieves 92%, 84% test accuracy on the labeled data and 82%, 70% average clustering accuracy on the unlabeled data for UCF101 and in-house dataset respectively. Our approaches outperform naive and single-stage NCD baselines in both datasets. We open-source our codes and data¹ for further research.

2. METHODOLOGY

2.1. Problem Formulation and Notations

The dataset consists of labeled data $\mathcal{D}^l = \{(x_i^l, y_i)\}_{i=1}^{N^l}$ where $y_i \in C_L = \{c_1, c_2, \dots, c_L\}$ from $L = |C_L|$ unique ordinary categories (instance belongs to the class) and unlabeled data $\mathcal{D}^u = \{(x_i^u)\}_{i=N^l+1}^{N^l+N^u}$ from $U = |C_U|$ latent novel classes $y_i \in C_U = \{c_{L+1}, c_{L+2}, \dots, c_{L+U}\}$ with non-overlapping categories $C_L \cap C_U = \emptyset$. Alternatively, we rewrite \mathcal{D}^u with multiple complementary labels (instance does not belong to complementary label classes) by $\{(x_i^u, \bar{y}_i)\}_{i=N^l+1}^{N^l+N^u}$, where $x_i^u \notin \bar{y}_i$ and $\bar{y}_i \in C_L$.

We denote parameterized 3D convolutional neural network (CNN) for video encoder f_θ , fully connected representation layer h_θ , and fully connected classification head layer g_θ . The g_θ consists of the concatenation of label head l_θ (0 to $L-1$ labeled categories) and unlabeled head u_θ (L to $L+U-1$ unlabeled categories) to represent the probabilities over the total dataset categories. The NCD objectives aim to develop DL models to classify the \mathcal{D}^l by comprehending the labeling semantics and clustering the \mathcal{D}^u counterpart based on their underlying latent class distribution.

2.2. NEV-NCD

Our proposed NEV-NCD method falls under the single-stage approach requiring the availability of both \mathcal{D}^l and \mathcal{D}^u concurrently. NEV-NCD jointly optimizes both the supervised and clustering objectives (Figure 2) as discussed in the following.

Ordinary Label Learning: The supervised training covers learning from both ordinary and multiple complementary labels. The ordinary labeled data \mathcal{D}^l are utilized to optimize the categorical cross-entropy loss 1 of $y_i \in C_L$ in the classification head. We denote network prediction $\hat{y}_i = \sigma_k(g_\theta(h_\theta(f_\theta(x_i))))$, where σ_k represents Softmax operation over the prediction containing $L+U$ neurons of g_θ .

$$\mathcal{L}_{ce} = - \mathbb{E}_{(x_i, y_i) \in \mathcal{D}^l} \sum_{j=0}^{L-1} y_{ij} \log \hat{y}_{ij} \quad (1)$$

¹ Codes and Dataset




Fig. 2: Overview of our NCD approach based on negative learning and variance constraint to identify novel clusters

y_{ij} represents a one-hot encoding vector of the target y_i .

Negative Learning: We utilize the disjoint condition between \mathcal{D}^u and \mathcal{D}^l latent classes to formulate the problem as learning from multiple complementary labels [14] for the \mathcal{D}^u and propose to incorporate negative-learning-based solution [10]. The negative learning utilizes the complementary labels to reduce their confidence on l_θ for the corresponding instances contrary to the ordinary labeled supervised learning. In our approach, we enforce zero (low-confidence) for the complementary labels by applying the loss defined in Equation 2.

$$\mathcal{L}_{nl} = - \mathbb{E}_{x_i \in \mathcal{D}^u} \left[\sum_{j=0}^{L-1} \log(1 - \hat{y}_{ij}) \right] \quad (2)$$

Entropy Regularization: The final classification layer SoftMax operation and negative learning loss function impose higher confidence in the unlabeled classification head u_θ for $x_i \in \mathcal{D}^u$. We employ the uniquely-labeled data condition and apply entropy constraint loss (Equation 3) to increase the confidence of the unlabeled head as the Dirac-delta distribution minimizes the entropy [15]. The entropy regularization further ensures gradient flow through the weight matrix of u_θ nodes.

$$\mathcal{L}_H = - \mathbb{E}_{x_i \in U} \left[\sum_{j=0}^{L+U-1} \hat{y}_{ij} \log \hat{y}_{ij} \right] \quad (3)$$

Contrastive and Consistency Loss: We adopt the Sup-Con [12] framework with NCD to enhance the representation quality by modifying the data sampling using the underlying data conditions. We optimize probabilistic InfoNCE contrastive objectives to learn both categorical and instance discrimination in the representation layer output denoted by $z_i = h_\theta(f_\theta(x_i))$.

Category Discrimination: We utilize the known categories of the \mathcal{D}^l and disjoint criterion to perform categorical contrastive learning. Here, we sample query and positive instances from $c_j \in C_L$ and take negative instances distributed from remaining classes and unlabeled data $C_L \setminus \{c_j\} \cup C_U$ to minimize the contrastive objective of 4 where $s_{i,j} = \text{cosine}(z_i, z_j)$, τ denotes cosine similarity between embedding and temperature.

$$\mathcal{L}_{cl} = \mathbb{E}_{c_l \sim C_L} \mathbb{E}_{x_i, x_j | c_l} - \log \frac{e^{(s_{i,j}/\tau)}}{e^{(s_{i,j}/\tau)} + \sum_{n=0}^N e^{(s_{i,kn}/\tau)}} \quad (4)$$

Instance Discrimination: Although there lacks information about the intra-class similarity between unlabeled instances, we have their complementary labels. For the query of unlabeled instances, we sample negatives from \mathcal{D}^l to avoid sampling

false negative keys and minimize 5. We also enforce consistency [16] between the representation of the instance x_i and their augmented version x'_i by equation 6. Here, we rely on variance regularization 2.2 instead of contrastive learning to discriminate between unlabeled latent classes.

$$\mathcal{L}_{cl} = \mathbb{E}_{\substack{x_i \in \mathcal{D}^u \\ x_{k,n} | c_l \in C_L}} -\log \frac{e^{(s_{i,i'}/\tau)}}{e^{(s_{i,i'}/\tau)} + \sum_{n=0}^N e^{(s_{i,kn}/\tau)}} \quad (5)$$

$$\mathcal{L}_{mse} = \mathbb{E}_{x_i \in \mathcal{D}^u} \|z - z'\|_2^2 \quad (6)$$

Variance Regularization: Uniform random sampling from a balanced data of k classes is equivalent to roll a fair k -face dice with probabilities, $P = \frac{1}{k}$, of each class. We can define a binary random variable, $X_j \in \{0, 1\}$, representing the events of j -th face, $j \in 1, 2, \dots, k$, outcome success and failure with probabilities, $P(X_j = 1) = \frac{1}{k}$, mean, $\mathbb{E}(X_j) = \frac{1}{k}$, and variance $\frac{k-1}{k^2}$. We explicitly apply this constraint as regularization 7 for the empirical distribution of the sharpened (sharpening hyper-parameter Sr) unlabeled prediction head $\tilde{y}_i = \sigma_k(\hat{y}_i/Sr)$ for relatively large batch size.

$$\mathcal{L}_{var} = \mathbb{E}_{x_i \in \mathcal{D}^u} \left| \text{var}(\tilde{y}_i) - \frac{k-1}{k^2} \right|^2 \quad (7)$$

Joint Optimization: We jointly optimize the learning objective 8 of a weighted (λ) sum of the discussed losses and regularization by utilizing both \mathcal{D}^l and \mathcal{D}^u in a single stage with supervised contrastive architecture framework.

$$\mathcal{L} = \lambda_{ce}\mathcal{L}_{ce} + \lambda_{cl}\mathcal{L}_{cl} + \lambda_{nl}\mathcal{L}_{nl} + \lambda_H\mathcal{L}_H + \lambda_{var}\mathcal{L}_{var} \quad (8)$$

3. EXPERIMENTS

We demonstrate the NEV-NCD effectiveness by experimenting with two video datasets and performing ablation studies.

3.1. Dataset

UCF101 [13]: We have preprocessed the UCF101 data with randomly selected 91 labeled actions and the remaining 10 as unlabeled actions. We balance the distribution of the instances from these unlabeled classes to match the NCD assumption.

In-house dataset: our collected multiview dataset contains three synchronous views from three cameras (action camera, handheld smartphone, and low-altitude drone) that captured the action from three realistic positions with varying camera settings, positions, and angles. It contains class-balanced ten regular micro-actions with both static (i.e. body poses: sitting, standing, lying with face up and down) and dynamic patterns (i.e. temporal patterns: walking, push up, waving hand, leg exercise, object carrying, and pick/drop). In total, we collect ≈ 14 hours of video from 11 volunteers under varying realistic lighting, environments, and backgrounds. We perform controlled experiments to study various NEV-NCD components, view-invariant, and action selection using our in-house dataset due to its scalable label space and available viewpoint meta-data.

3.2. Implementation Details

Network Architecture: We experimented with three different CNN architectures (SlowFast, R2+1D [17], and P3D [18]) that

have demonstrated their capabilities in VAR as encoder design choices. We experimented with both random initialization and weight transfer from the kinetics dataset pre-trained model weights available in PyTorch Modelzoo and remove their classification head. We placed a 1000-dimensional embedding layer with ReLU activation and classification head with SoftMax activation on top of encoder features.

Losses: Inspired by [19], we applied adaptive weights for the contrastive learning, negative learning, variance loss, and cross-entropy loss components with the epoch number, n_{ep} , throughout the learning phase as shown in equation 9 and 10. In our experimentations, We fixed entropy loss weights ($\lambda_H = 1.0$) and other hyper-parameters ($\tau = 0.05$, $Sr = 0.1$). To calculate the empirical distribution variance, we traded off between computational cost and the law of large numbers and chose the batch size of $10 \times U$ instances.

$$\lambda_{cl} = \lambda_{nl} = \lambda_{var} = 0.2 + 0.5 \times n_{ep} \quad (9)$$

$$\lambda_{ce} = \max(0, 1 - 0.01 \times n_{ep}) + 0.5 \quad (10)$$

3.3. Baseline Comparison and Ablation Studies

We compare our methodology with three naive approaches – (i) the supervised training using only the labeled parts, (ii) instance discrimination-based contrastive learning utilizing a complete dataset, and (iii) the SupCon approach combining both labeled and unlabeled data. Additionally, we compare the performance of NEV-NCD with the most relevant single-stage joint optimization-based state-of-the-art NCD method UNO [8] with single head Sinkhorn-Knopp clustering (replacing the swapped prediction with self-labeling and retraining [9]) by adapting it for our VAR scenario.

We run numerous experiments to study the impacts of different design choices toward better understanding the proposed NCD VAR system. Firstly, we investigate the impacts and the relative importance of each loss component of our proposed joint loss. We also vary the hyper-parameters of the loss components (λ_{nl} , λ_{nl} , negative sample number) to study their robustness. Secondly, we investigate the impacts of choosing different 3D CNN architectures (i.e. P3D, R2+1D, SlowFast) as the parameterized feature encoders for our VAR model. Finally, we inspect the effect of multiview properties by varying members of C_L, C_U , and multiple viewpoint variations for the actions using our in-house dataset.

4. RESULTS

We evaluated NEV-NCD in left-out labeled and unlabeled test clips by both quantitative and qualitative metrics for both datasets. To scrutinize the quality (i.e. separability of classes) of the learned embeddings, we employed low-dimensional t-SNE visualization. We quantified the performance over the labeled and unlabeled data by deploying standard top-1 accuracy (ACR) and average clustering accuracy (ACC) (equation 11), respectively.

$$ACC = \max_{p \in Pe} \frac{1}{N} \sum_{i=1}^N \mathbb{1}\{y_i = p(\hat{y}_i)\} \quad (11)$$

Here $y_i \in C_U$ and \hat{y}_i represent the ground-truth and cluster prediction of a sample $x_i \in \mathcal{D}^u$, respectively. Pe denotes the set

of all permutations and can be efficiently computed $\forall x_i$ [20].

Table 1: Summarized Result on UCF101 Dataset (90 labeled, R2+1d)

Methods	NEV-NCD	UNO	SpCn	Sup
ACR %	92.4	92.1	93.2	92.6
ACC %	82.7	79.8	24.1	16.4

Table 2: Overall Result on our in-house video dataset

Methods	Loss Components			Labeled Classes	Metrics	
	$\mathcal{L}_{ce}, \mathcal{L}_{cl}$	SK	$\mathcal{L}_{nl}, \mathcal{L}_{var}, \mathcal{L}_H$		ACR (%)	ACC (%)
Ours (R2+1d)	Y	N	Y	Uniform	82.9	67.3
				Dynamics	81.8	64.2
				Static	87.1	53.0
Ours (SlowFast)	Y	N	Y	Uniform	84.1	70.5
				Dynamics	82.7	68.2
				Static	87.8	53.2
Ours (P3D)	Y	N	Y	Uniform	76.5	62.3
				Dynamics	72.9	59.3
				Static	78.6	45.7
Sup	\mathcal{L}_{ce}	N	N	Uniform	84.6	12.5
UNO [8]	\mathcal{L}_{ce}	Y	N	Uniform	83.4	62.1
SpCn [12]	Y	N	N	Uniform	85.2	12.8
CL [21]	Fine tuned			Uniform	80.6	13.4
Ablation Studies	Y	N	\mathcal{L}_{nl}	Uniform	83.9	25.3
	Y	N	\mathcal{L}_{var}	Uniform	85.1	15.2
	Y	N	\mathcal{L}_H	Uniform	84.8	25.6
	Y	N	$\mathcal{L}_{nl}, \mathcal{L}_H$	Uniform	83.9	30.2
	Y	N	$\mathcal{L}_{nl}, \mathcal{L}_{var}$	Uniform	83.7	60.5

We report quantitative results over the left-out test data of the UCF101 and the in-house dataset in Table 1 and 2 respectively. In both cases, NEV-NCD outperforms others in clustering the \mathcal{D}^u (high ACC) while classifying the \mathcal{D}^l (high ACR).

The naive baselines fail to cluster the \mathcal{D}^u separately (low ACC) from the labeled data embedding without external constraints. NEV-NCD outperforms the comparable UNO (self-labeling via Sinkhorn-Knopp (SK) clustering and retraining on pseudo-label) in the ACC metrics without retraining. Simultaneously, we analyze the 2D t-SNE plot for the embedding layer output. As depicted in figure 3, we observe that NEV-NCD learns distinguished features for both labeled and unlabeled instances.

The NCD performance drops significantly in the absence of \mathcal{L}_{var} and \mathcal{L}_{nl} as reported in Table 2. The \mathcal{L}_{nl} enforces a boundary between labeled and unlabeled classes by pushing away the unlabeled embedding from the labeled embedding. However, the \mathcal{L}_{nl} results in collapsing unlabeled classes into a single cluster without additional constraints. The \mathcal{L}_{var} provides the equipartition cluster constraint and avoids trivial collapses for the embedding of unlabeled instances. The \mathcal{L}_H loss increase inter-class separation by enforcing unique decision for the unlabeled classification heads. Additionally, we study and report the impact of relative weights of the loss components in Table 3.

We also observe the impact of labeled class notions in unlabeled data clustering. In our VAR experiments, labeled class sets with only static categories (sitting, standing, etc.) reduce the performance of clustering of dynamic actions (walking, push-ups, etc.). Besides, we find that NCD methods learn the invariant representations of actions from different viewpoints of unlabeled data by providing multiview labeled instances. Finally, we notice that the strong function class improves the overall performance in both labeled and unlabeled data (Table 2).

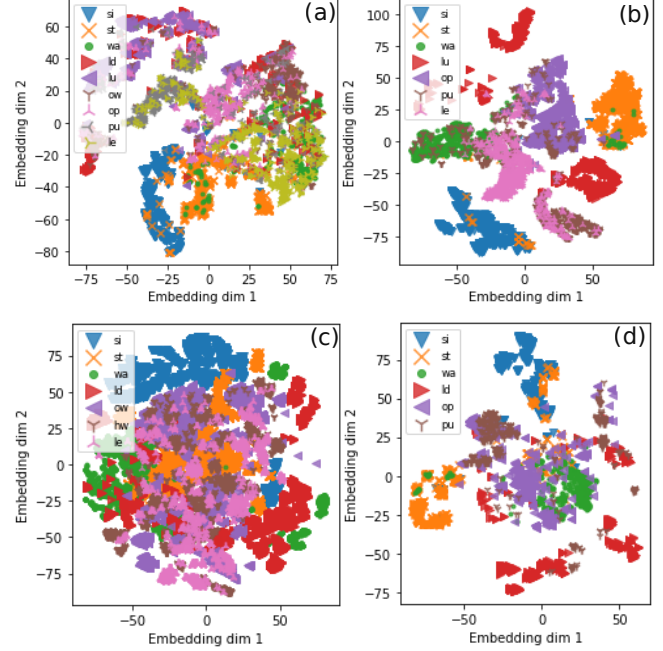


Fig. 3: 2D t-SNE plot for the representation of the embedding layer (a) UNO (b) NEV-NCD (c) Supervised Baseline (d) Negative Learning

Table 3: Ablation Studies using in-house dataset for different constant λ_{nl} and λ_{var} with SlowFast model (fixed other parameters).

Weights		Metrics		Comment
λ_{nl}	λ_{var}	ACR (%)	ACC (%)	
0.1 ~ 1	1	84	65 ~ 70	Slower convergence for smaller values
1	0.1 ~ 1	84	64 ~ 70	

5. DISCUSSION AND FUTURE WORKS

Our current method leverages three typical NCD data assumptions: (i) prior knowledge about the number of latent classes for the unlabeled data (unknown target cluster quantity), (ii) disjoint set condition between the C_L and C_U (negative learning), and (iii) the uniform distribution among unlabeled classes (variance regularization). Moreover, the empirical variance estimation requires a large batch size relative to U randomly sampled from unlabeled data as per the strong law of large numbers.

In the future, we aim to challenge the NCD assumptions and investigate the minimal constraint required for the unlabeled data to solve the NCD clustering. We plan to extend our works towards generalized categories discovery [22] and explore over-clustering, self-training, and distribution aligning solutions. We also plan to integrate and investigate the impact of attention-based learners (vision transformers) in NCD VAR applications.

6. CONCLUSION

In this work, we introduced and studied the negative learning, entropy, and variance regularization in NCD research to propose a novel single-stage NCD method NEV-NCD removing additional pseudo-label generation and retraining steps. We validated NEV-NCD by performing extensive experiments with the VAR domain using 3D CNN-based networks. The qualitative and quantitative results demonstrate the impacts of NEV-NCD components and compatibility with state-of-the-art NCD works.

7. REFERENCES

- [1] Kai Han, Andrea Vedaldi, and Andrew Zisserman, "Learning to discover novel visual categories via deep transfer clustering," in *ICCV*, 2019, pp. 8401–8409.
- [2] Kai Han, Sylvestre-Alvise Rebuffi, Sebastien Ehrhardt, Andrea Vedaldi, and Andrew Zisserman, "Autonovel: Automatically discovering and learning novel visual categories," *IEEE TPAMI*, 2021.
- [3] Kai Han, Sylvestre-Alvise Rebuffi, Sebastien Ehrhardt, Andrea Vedaldi, and Andrew Zisserman, "Automatically discovering and learning new visual categories with ranking statistics," *arXiv:2002.05714*, 2020.
- [4] KJ Joseph, Sujoy Paul, Gaurav Aggarwal, Soma Biswas, Piyush Rai, Kai Han, and Vineeth N Balasubramanian, "Spacing loss for discovering novel categories," in *CVPR*, 2022, pp. 3761–3766.
- [5] Yen-Chang Hsu, Zhaoyang Lv, and Zsolt Kira, "Learning to cluster in order to transfer across domains and tasks," *arXiv:1711.10125*, 2017.
- [6] Xuhui Jia, Kai Han, Yukun Zhu, and Bradley Green, "Joint representation learning and novel category discovery on single-and multi-modal data," in *ICCV*, 2021, pp. 610–619.
- [7] Lu Zhang, Lu Qi, Xu Yang, Hong Qiao, Ming-Hsuan Yang, and Zhiyong Liu, "Automatically discovering novel visual categories with self-supervised prototype learning," *arXiv:2208.00979*, 2022.
- [8] Enrico Fini, Enver Sangineto, Stéphane Lathuilière, Zhun Zhong, Moin Nabi, and Elisa Ricci, "A unified objective for novel class discovery," in *ICCV*, 2021, pp. 9284–9292.
- [9] Yuki Markus Asano, Christian Rupprecht, and Andrea Vedaldi, "Self-labelling via simultaneous clustering and representation learning," *arXiv:1911.05371*, 2019.
- [10] Youngdong Kim, Junho Yim, Juseung Yun, and Junmo Kim, "Nlnl: Negative learning for noisy labels," in *ICCV*, 2019, pp. 101–110.
- [11] Adrien Bardes, Jean Ponce, and Yann LeCun, "Vicreg: Variance-invariance-covariance regularization for self-supervised learning," *arXiv:2105.04906*, 2021.
- [12] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan, "Supervised contrastive learning," *NeurIPS*, vol. 33, pp. 18661–18673, 2020.
- [13] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah, "Ucf101: A dataset of 101 human actions classes from videos in the wild," *arXiv preprint arXiv:1212.0402*, 2012.
- [14] Lei Feng, Takuo Kaneko, Bo Han, Gang Niu, Bo An, and Masashi Sugiyama, "Learning with multiple complementary labels," in *ICML*. PMLR, 2020, pp. 3072–3081.
- [15] Yves Grandvalet and Yoshua Bengio, "Semi-supervised learning by entropy minimization," *NeurIPS*, vol. 17, 2004.
- [16] Muli Yang, Yuehua Zhu, Jiaping Yu, Aming Wu, and Cheng Deng, "Divide and conquer: Compositional experts for generalized novel class discovery," in *CVPR*, 2022, pp. 14268–14277.
- [17] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri, "Learning spatiotemporal features with 3d convolutional networks," in *ICCV*, 2015, pp. 4489–4497.
- [18] Zhaofan Qiu, Ting Yao, and Tao Mei, "Learning spatio-temporal representation with pseudo-3d residual networks," in *ICCV*, 2017, pp. 5533–5541.
- [19] Yu Liu and Tinne Tuytelaars, "Residual tuning: Toward novel category discovery without labels," *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- [20] Harold W Kuhn, "The hungarian method for the assignment problem," *Naval research logistics quarterly*, vol. 2, no. 1-2, pp. 83–97, 1955.
- [21] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton, "A simple framework for contrastive learning of visual representations," in *ICML*. PMLR, 2020, pp. 1597–1607.
- [22] Sagar Vaze, Kai Han, Andrea Vedaldi, and Andrew Zisserman, "Generalized category discovery," in *CVPR*, 2022, pp. 7492–7501.