Exploring Structural Degradation in Dense Representations for Self-supervised Learning

Siran Dai 1,2 Qianqian Xu 3,4 * Peisong Wen 5 Yang Liu 5 Qingming Huang 5,3*

¹ Institute of Information Engineering, Chinese Academy of Sciences
 ² School of Cyber Security, University of Chinese Academy of Sciences
 ³ State Key Laboratory of AI Safety, Institute of Computing Technology, CAS
 ⁴ Peng Cheng Laboratory

⁵ School of Computer Science and Tech., University of Chinese Academy of Sciences daisiran@iie.ac.cn xuqianqian@ict.ac.cn

{wenpeisong, qmhuang}@ucas.ac.cn liuyang232@mails.ucas.ac.cn

Abstract

In this work, we observe a counterintuitive phenomenon in self-supervised learning (SSL): longer training may impair the performance of dense prediction tasks (e.g., semantic segmentation). We refer to this phenomenon as Self-supervised Dense Degradation (SDD) and demonstrate its consistent presence across sixteen state-of-the-art SSL methods with various losses, architectures, and datasets. When the model performs suboptimally on dense tasks at the end of training, measuring the performance during training becomes essential. However, evaluating dense performance effectively without annotations remains an open challenge. To tackle this issue, we introduce a Dense representation Structure Estimator (DSE), composed of a class-relevance measure and an effective dimensionality measure. The proposed DSE is both theoretically grounded and empirically validated to be closely correlated with the downstream performance. Based on this metric, we introduce a straightforward yet effective model selection strategy and a DSE-based regularization method. Experiments on sixteen SSL methods across four benchmarks confirm that model selection improves mIoU by 3.0% on average with negligible computational cost. Additionally, DSE regularization consistently mitigates the effects of dense degradation. Code is available at https://github.com/EldercatSAM/SSL-Degradation.

1 Introduction

Self-Supervised Learning (SSL) has greatly benefited from advancements in training algorithms, larger datasets and models, and extended training periods, leading to significant success in image-level representation learning [10, 16, 3, 29]. However, dense (patch or pixel-level) representation learning remains challenging with only slight improvements [70, 5].

While training models for extended periods to extract high-quality representations has been a common practice in SSL, we identify and study a counterintuitive phenomenon, named Self-supervised Dense Degradation (SDD), which helps explain the challenges in self-supervised dense representation learning. As illustrated in Fig. 1, although the training loss converges and classification performance steadily improves, the dense performance declines at the later stages of training. Consequently, the final checkpoint exhibits a significant performance gap compared to the best performance observed during training.

^{*}Corresponding authors.

Extensive experiments on sixteen state-of-theart methods across four benchmarks confirm that the SDD phenomenon consistently appears across diverse training approaches and evaluation protocols. More importantly, it persists even when training and evaluation are conducted on the same dataset. This demonstrates that SDD highlights the performance inconsistency between different tasks rather than overfitting to the data distribution, introducing a new challenge to the SSL community.

Because of SDD, the common practice of training until convergence results in suboptimal dense performance. Finding a metric to predict downstream dense performance thus becomes essential for both understanding the cause of

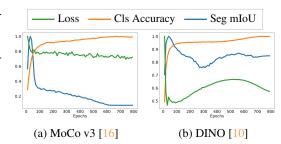


Figure 1: The Self-supervised Dense Degradation (SDD) phenomenon. While the training loss converges and classification performance gradually increases, the segmentation performance declines and the final checkpoint is not optimal. The metrics are normalized to 1 for a better illustration.

SDD and reducing its negative impact. Typically, models are evaluated on a labeled validation set; however, the cost of evaluating even a single checkpoint can exceed that of one full epoch of pretraining, making this impractical. Moreover, SSL typically lacks access to downstream data or labels. Although previous studies have tried estimating SSL downstream performance in an unsupervised manner [2, 24, 69, 68, 36, 60], these metrics mainly focus on image-level tasks and are negatively correlated with dense-level performance in our experiments.

To guide the development of a better metric, we first provide a theoretical analysis based on error rate decomposition. We show that the downstream error rate is small if 1) intra-class representation radius is smaller than inter-class distance, and 2) effective dimensionality of representations is large.

Based on this analysis, we propose a Dense Representation Structure Estimator (DSE), which consists of measures for class-separability and effective dimensionality, directly corresponding to our theoretical findings. Using the DSE metric, we propose two strategies to counteract SDD. In the off-the-shelf setting, we suggest a simple yet effective method of selecting the checkpoint with the highest local DSE. In the online training setting, we integrate the DSE metric as a regularizer.

Empirical evaluations demonstrate that the proposed DSE accurately predicts downstream performance, significantly outperforming existing metrics. By visualizing components of the DSE, we explain that the cause of the SDD phenomenon is either a loss of class separability or dimensional collapse. For example, in Fig. 1, MoCo v3's performance drop is due to dense dimensional collapse, while DINO's degradation arises from decreased class separability. Moreover, our proposed checkpoint selection approach improves the mIoU by 3.0% on average, and incorporating DSE into training further improves both DSE scores and downstream dense performance, effectively eliminating SDD's negative effects.

We summarize our contributions as follows:

- We identify the SDD phenomenon in SSL, revealing an inconsistency between image-level and dense-level performance, which is prevalent in state-of-the-art methods and negatively impacts dense tasks.
- We introduce the Dense Representation Structure Estimator (DSE), which accurately and efficiently predicts downstream dense performance without relying on downstream data.
- Using the DSE metric, we propose model selection and regularization strategies that effectively reduce the negative impact of the SDD phenomenon.
- Extensive experiments on sixteen leading SSL methods across four benchmarks demonstrate the precision of DSE and the effectiveness of the proposed approaches in addressing the SDD phenomenon.

2 Related Work

2.1 Self-supervised Learning

Contrastive Learning. Contrastive self-supervised methods have shown significant progress in recent years [11, 73, 28, 61, 62]. These methods typically construct positive examples (different augmented views of an image) and negative examples (samples from other images), then use contrastive loss [49] to train models to differentiate between them [12, 15, 78, 4, 32, 27, 30].

Non-contrastive Learning. Recent advances in self-supervised learning avoid explicit negative examples. These methods often employ siamese networks and aim to achieve cross-view consistency by aligning representations between teacher and student networks [26, 10, 83, 50, 14, 9].

Masked Image Modeling. Masked image modeling can be framed as a generative task. Models are trained to reconstruct original images from masked inputs [29, 74, 3, 13]. Recent works [83, 50] achieved great success by combining latent-space reconstruction with self-distillation.

Dense Representation Learning. Another line of work focuses on learning dense representations. Using techniques such as dense alignment [86, 70, 31, 71], clustering [84, 58, 40, 59, 5], and reconstruction [83, 50], these methods achieve strong results on dense prediction tasks. However, despite optimizing dense representations directly, performance degradation is still observedduring training.

2.2 Unsupervised Transferability Estimation

Recently, α -REQ [2] uses the parameter of power-law distributions of covariance matrix singular values as a metric. RankMe [24] employs the effective rank [53] to estimate the transferability performance, and Lidar [60] further improves it by introducing linear discriminative analysis. Other approaches analyze feature activation statistics [36], coding rate reduction [77] or model memorization effects [69, 68]. Despite progress in image-level tasks, evaluating dense representations remains an open challenge. Due to space limitation, more related works are discussed in Appendix B.

3 The Self-supervised Dense Degradation Phenomenon

In self-supervised learning, models are typically trained for long periods. It is widely recognized that downstream classification performance generally improves as training loss converges [10, 14, 16]. However, we observe that dense performance actually degrades during pretraining, with the performance at the final checkpoint being significantly worse than that of the best model. This observation contradicts previous intuitions. We term this phenomenon **Self-supervised Dense Degradation (SDD)**, and empirical and theoretical analyze this phenomenon in this section.

3.1 Empirical Observations of the SDD Phenomenon

SDD is a General and Harmful Phenomenon. To investigate whether SDD occurs broadly, we conduct extensive experiments. The main findings are summarized in Fig. 2. These experiments confirm that SDD occurs consistently across 1) Various Pre-training Approaches: SDD exists in sixteen state-of-the-art methods across different types of training loss, model architecture and optimization strategies, and 2) Various Evaluation Protocols: SDD is evident in both linear probing and transfer learning scenarios (where the backbone is not frozen), spanning diverse downstream datasets, evaluation hyperparameters, and tasks. Detailed results can be found in Appendix E. SDD is Not Caused by Overfitting the Training Data. We further investigate whether SDD stems from memorizing the training data. To test this, we train and evaluate DINO [10] on the same COCO dataset. The trend of dense performance mirrors that observed for DINO trained on ImageNet, with the final checkpoint experiencing a significant degradation of 4.0% in mIoU. This result indicates that SDD is not due to overfitting. Further details are presented in Appendix E.3.

Since SDD occurs broadly and is not related to dataset overfitting, identifying a suitable metric to predict downstream performance would be valuable for understanding and mitigating this degradation. However, defining such a metric is challenging. In the following subsection, we seek theoretical insights to help develop such a performance measure.

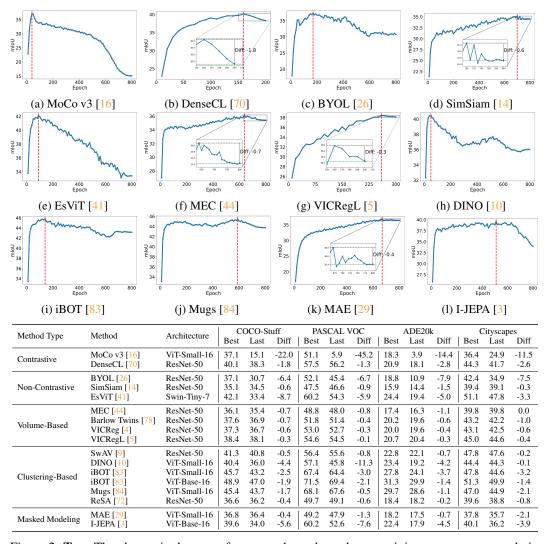


Figure 2: **Top:** The change in dense performance throughout the pretraining process, assessed via linear segmentation on the COCO-Stuff dataset. Performance degradation is consistently observed across all methods. **Bottom:** A performance gap between the best and the last models is present across all datasets and methods.

3.2 Theoretical Analysis of the SDD Phenomenon

Our goal is to establish a theoretically grounded metric for downstream performance. To achieve this, we analyze the downstream performance and identify two crucial factors influencing it: 1) class separability, quantified by the difference between inter-class distance and intra-class radius (formally presented in Thm. 2), and 2) the dimensionality of the representations (Cor. 5). With these insights, readers interested primarily in methodology and experiments may skip the remainder of this section without any loss of continuity.

3.2.1 Problem Formulation

Given that SDD appears across various methods, analyzing their training processes within a single unified framework is challenging. Therefore, we mainly focus on the linear probing approach, which aims to train a classifier using fixed (dense) representations. Specifically, given a downstream dataset $\mathcal{D} = \{X_i\}_{i=1}^{\bar{N}}$, consisting of \bar{N} images, a fixed encoder $f_{\theta}: \mathcal{X} \to \mathbb{R}^{N \times d}$ produces N dense representations $\{z_i\}_{i=1}^{N}$ for each image, where d is the dimension of the representations. To simplify, we formulate dense linear probing as a classification problem, where each representation

 z_i is assigned to one of K latent classes, represented as $y(z_i)$. The aim of linear probing is to train a classifier G(z) (for example, a linear head) that accurately maps each z to its correct latent class. Following the analysis in [34], we choose a simple Nearest Neighbor (NN) classifier:

$$G(\boldsymbol{z}) = \arg\min_{k \in [K]} ||\boldsymbol{z} - \boldsymbol{\mu}_k||,$$

where $\mu_k = \mathbb{E}_{z:y(z)=k}[z]$ denotes the center of the representations for the k-th class. The error rate of the fixed encoder f_{θ} on the downstream dataset is given by:

$$\operatorname{Err}_{\mathcal{D}}(f_{\boldsymbol{\theta}}) = \mathbb{E}_{\boldsymbol{x} \in \mathcal{D}} \left[\mathbb{P}_{\boldsymbol{z} \in f_{\boldsymbol{\theta}}(\boldsymbol{x})} \left[y(\boldsymbol{z}) \neq G(\boldsymbol{z}) \right] \right].$$

Since the NN classifier can be viewed as a special case of any linear classifier, its error rate naturally serves as an upper bound for all classifiers.

3.2.2 Decomposing the Downstream Error Rate

Next, we decompose downstream performance and identify the factors that influence downstream accuracy. For the NN classifier G, a representation can be correctly classified if the following condition holds:

$$\underbrace{||z - \mu_{y(z)}||}_{Intra-class \ distance} - \underbrace{\min_{k \in [K] \setminus y(z)} ||z - \mu_k||}_{Inter-class \ distance} \le 0. \tag{1}$$

Inspired by this relationship, downstream performance can be expressed in terms of intra and interclass distances. However, directly measuring these distances is not feasible. The main challenge arises from the fact that, in a self-supervised setting, class labels y(z) are unavailable. While techniques such as k-means clustering can be utilized to generate pseudo labels, we find that instance-wise distance measures still face a critical issue, resulting in meaningless predictions:

Proposition 1. The instance-wise intra-class distance is always smaller than the inter-class distance when using k-means pseudo-labels. Thus, the estimated accuracy is always 1, regardless of the actual situation.

Proof can be found in the Appendix A. The fundamental reason for this issue lies in the fact that the instance-wise distance measure tends to underestimate the intra-class distance for those examples near the decision boundary. We present an illustration of this issue in Fig. 3. To address this issue, we replace the instance-wise measure with a class-wise radius. As a result, we reformulate the condition that z could be correctly classified as:

$$\underbrace{\mathcal{R}_{y(z)}}_{\textit{Intra-class radius}} - \underbrace{\min_{k \in [K] \setminus y(z)} ||z - \mu_k||}_{\textit{Inter-class distance}} \le 0. \quad (2)$$

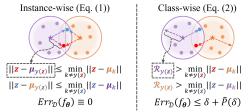


Figure 3: **Left:** Instance-wise condition (Eq. 1) predicts all examples in the intersection area as **correctly** classified, leading to an inaccurate error rate estimation. **Right:** Our class-wise condition (Eq. 2). Examples in the intersection area are accurately predicted as **misclassified**.

Building on this idea, we further demonstrate that under the assumption that the representations within each class are concentrated (e.g., following a sub-Gaussian distribution), the downstream performance can be guaranteed.

Theorem 2 (Class-relevant Measure for Downstream Performance). Let $Z^j = \{Z : y(Z) = j\}$ be the set of examples in the j-th class with $|Z^j| = N_j$. Assume that for all $j \in [K]$, the examples $\{z_i^j\}_{i=1}^{N_j}$ in Z^j are i.i.d. R-sub-Gaussian random vectors in \mathbb{R}^d . Denote $\bar{z}^j = \frac{1}{N_j} \sum_{i=1}^{N_j} z_i^j$ and $Z^j = [z_1^j - \bar{z}^j, \cdots, z_N^j - \bar{z}^j]$ as the centered embedding matrix for Z^j . Then, for any $\delta > 0$:

$$Z_{c}^{j} = \begin{bmatrix} \mathbf{z}_{1}^{j} - \bar{\mathbf{z}}^{j}, \cdots, \mathbf{z}_{N}^{j} - \bar{\mathbf{z}}^{j} \end{bmatrix} \text{ as the centered embedding matrix for } Z^{j}. \text{ Then, for any } \delta > 0:$$

$$Err_{\mathcal{D}}(f_{\theta}) \leq \delta + \mathbb{P}_{\mathbf{z}} \left(D_{\min}^{\mathbf{z}} - \frac{\sum_{i=1}^{d} \sigma_{i}(Z_{c}^{y(\mathbf{z})})}{\sqrt{N_{y(\mathbf{z})} - 1}} < C_{\delta} \right). \tag{3}$$

$$Downstream error rate \qquad Inter-class distance \qquad Estimated intra-class radius$$

Here, $\sigma_i(\cdot)$ represents the *i*-th singular value, C_δ is a margin term that jointly determined by R, δ , and N_j (please refer to Appendix A.2.3 for exact formulation), and $D_{\min}^{\mathbf{z}} = \min_{k \in [K] \setminus y(\mathbf{z})} ||\mathbf{z} - \boldsymbol{\mu}_k||$ denotes the minimal inter-class distance.

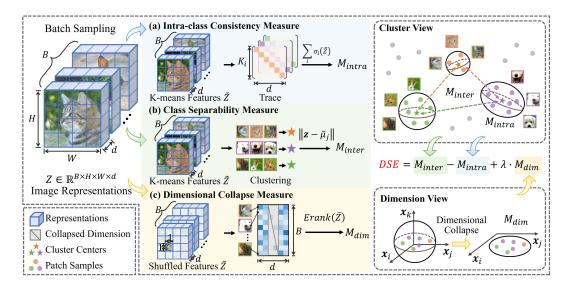


Figure 4: The proposed Dense representation Structure Estimator (DSE) consists of three components. The intra-class consistency measure and inter-class distance measure jointly conduct the class-separability measure, which is motivated by the results in Thm 2. The dimensional collapse measure estimates the effective dimensionality, which corresponds to the analysis in Cor. 5.

Remark 3. In this theorem, we estimate the intra-class radius $\mathcal{R}_{y(z)}$ with the normalized trace of the representation matrix. When this radius (plus a margin term C_{δ}) is smaller than the inter-class distance, a simple NN classifier would be enough to separate these representations.

Remark 4. For simplicity, we ignore the label estimation error in this theorem. A complete version, including the effect of k, is provided in Appendix A.4. Briefly, the bound is tightest when k equals the true number of classes, but it generally remains valid when k exceeds the actual number of classes.

Next, we reveal another key factor affecting the downstream error rate: the dimensionality of representations.

Corollary 5 (Error Rate Decay with Dimensionality). *Under Thm.* 2's assumptions with $\min_{k \neq y(z)} \|\mu_{y(z)} - \mu_k\| > \sqrt{d}R \left(2 + \sqrt{\frac{\log(8/\delta)}{N_j}} + \sqrt{3}\right)$, for any $\delta > 0$:

$$Err_{\mathcal{D}}(f_{\theta}) \le \delta + 2K \exp\left(-\tilde{C}_{\delta} \cdot d\right),$$
 (4)

where $\tilde{C}_{\delta} > 0$ is a constant.

Remark 6. This corollary establishes a connection between our analysis and the broadly studied dimensional collapse phenomenon in SSL [24, 85]. When d is small, the downstream performance experiences significant degradation. These findings highlight the critical need to jointly assess the dimensionality of representations.

4 Addressing SDD via Dense Representation Structure Estimator

4.1 Dense Representation Structure Estimator

Inspired by the analysis in the previous section, we propose a metric called Dense representation Structure Estimator (DSE) with the following formulation:

$$\text{DSE} = \underbrace{\mathbb{E}_{\boldsymbol{z}} \left[D_{\min}^{\boldsymbol{z}} - \frac{\sum_{i=1}^{d} \sigma_i(Z_c^{y(\boldsymbol{z})})}{\sqrt{(N_{y(\boldsymbol{z})} - 1)}} \right]}_{\text{Class separability measure}} + \underbrace{\lambda \cdot M_{dim.}}_{\text{Effective dimensionality}}$$

The first term is derived from the result of Thm. 2. Since the cumulative density function in Eq. 3 decreases monotonically with the deviation, the deviation between two terms can be treated as a

Algorithm 1 DSE-based Model Selection

Input: Training dataset X, checkpoints $\{f_{\theta}^i\}_{i=1}^N$, maximum number of candidates T.

for i = 1 to N do

Sample a batch of data \bar{X} from X and calculate the dense representations $Z = f^i_{\theta}(X)$ Calculate the metric P_i based on Eq. 5

end for

Select the local maximum points by $\bar{C} = \{i: i = \arg\max_{j \in [i-2, i+2]} P_j \}$

Keep the indices in \bar{C} with the top-T metric and obtain $C = \{f_{\theta}^i : i \in \bar{C}\}$.

Output: Model candidates C.

measure of downstream performance. The second term corresponds to the analysis in Cor. 5. For an intuitive understanding, readers can refer to Fig. 4.

Measuring Class Separability. Given a batch of dense representations $Z = \{z_i\}_{i=1}^{B \times N}$, we first calculate the k-means on all $B \times N$ representations to obtain a pseudo-label $\tilde{y}(z) \in [k]$ for all representations z. Let $\tilde{Z}^j = \{z \in Z : y(z) = j\}$ denotes the representations in the j-th cluster and $\tilde{N}_j = |\tilde{Z}^j|$ represents the number of representations, the intra-class radius and inter-class distance are calculated as:

$$M_{intra} = \frac{1}{k} \sum_{j=1}^{k} \frac{\sum_{i=1}^{\min{\{\tilde{N}_{j},d\}}} \sigma_{i}(\tilde{Z}_{c}^{j})}{\sqrt{(\tilde{N}_{j}-1)}}, \quad M_{inter} = \frac{1}{k} \sum_{j=1}^{k} \frac{1}{N_{j}} \sum_{\boldsymbol{z} \in \tilde{Z}^{j}} \min_{i \neq j} \|\boldsymbol{z} - \tilde{\boldsymbol{\mu}}_{i}\|_{2}.$$

where $\tilde{Z}_c^j = \tilde{Z}^j - \mathbf{1} \frac{1}{\tilde{N}_j} \sum_{i=1}^{\tilde{N}_j} \boldsymbol{z}_i^T$ is the centered representation matrix, and $\tilde{\boldsymbol{\mu}}_j = \frac{1}{|\tilde{Z}^j|} \sum_{\boldsymbol{z} \in \tilde{Z}^j} \boldsymbol{z}$ represents the center of the j-th cluster.

Measuring Dimensional Collapse. As discussed in Cor. 5, the dimensionality of representations affects their separability, and thus it should also be considered. Building on previous work [24], we first randomly sample B' dense representations from different images (ensuring their independence) and concatenate them to a $B' \times d$ matrix \bar{Z} . By setting $B' \gg d$, the rank of \bar{Z} reflects the number of non-collapsed dimensions of representations. Thus, we compute its effective rank [53]:

$$M_{dim} = \operatorname{Erank}(\bar{Z}) = \exp\bigg(-\sum_{i=1}^{d} p_i \log p_i\bigg),$$

where $p_i=rac{\sigma_i(ar{Z})}{||\sigma_i(ar{Z})||_1}$ is the *i*-th normalized singular value of $ar{Z}$.

Final Formulation of DSE. DSE is calculated by:

$$DSE = M_{inter} - M_{intra} + \lambda \cdot M_{dim}, \tag{5}$$

where λ is a parameter that rescales the measure of effective dimensionality to the same amplitude of class-separability statistics. In practice, it is taken as:

$$\lambda = \frac{\operatorname{Std}(M_{inter} - M_{intra})}{\operatorname{Std}(M_{dim})},$$

where $Std(\cdot)$ denotes the standard deviation calculated across all checkpoints.

4.2 Mitigating SDD Phenomenon with the DSE Metric

DSE-Guided Off-the-shelf Model Selection. When modifying the training process is not feasible, we propose selecting the best model from the saved checkpoints using the DSE metric to reduce the negative impact of the SDD phenomenon. When comparing two models, the one with a higher DSE indicates better class separability and effective dimensionality, and is therefore expected to perform better according to our theory. Based on this idea, we first compute the DSE metric and select checkpoints corresponding to local maxima as potential candidates. To reduce computational costs, we then choose the top T (T=3) checkpoints with the highest DSE values from this candidate set as our final models. The complete procedure is shown in Alg. 1.

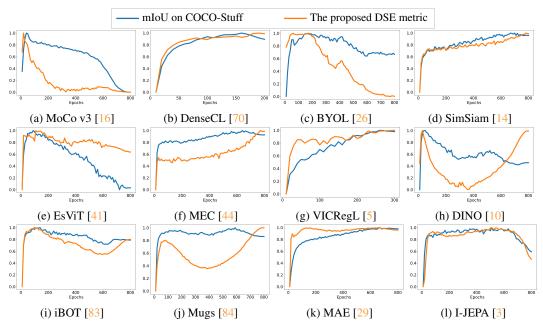


Figure 5: The proposed DSE metric precisely predicts the downstream performance.

DSE-regularized Online Optimization. Since the DSE metric represents a lower bound on performance and all operations involved in computing DSE are differentiable, we also consider directly optimizing the DSE as an explicit regularizer. Specifically, we add the negative DSE metric to the original loss function for each learning framework:

$$\mathcal{L} = \mathcal{L}_{original} - \beta \cdot DSE.$$

In our experiments, we set λ in DSE to 1 and β to 0.001. Empirically, we find that training for 10 epochs starting from the checkpoint with the best initial performance effectively mitigates the SDD phenomenon and enhances downstream performance. More details are presented in Appendix C.

5 Empirical Studies

5.1 DSE Metric is a Precise Estimator for Dense Performance

Experiment Setup. To assess the effectiveness of the proposed DSE metric, we examine the correlation between DSE and downstream segmentation performance across sixteen SSL methods and four datasets. We leverage linear transfer learning as the downstream task, with settings remain the same as introduced in the previous section. To save time, only 2048 images ($\sim 0.16\%$) of the **training dataset** (ImageNet-1k) are used to compute the metric, and both the metric and dense performance are evaluated every 10 epochs.

To validate the correlation between the proposed metric and downstream performance, the results are evaluated using Kendall's τ coefficient [37]. The coefficient ranges from [-1,1]. When $\tau=1$, the metric is perfectly aligned with the downstream performance, and $\tau=-1$ represents that they are inversely correlated. We provide a detailed explanation in the App. D.

Analysis of the Empirical Results. We draw two main conclusions from Fig. 5 and Tab. 1. 1) The DSE metric accurately reflects downstream performance. The metric curve consistently aligns with downstream performance across different datasets and methods. Hypothesis tests yield an average Kendall's τ coefficient of 0.57, confirming the reliability of DSE. 2) Compared to existing estimators, DSE is better suited for dense tasks. Current estimators are ineffective for dense performance evaluation due to the SDD phenomenon. Even if these estimators are adapted for dense representations (specifically, the adapted RankMe [24] is equivalent to our M_{dim}), DSE still significantly outperforms them. The advantage of DSE stems from its ability to more comprehensively characterize dense performance, as analyzed in detail in the following subsection. Additional empirical analyses, omitted here due to space constraints, are provided in Appendix F.

Table 1: Left: Kendall's τ coefficient of the DSE metric. We denote the insignificant results with * (p > 0.05), otherwise, the results are significant with p < 0.005. Right: Comparison of Kendall's τ coefficients for different estimators. Methods with † are adapted to dense representations.

Method	COCO	PVOC	ADE20k	Cityscapes
MoCo v3 [16]	0.55	0.58	0.60	0.45
DenseCL [70]	0.70	0.81	0.52	0.63
BYOL [26]	0.61	0.48	0.79	0.78
SimSiam [14]	0.82	0.84	0.51	0.76
EsViT [41]	0.70	0.58	0.65	0.58
MEC [44]	0.68	0.65	0.15*	0.63
Barlow Twins [78]	0.55	0.57	0.53	0.61
VICReg [4]	0.73	0.75	0.74	0.76
VICRegL [5]	0.72	0.74	0.72	0.72
SwAV [9]	0.90	0.91	0.89	0.88
DINO [10]	0.00*	0.07^{*}	0.15*	0.42
iBOT [83]	0.68	0.64	0.49	0.07*
Mugs [84]	0.01*	0.47	0.46	0.20
ReSA [72]	0.74	0.72	0.71	0.64
MAE [29]	0.38	0.44	0.46	0.21
I-JEPA [3]	0.49	0.38	0.59	0.28

Estimator	Images ↓	COCO	VOC	ADE	City	Avg
α-ReQ [2]	25600	-0.07	-0.05	-0.05	0.09	-0.02
RankMe [24]	25600	-0.10	-0.09	-0.14	0.00	-0.08
Lidar [60]	10000	-0.37	-0.36	-0.26	-0.21	-0.30
α -ReQ [†] [2]	2048	0.17	0.19	0.11	0.10	0.14
RankMe [†] [24]	2048	0.25	0.26	0.22	0.23	0.24
Lidar [†] [60]	2048	0.38	0.37	0.33	0.23	0.33
DSE (Ours)	2048	0.58	0.60	0.56	0.49	0.57

It is noteworthy that the DSE metric is theoretically derived from the class-relevance downstream tasks like semantic segmentation. To see if DSE can accurately predict the dense performance beyound segmentation, we present more results on depth estimation task in the Appendix F.1.

5.2 DSE Metric Demystifies the SDD Phenomenon

Based on these observations, we provide insights into the causes of the SDD phenomenon from two perspectives: 1) SDD arises from insufficient class separability, reduced effective dimensionality, or both. While SSL aims to balance semantic alignment and the effective dimensionality of representations, these two objectives often involve a trade-off. For instance, fully uniform representations have high effective dimensionality but may lack separability; conversely, focusing too much on semantic alignment can lead to dimension collapse. We argue

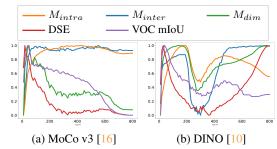


Figure 6: Different components of DSE metric.

that the degradation seen in different methods is due to the failure to maintain this trade-off during training, which causes a bias toward one objective. 2) The reasons for degradation depend specifically on the method. As illustrated in Fig. 6, MoCo v3's performance degradation is due to dimensional collapse in dense features, which reduces representation separability. Additionally, the unusual performance drop observed in DINO at around 300 epochs corresponds to a slower reduction in intra-class distances relative to inter-class distances, decreasing overall class separability.

These insights also explain why DSE improves over other estimators. Although other estimators accurately measure effective dimensionality, they fail to predict collapses caused by reduced class separability. DSE metric addresses this by providing a more complete measure of dense performance.

5.3 DSE-based Approaches Effectively Mitigate the Negative Impact of SDD

DSE-based Model Selection is Accurate and Efficient. We validate the effectiveness of our proposed model selection method on four benchmark datasets. As shown in Tab. 2, our model selection consistently improves mIoU and accuracy across all methods and datasets, achieving an average improvement of 3.0% in mIoU. Com-

DSE-based Model Selection is Accu- rate and Efficient. We validate the effectiveness of our proposed model

Table 3: Our approach achieves a performance gain comparable to supervised oracle, without requiring testing data or labels, and with negligible computational cost.

Estimator Data Label Δ mIoU \uparrow GPU hours \downarrow							
Loss Supervised DSE	√	√	-1.0 +3.6 +3.0	0.0 2.43 $0.025 (\sim 0.01 \times)$			

pared to the previous state-of-the-art method, iBOT, our approach further improves the best mIoU by an average of 2.5%.

Table 2: The overall	performance of DSE-based model selection	ı.

Method	Architecture	COCC)-Stuff	PASCA	PASCAL VOC		ADE20k		Cityscapes		
Method	Architecture	mIoU	Acc	mIoU	Acc	mIoU	Acc	mIoU	Acc		
MoCo v3 [16]	ViT-Small-16	15.1	53.9	5.9	75.5	3.9	49.7	24.9	83.5		
+MS	ViT-Small-16	30.9 (+15.8)	69.4 (+15.5)	42.0 (+36.1)	85.5 (+10.0)	16.0 (+12.1)	63.2 (+13.5)	34.8 (+9.9)	86.8 (+3.3)		
DenseCL [70]	ResNet-50	38.3	72.0	56.2	89.2	18.1	64.5	41.7	87.1		
+MS	ResNet-50	39.7 (+1.4)	72.8 (+0.8)	56.9 (+0.7)	89.5 (+0.3)	20.5 (+2.4)	66.0 (+1.5)	43.4 (+1.7)	87.8 (+0.7)		
BYOL [26]	ResNet-50	30.7	65.2	45.4	85.8	10.9	57.8	34.9	84.5		
+MS	ResNet-50	37.1 (+6.4)	70.3 (+5.1)	51.1 (+5.7)	87.9 (+2.1)	18.7 (+7.8)	63.5 (+5.7)	42.2 (+7.3)	87.3 (+2.8)		
SimSiam [14]	ResNet-50	34.5	67.8	46.6	86.7	14.4	59.7	39.1	85.4		
+MS	ResNet-50	35.0 (+0.5)	68.1 (+0.3)	47.0 (+0.4)	86.9 (+0.2)	15.6 (+1.2)	60.5 (+0.8)	39.3 (+0.2)	85.6 (+0.2)		
EsViT [41]	Swin-Tiny-7	33.4	66.3	54.3	87.6	19.4	61.3	47.8	88.9		
+MS	Swin-Tiny-7	41.6 (+8.2)	73.6 (+7.3)	59.8 (+5.5)	89.7 (+2.1)	24.4 (+5.0)	67.5 (+6.2)	50.8 (+3.0)	89.5 (+0.6)		
MEC [44]	ResNet-50	35.4	67.8	48.0	87.1	16.3	60.3	39.8	85.4		
+MS	ResNet-50	35.6 (+0.2)	68.0 (+0.2)	48.5 (+0.5)	87.2 (+0.1)	17.3 (+1.0)	60.7 (+0.4)	39.7 (-0.1)	85.4 (+0.0)		
Barlow Twins [78]	ResNet-50	36.9	69.0	51.4	88.1	19.6	62.4	42.3	86.8		
+MS	ResNet-50	37.4 (+0.5)	69.3 (+0.3)	51.6 (+0.2)	88.2 (+0.1)	19.9 (+0.3)	62.6 (+0.2)	42.8 (+0.5)	87.0 (+0.2)		
VICReg [4]	ResNet-50	36.7	69.2	52.7	88.1	19.6	62.5	42.5	86.7		
+MS	ResNet-50	37.2 (+0.5)	69.5 (+0.3)	53.0 (+0.3)	88.3 (+0.2)	19.8 (+0.2)	62.6 (+0.1)	42.8 (+0.3)	86.9 (+0.2)		
VICRegL [5]	ResNet-50	38.1	71.5	54.5	88.7	20.4	64.6	44.6	87.9		
+MS	ResNet-50	38.3 (+0.2)	71.6 (+0.1)	54.6 (+0.1)	88.8 (+0.1)	20.7 (+0.3)	64.8 (+0.2)	44.8 (+0.2)	88.0 (+0.1)		
SwAV [9]	ResNet-50	40.8	72.5	55.6	89.0	22.1	65.6	47.6	88.7		
+MS	ResNet-50	41.0 (+0.2)	72.5 (+0.0)	56.0 (+0.4)	89.0 (+0.0)	22.8 (+0.7)	65.7 (+0.1)	47.6 (+0.0)	88.7 (+0.0)		
DINO [10]	ViT-Small-16	36.0	69.7	45.8	87.3	19.2	64.8	44.3	89.1		
+MS	ViT-Small-16	40.1 (+4.1)	74.5 (+4.8)	56.3 (+10.5)	89.8 (+2.5)	22.7 (+3.5)	68.3 (+3.5)	44.3 (+0.0)	89.1 (+0.0)		
iBOT [83]	ViT-Small-16	43.2	73.6	64.4	91.7	24.1	69.6	44.6	89.2		
+MS	ViT-Small-16	45.4 (+2.2)	76.5 (+2.9)	66.8 (+2.4)	92.5 (+0.8)	27.5 (+3.4)	71.5 (+1.9)	46.5 (+1.9)	89.7 (+0.5)		
MAE [29]	ViT-Small-16	36.4	71.7	47.9	88.2	17.5	65.1	35.7	87.0		
+MS	ViT-Small-16	36.7 (+0.3)	71.9 (+0.2)	48.9 (+1.0)	88.2 (+0.0)	18.2 (+0.7)	65.5 (+0.4)	37.2 (+1.5)	87.6 (+0.6)		
Mugs [84]	ViT-Small-16	43.7	74.9	67.6	92.3	28.6	70.7	44.9	89.2		
+MS	ViT-Small-16	44.6 (+0.9)	75.9 (+1.0)	67.6 (+0.0)	92.3 (+0.0)	28.6 (+0.0)	70.7 (+0.0)	45.5 (+0.6)	89.0 (-0.2)		
ReSA [72]	ResNet-50	36.2	68.2	49.1	87.0	18.2	61.2	38.8	84.7		
+MS	ResNet-50	36.6 (+0.4)	68.5 (+0.3)	49.5 (+0.4)	87.4 (+0.4)	18.3 (+0.1)	61.2 (+0.0)	39.6 (+0.8)	85.0 (+0.3)		
I-JEPA [3]	ViT-Base-16	34.0	68.6	52.6	88.3	17.9	63.4	36.2	86.3		
+MS	ViT-Base-16	39.6 (+5.6)	72.6 (+4.0)	59.3 (+6.7)	89.7 (+1.4)	22.4 (+4.5)	66.9 (+3.5)	38.6 (+2.4)	87.4 (+1.1)		
			With DSE	Regulariza	tion —	Baseline					
67.50		1.526					0.272				
67.50		-1.526			51.1		-0.273				
	+0.4			+0.029		+1.0	,		+0.019		
67.30	_	-1.542		•	50.6		-0.283				
				/ _		0.2					
은 67.10		ద్ద-1.558	+0.013	ૄ€	50.1		ਲ -0.292	+0.00	1		
		_									
66.90		-1.575			9.6		-0.302				
66.70		-1.591		150	59.1	505	-0.311				
140 1 Ep	45 150 och	140	145 Epoch	150	590	595 € Epoch	500 -0.311	590 59 Epo	95 600 och		

Figure 7: Effect of DSE Regularization on iBOT [83] and I-JEPA [3].

(b) I-JEPA

In Tab. 3, we compare our model selection method with two baseline estimators: training loss and supervised downstream performance. Loss-based selection fails to track dense performance under the SDD phenomenon, and supervised selection is impractical due to high computational cost. In contrast, our DSE-based selection achieves competitive results with approximately $97.2\times$ speed-up, highlighting both efficiency and effectiveness.

DSE Regularization Improves Dense Performance. As shown in Fig. 7, DSE regularization consistently enhances model performance and the DSE metric. More importantly, it reverses the trend of dense degradation, demonstrating its fundamental capability to mitigate the SDD phenomenon. Results for additional methods and component ablation studies are provided in the Appendix F.9.

6 Conclusion

(a) iBOT

This paper identifies a widespread and detrimental phenomenon in self-supervised learning, termed Self-supervised Dense Degradation (SDD). Specifically, SDD occurs when dense task performance degrades over the course of training, causing the final checkpoint to exhibit a significant performance gap compared to the best intermediate checkpoint. To address SDD, we propose a Dense Representation Structure Estimator (DSE) that evaluates representation quality for downstream tasks by quantifying class separability and effective dimensionality. Our DSE is theoretically justified and empirically shown to correlate strongly with downstream task performance. To eliminate the harm of SDD, based on DSE, we introduce a checkpoint selection method for the off-the-shelf setting, and also optimize DSE directly as a regularizer. Experiments of sixteen SSL methods across four benchmark datasets confirm that these approaches effectively mitigate the negative effects of SDD.

Acknowledgments

This work was supported in part by National Natural Science Foundation of China: 62525212, 62236008, 62441232, U21B2038, and U23B2051, in part by Youth Innovation Promotion Association CAS, in part by the Strategic Priority Research Program of the Chinese Academy of Sciences, Grant No. XDB0680201, in part by the China National Postdoctoral Program for Innovative Talents, Grant No. BX20250377.

References

- [1] Andrea Agostinelli, Jasper Uijlings, Thomas Mensink, and Vittorio Ferrari. Transferability metrics for selecting source model ensembles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7936–7946, 2022.
- [2] Kumar K Agrawal, Arnab Kumar Mondal, Arna Ghosh, and Blake Richards. α-ReQ: Assessing representation quality in self-supervised learning by measuring eigenspectrum decay. Advances in Neural Information Processing Systems, 35:17626–17638, 2022.
- [3] Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. Self-supervised learning from images with a joint-embedding predictive architecture. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 15619–15629, 2023.
- [4] Adrien Bardes, Jean Ponce, and Yann LeCun. Vicreg: Variance-invariance-covariance regularization for self-supervised learning. *arXiv* preprint arXiv:2105.04906, 2021.
- [5] Adrien Bardes, Jean Ponce, and Yann LeCun. Vicregl: Self-supervised learning of local visual features. *Advances in Neural Information Processing Systems*, 35:8799–8810, 2022.
- [6] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. IEEE transactions on pattern analysis and machine intelligence, 35(8):1798–1828, 2013.
- [7] Daniel Bolya, Rohit Mittapalli, and Judy Hoffman. Scalable diverse model selection for accessible transfer learning. *Advances in Neural Information Processing Systems*, 34:19301–19312, 2021.
- [8] Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. Coco-stuff: Thing and stuff classes in context. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1209–1218, 2018.
- [9] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in neural information processing systems*, 33:9912–9924, 2020.
- [10] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021.
- [11] 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.
- [12] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E Hinton. Big self-supervised models are strong semi-supervised learners. Advances in neural information processing systems, 33:22243–22255, 2020.
- [13] Xiaokang Chen, Mingyu Ding, Xiaodi Wang, Ying Xin, Shentong Mo, Yunhao Wang, Shumin Han, Ping Luo, Gang Zeng, and Jingdong Wang. Context autoencoder for self-supervised representation learning. International Journal of Computer Vision, 132(1):208–223, 2024.
- [14] 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.
- [15] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv* preprint arXiv:2003.04297, 2020.
- [16] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9640–9649, 2021.

- [17] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 3213–3223, 2016.
- [18] Siran Dai, Qianqian Xu, Peisong Wen, Yang Liu, and Qingming Huang. Exploring non-contrastive self-supervised representation learning for image-based profiling. arXiv preprint arXiv:2506.14265, 2025.
- [19] Aditya Deshpande, Alessandro Achille, Avinash Ravichandran, Hao Li, Luca Zancato, Charless Fowlkes, Rahul Bhotika, Stefano Soatto, and Pietro Perona. A linearized framework and a new benchmark for model selection for fine-tuning. arXiv preprint arXiv:2102.00084, 2021.
- [20] Kshitij Dwivedi and Gemma Roig. Representation similarity analysis for efficient task taxonomy & transfer learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 12387–12396, 2019.
- [21] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88:303–338, 2010.
- [22] Ronald A Fisher. The use of multiple measurements in taxonomic problems. *Annals of eugenics*, 7(2): 179–188, 1936.
- [23] Keinosuke Fukunaga. Introduction to statistical pattern recognition. Elsevier, 2013.
- [24] Quentin Garrido, Randall Balestriero, Laurent Najman, and Yann LeCun. Rankme: Assessing the downstream performance of pretrained self-supervised representations by their rank. In *International Conference on Machine Learning*, pages 10929–10974. PMLR, 2023.
- [25] Mohsen Gholami, Mohammad Akbari, Xinglu Wang, Behnam Kamranian, and Yong Zhang. Etran: Energy-based transferability estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18613–18622, 2023.
- [26] 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.
- [27] 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.
- [28] 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.
- [29] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 16000–16009, 2022.
- [30] Olivier Henaff. Data-efficient image recognition with contrastive predictive coding. In *International conference on machine learning*, pages 4182–4192. PMLR, 2020.
- [31] Olivier J Hénaff, Skanda Koppula, Jean-Baptiste Alayrac, Aaron Van den Oord, Oriol Vinyals, and Joao Carreira. Efficient visual pretraining with contrastive detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10086–10096, 2021.
- [32] 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. *arXiv preprint arXiv:1808.06670*, 2018.
- [33] Long-Kai Huang, Junzhou Huang, Yu Rong, Qiang Yang, and Ying Wei. Frustratingly easy transferability estimation. In *International conference on machine learning*, pages 9201–9225. PMLR, 2022.
- [34] Weiran Huang, Mingyang Yi, Xuyang Zhao, and Zihao Jiang. Towards the generalization of contrastive self-supervised learning. arXiv preprint arXiv:2111.00743, 2021.
- [35] Yiding Jiang, Behnam Neyshabur, Hossein Mobahi, Dilip Krishnan, and Samy Bengio. Fantastic generalization measures and where to find them. arXiv preprint arXiv:1912.02178, 2019.

- [36] Neha Kalibhat, Kanika Narang, Hamed Firooz, Maziar Sanjabi, and Soheil Feizi. Measuring self-supervised representation quality for downstream classification using discriminative features. In *Proceedings of the* AAAI Conference on Artificial Intelligence, volume 38, pages 13031–13039, 2024.
- [37] Maurice G Kendall. A new measure of rank correlation. Biometrika, 30(1-2):81-93, 1938.
- [38] Diederik P Kingma. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014.
- [39] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. Advances in neural information processing systems, 25, 2012.
- [40] Tim Lebailly, Thomas Stegmüller, Behzad Bozorgtabar, Jean-Philippe Thiran, and Tinne Tuytelaars. Cribo: Self-supervised learning via cross-image object-level bootstrapping. arXiv preprint arXiv:2310.07855, 2023.
- [41] Chunyuan Li, Jianwei Yang, Pengchuan Zhang, Mei Gao, Bin Xiao, Xiyang Dai, Lu Yuan, and Jianfeng Gao. Efficient self-supervised vision transformers for representation learning. *International Conference on Learning Representations (ICLR)*, 2022.
- [42] Xiaotong Li, Zixuan Hu, Yixiao Ge, Ying Shan, and Ling-Yu Duan. Exploring model transferability through the lens of potential energy. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5429–5438, 2023.
- [43] Yandong Li, Xuhui Jia, Ruoxin Sang, Yukun Zhu, Bradley Green, Liqiang Wang, and Boqing Gong. Ranking neural checkpoints. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2663–2673, 2021.
- [44] Xin Liu, Zhongdao Wang, Ya-Li Li, and Shengjin Wang. Self-supervised learning via maximum entropy coding. Advances in Neural Information Processing Systems, 35:34091–34105, 2022.
- [45] Yang Liu, Qianqian Xu, Peisong Wen, Siran Dai, and Qingming Huang. Not all pairs are equal: Hierarchical learning for average-precision-oriented video retrieval. In *Proceedings of the 32nd ACM International* Conference on Multimedia, pages 3828–3837, 2024.
- [46] Yang Liu, Qianqian Xu, Peisong Wen, Siran Dai, and Qingming Huang. When the future becomes the past: Taming temporal correspondence for self-supervised video representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 24033–24044, June 2025.
- [47] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor segmentation and support inference from rgbd images. In ECCV, 2012.
- [48] Cuong Nguyen, Tal Hassner, Matthias Seeger, and Cedric Archambeau. Leep: A new measure to evaluate transferability of learned representations. In *International Conference on Machine Learning*, pages 7294–7305. PMLR, 2020.
- [49] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [50] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. arXiv preprint arXiv:2304.07193, 2023.
- [51] Michal Pándy, Andrea Agostinelli, Jasper Uijlings, Vittorio Ferrari, and Thomas Mensink. Transferability estimation using bhattacharyya class separability. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9172–9182, 2022.
- [52] Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alex Sorkine-Hornung, and Luc Van Gool. The 2017 davis challenge on video object segmentation. arXiv preprint arXiv:1704.00675, 2017.
- [53] Olivier Roy and Martin Vetterli. The effective rank: A measure of effective dimensionality. In 2007 15th European signal processing conference, pages 606–610. IEEE, 2007.
- [54] Ana Sanchez-Fernandez, Elisabeth Rumetshofer, Sepp Hochreiter, and Günter Klambauer. Cloome: contrastive learning unlocks bioimaging databases for queries with chemical structures. *Nature Communications*, 14(1):7339, 2023.

- [55] Jonathan Sauder and Bjarne Sievers. Self-supervised deep learning on point clouds by reconstructing space. *Advances in neural information processing systems*, 32, 2019.
- [56] Wenqi Shao, Xun Zhao, Yixiao Ge, Zhaoyang Zhang, Lei Yang, Xiaogang Wang, Ying Shan, and Ping Luo. Not all models are equal: Predicting model transferability in a self-challenging fisher space. In European Conference on Computer Vision, pages 286–302. Springer, 2022.
- [57] Oriane Siméoni, Huy V Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, et al. Dinov3. arXiv preprint arXiv:2508.10104, 2025.
- [58] Thomas Stegmüller, Tim Lebailly, Behzad Bozorgtabar, Tinne Tuytelaars, and Jean-Philippe Thiran. Croc: Cross-view online clustering for dense visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7000–7009, 2023.
- [59] Qing Su, Anton Netchaev, Hai Li, and Shihao Ji. Flsl: Feature-level self-supervised learning. Advances in Neural Information Processing Systems, 36:6568–6581, 2023.
- [60] Vimal Thilak, Chen Huang, Omid Saremi, Laurent Dinh, Hanlin Goh, Preetum Nakkiran, Joshua M Susskind, and Etai Littwin. Lidar: Sensing linear probing performance in joint embedding ssl architectures. arXiv preprint arXiv:2312.04000, 2023.
- [61] 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.
- [62] Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What makes for good views for contrastive learning? Advances in neural information processing systems, 33:6827–6839, 2020.
- [63] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. Advances in neural information processing systems, 35: 10078–10093, 2022.
- [64] Anh T Tran, Cuong V Nguyen, and Tal Hassner. Transferability and hardness of supervised classification tasks. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 1395–1405, 2019.
- [65] Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- [66] Martin J Wainwright. High-dimensional statistics: A non-asymptotic viewpoint, volume 48. Cambridge university press, 2019.
- [67] Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International conference on machine learning*, pages 9929–9939. PMLR, 2020.
- [68] Wenhao Wang, Adam Dziedzic, Michael Backes, and Franziska Boenisch. Localizing memorization in ssl vision encoders. *arXiv preprint arXiv:2409.19069*, 2024.
- [69] Wenhao Wang, Muhammad Ahmad Kaleem, Adam Dziedzic, Michael Backes, Nicolas Papernot, and Franziska Boenisch. Memorization in self-supervised learning improves downstream generalization. arXiv preprint arXiv:2401.12233, 2024.
- [70] Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, and Lei Li. Dense contrastive learning for self-supervised visual pre-training. In *Proceedings of the IEEE/CVF conference on computer vision and* pattern recognition, pages 3024–3033, 2021.
- [71] Peisong Wen, Qianqian Xu, Siran Dai, Runmin Cong, and Qingming Huang. Semantic concentration for self-supervised dense representations learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–17, 2025. doi: 10.1109/TPAMI.2025.3609758.
- [72] Xi Weng, Jianing An, Xudong Ma, Binhang Qi, Jie Luo, Xi Yang, Jin Song Dong, and Lei Huang. Clustering properties of self-supervised learning. *arXiv preprint arXiv:2501.18452*, 2025.
- [73] 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.

- [74] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9653–9663, 2022.
- [75] Huiwen Xu and U Kang. Fast and accurate transferability measurement by evaluating intra-class feature variance. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11474– 11482, 2023.
- [76] Kaichao You, Yong Liu, Jianmin Wang, and Mingsheng Long. Logme: Practical assessment of pre-trained models for transfer learning. In *International Conference on Machine Learning*, pages 12133–12143. PMLR, 2021.
- [77] Yaodong Yu, Kwan Ho Ryan Chan, Chong You, Chaobing Song, and Yi Ma. Learning diverse and discriminative representations via the principle of maximal coding rate reduction. Advances in neural information processing systems, 33:9422–9434, 2020.
- [78] 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.
- [79] Kexin Zhang, Qingsong Wen, Chaoli Zhang, Rongyao Cai, Ming Jin, Yong Liu, James Y Zhang, Yuxuan Liang, Guansong Pang, Dongjin Song, et al. Self-supervised learning for time series analysis: Taxonomy, progress, and prospects. *IEEE transactions on pattern analysis and machine intelligence*, 46(10):6775–6794, 2024.
- [80] Xiang Zhang, Ziyuan Zhao, Theodoros Tsiligkaridis, and Marinka Zitnik. Self-supervised contrastive pre-training for time series via time-frequency consistency. Advances in neural information processing systems, 35:3988–4003, 2022.
- [81] Zaiwei Zhang, Rohit Girdhar, Armand Joulin, and Ishan Misra. Self-supervised pretraining of 3d features on any point-cloud. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10252–10263, 2021.
- [82] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 633–641, 2017.
- [83] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. ibot: Image bert pre-training with online tokenizer. *arXiv preprint arXiv:2111.07832*, 2021.
- [84] Pan Zhou, Yichen Zhou, Chenyang Si, Weihao Yu, Teck Khim Ng, and Shuicheng Yan. Mugs: A multi-granular self-supervised learning framework. *arXiv preprint arXiv:2203.14415*, 2022.
- [85] Zhijian Zhuo, Yifei Wang, Jinwen Ma, and Yisen Wang. Towards a unified theoretical understanding of non-contrastive learning via rank differential mechanism. arXiv preprint arXiv:2303.02387, 2023.
- [86] Adrian Ziegler and Yuki M Asano. Self-supervised learning of object parts for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14502–14511, 2022.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The contributions and scope of the paper are accurately summarized in the abstract and introduction.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the
 contributions made in the paper and important assumptions and limitations. A No or
 NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We discuss the limitations of the work in Appendix. G.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: The assumptions are provided in Sec. 3.2 and the proofs are provided in Appendix A.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: All details are introduced in Sec. 4,.Sec. 5, and Appendix D. The code is provided in the supplementary material.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived
 well by the reviewers: Making the paper reproducible is important, regardless of
 whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
- (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We use open-source datasets for evaluation, the code is provided in the supplementary material.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be
 possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not
 including code, unless this is central to the contribution (e.g., for a new open-source
 benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how
 to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We following the default data splits for the datasets used. Other details are provided in Appendix D.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No

Justification: Though we report statistical significance tests in the paper, we do not report error bars due to the computational cost of the experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error
 of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We provide the information on the computer resources in Appendix D.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: We follow the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: In this paper, we focus on foundational research for self-supervised learning, no societal impact of the work performed.

Guidelines:

• The answer NA means that there is no societal impact of the work performed.

- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We properly credit and mention the license and terms of use of the assets used in the paper.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the
 package should be provided. For popular datasets, paperswithcode.com/datasets
 has curated licenses for some datasets. Their licensing guide can help determine the
 license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: We provide the code in the supplementary material.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects. Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The paper does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.

Appendix Contents

A	Proc	ofs	24
	A. 1	Proof of Proposition 1	24
	A.2	Proof of Theorem 2	24
		A.2.1 Preliminaries	24
		A.2.2 Important Lemmas for the Proof	25
		A.2.3 Proof of the Main Theorem	28
	A.3	Proof of Corollary 5	29
	A.4	Analysis of the Effect of k	30
В	Add	itional Related Works	32
C	Deta	niled Methodology for DSE-regularized Online Optimization	33
D	Deta	ailed Settings	34
	D.1	Pretraining	34
	D.2	Evaluation	35
	D.3	Kendall's τ Coefficient	35
	D.4	Details in DSE Calculation	35
	D.5	Computational Resources	35
E	Add	itional Experiment Results of the SDD Phenomenon	36
	E. 1	The SDD Phenomenon Exists Across Datasets	36
	E.2	The SDD Phenomenon Exists Under Varying Evaluation Protocols	38
		E.2.1 The SDD Phenomenon Exists when Backbone is Not Frozen	38
		E.2.2 The SDD Phenomenon Exists in Varying Evaluation Hyperparameters	38
	E.3	The SDD Phenomenon is not Caused by Dataset Overfitting	39
	E.4	The SDD Phenomenon Exists in Varying Downstream Tasks	39
F	Add	itional Experiment Results of the DSE Metric	40
	F.1	DSE Metric is a Precise Estimator for Depth Estimation Task	40
	F.2	DSE Precisely Indicates Performance Trends Across Unseen Datasets	40
	F.3	Ablation Study of Different Components	42
	F.4	Sensitivity Analysis on Number of Images and Clusters	42
	F.5	Analysis of Bias Introduced by Data Distribution Shift and Label Estimation Error	44
	F.6	Understanding the Reason behind SDD Phenomenon	46
	F.7	Extending the DSE Metric toward Image-level Performance Estimation	46
	F.8	Visualizations	47
	F.9	Ablation Study of DSE Regularization	49
G	Λdd	itional Discussions	50

A Proofs

A.1 Proof of Proposition 1

Proposition 7 (Restate of Prop. 1). Let $\mathcal{Z} = \{z_1, z_2, \dots, z_n\}$ be a set of points in \mathbb{R}^d , and let $\mathcal{C} = \{c_1, c_2, \dots, c_k\}$ be the set of cluster centers obtained by the K-means algorithm. For each point $z \in \mathcal{Z}$, let $c(z) \in \mathcal{C}$ denote the cluster center to which z is assigned by k-means. Then, for every $z \in \mathcal{Z}$ and for all $c_i \in \mathcal{C}$ with $c_i \neq c(z)$, the Euclidean distance satisfies

$$\|z - c(z)\|_2 \le \|z - c_i\|_2$$
.

As a result, the estimated error rate of the NN classifier is always 0:

$$\widetilde{Err}_{\mathcal{D}}(f_{\theta}) = \mathbb{E}_{\boldsymbol{x} \in \mathcal{D}}\left[\mathbb{P}_{\boldsymbol{z} \in f_{\theta}(\boldsymbol{x})}\left[||\boldsymbol{z} - \boldsymbol{c}(\boldsymbol{z})|| > \min_{\boldsymbol{c}_i \neq \boldsymbol{c}(\boldsymbol{z})}||\boldsymbol{z} - \boldsymbol{c}_i||\right]\right] \equiv 0.$$

Proof. By the definition of the K-means clustering algorithm, each point $z \in \mathcal{Z}$ is assigned to the cluster with the nearest center. Specifically, c(z) is chosen to minimize the squared Euclidean distance to z among all cluster centers in C. Formally,

$$c(z) = \arg\min_{c_j \in \mathcal{C}} \|z - c_j\|_2^2$$

Assume, for contradiction, that there exists a point $z \in Z$ and a cluster center $c_i \in C$ with $c_i \neq c(z)$ such that

$$\|z - c(z)\|_2 > \|z - c_i\|_2.$$

Squaring both sides (since the Euclidean distance is non-negative), we obtain

$$\|z - c(z)\|_2^2 > \|z - c_i\|_2^2$$
.

However, this contradicts the definition of c(z) as the cluster center that minimizes the squared distance to z. Therefore, our assumption must be false, and it must hold that

$$\|z - c(z)\|_2 \le \|z - c_i\|_2$$

for all $c_i \in \mathcal{C}$ with $c_i \neq c(z)$. It yields that:

$$\mathbb{P}_{\boldsymbol{z} \in f_{\boldsymbol{\theta}}(\boldsymbol{x})} \left[||\boldsymbol{z} - \boldsymbol{c}(\boldsymbol{z})|| > \min_{\boldsymbol{c}_i \neq \boldsymbol{c}(\boldsymbol{z})} ||\boldsymbol{z} - \boldsymbol{c}_i|| \right] \equiv 0.$$

This completes the proof.

A.2 Proof of Theorem 2

Proof Scratch. We first present a proof scratch for better understanding. The idea is straightforward: by assuming representation in each class is R-sub-Gaussian, by the property of concentration, the representations lie in the main part of the distribution with probability of at least $1-\delta$ for any $\delta>0$. If the radius is smaller than the minimal inter-class distance, all the representations in the main part would be correctly classified, leading to a final error rate smaller than $1-\delta$.

A.2.1 Preliminaries

We first list some important properties used in the derivation.

Lemma 8 (Sub-Gaussian Property). A random variable X is R-sub-Gaussian if its moment generating function satisfies:

$$\mathbb{E}[\exp(\lambda X)] \leq \exp\left(\frac{\lambda^2 R^2}{2}\right), \quad \forall \lambda \in \mathbb{R}.$$

Lemma 9 (Sub-Exponential Norm Bound). If $X \in \mathbb{R}^d$ is an R-sub-Gaussian vector with independent coordinates, then $||X||^2 = \sum_{i=1}^d X_i^2$ is sub-exponential. Specifically, the sub-exponential norm $||X^2||_{\varphi_1}$ satisfies:

$$||X^2||_{\varphi_1} \le CR^2d,$$

where C > 0 is an absolute constant.

Proof. Since each X_i is R-sub-Gaussian, X_i^2 is sub-exponential with $||X_i^2||_{\varphi_1} \leq CR^2$ (by [65], Prop 2.7.1). The sum of d independent sub-exponential variables has a norm of at most CR^2d . \square

Lemma 10 (Chernoff Bound for Sub-Exponential Tails). Let $Y = ||X||^2$ where $X \in \mathbb{R}^d$ is an R-sub-Gaussian vector with independent coordinates. Then Y is sub-exponential with $||Y||_{\psi_1} \leq CR^2d$. For any t > 0:

$$P\left(Y \ge \mathbb{E}[Y] + C_1 R^2 d(t + \sqrt{t})\right) \le e^{-t}.$$

Proof. From Lemma 9, $||Y||_{\psi_1} \leq CR^2d =: \alpha$. Using the sub-exponential tail bound:

$$P(Y - \mathbb{E}[Y] \ge t) \le \exp\left(-c\min\left(\frac{t^2}{\alpha^2}, \frac{t}{\alpha}\right)\right).$$

Set $t = \alpha(s + s^{1/2})$ for s > 0. Then:

$$\min\left(\frac{t^2}{\alpha^2}, \frac{t}{K}\right) = \min\left(s^2 + 2s^{3/2} + s, \ s + s^{1/2}\right) \ge s.$$

Thus $P(Y \geq \mathbb{E}[Y] + CR^2d(s + \sqrt{s})) \leq e^{-cs}$. Rename $s \to t/c$ for absolute constant $C_1 = C \max\{\frac{1}{c}, \frac{1}{\sqrt{c}}\}$ yields the proof.

Lemma 11 (Bernstein's Inequality [65]). Let Y_1, \ldots, Y_n be be independent sub-exponential random variables with $||Y_i||_{\varphi_1} \leq K$. For any $t \geq 0$,

$$P\left(\left|\sum_{i=1}^{n} (Y_i - \mathbb{E}[Y])\right| \ge t\right) \le 2\exp\left(-C_2 n \min\left(\frac{t^2}{K^2}, \frac{t}{K}\right)\right),$$

where $C_2 > 0$ is an absolute constant.

Lemma 12 (Norm Concentration for Sub-Gaussian Variables [65]). If $X \in \mathbb{R}^d$ is R-sub-Gaussian, then for all t > 0,

$$P\left(\|X\| \ge C_3 R(\sqrt{d} + t)\right) \le \exp(-t^2).$$

where $C_3 > 0$ is an absolute constant.

A.2.2 Important Lemmas for the Proof

Next, we proof some crucial lemmas for the proof.

Lemma 13 (Trace Concentration for Sub-Gaussian Random Vectors). Let $X \in \mathbb{R}^d$ be a mean-zero R-sub-Gaussian random variable; i.e., for all $\alpha \in \mathbb{R}^d$,

$$\mathbb{E}\big[e^{\alpha^\top X}\big] \le \exp\!\left(\frac{R^2 \|\alpha\|^2}{2}\right).$$

Let $\mu = \mathbb{E}[X]$ and $\Sigma = \text{Cov}(X)$. Suppose $\{X_i\}_{i=1}^N$ are i.i.d. copies of X, and define

$$\widehat{\Sigma} = \frac{1}{N-1} \sum_{i=1}^{N} (X_i - \overline{X}) (X_i - \overline{X})^{\top}, \quad \overline{X} = \frac{1}{N} \sum_{i=1}^{N} X_i.$$

Then with probability at least $1 - \frac{\delta}{2}$,

$$\left|\operatorname{tr}(\widehat{\Sigma}) - \operatorname{tr}(\Sigma)\right| \leq \tilde{C} R^2 \left(d\sqrt{\frac{\log(8/\delta)}{N}} + \frac{d + \log(\frac{8}{\delta})}{N}\right).$$

where $\tilde{C} > 0$ is a constant.

Proof. We know

$$tr(\Sigma) = \mathbb{E}[||X||^2] - ||\mu||^2.$$

$$\operatorname{tr}(\widehat{\Sigma}) = \frac{1}{N-1} \sum_{i=1}^{N} \left[\|X_i\|^2 \right] - \frac{N}{N-1} \|\overline{X}\|^2.$$

Rearrange the terms, we have

$$\operatorname{tr}(\widehat{\Sigma}) - \operatorname{tr}(\Sigma) \; = \; \frac{N}{N-1} \Big[\underbrace{\left(\frac{1}{N} \sum_{i=1}^N \|X_i\|^2 - \mathbb{E}[\|X\|^2]\right)}_{\operatorname{Term} \, \mathbf{A}} \; - \; \underbrace{\left(\|\overline{X}\|^2 - \|\mu\|^2\right)}_{\operatorname{Term} \, \mathbf{B}} \Big].$$

where the factor $\frac{N}{N-1} \approx 1$ only introduces a constant factor. Next, we bound *Term A* and *Term B* separately.

1. Bounding Term A. Define

Term A =
$$\frac{1}{N} \sum_{i=1}^{N} ||X_i||^2 - \mathbb{E}[||X||^2].$$

Since X is R-sub-Gaussian, Lem. 9 imply that $||X||^2$ is sub-exponential with $||X^2||_{\varphi_1} \leq CR^2$. Let $Y_i = X_i^2$, $K = CR^2$, applying Bernstein's inequality (Lem. 11) gives:

$$\mathbb{P}\Big(\big|\|X\|^2 - \mathbb{E}[\|X\|^2]\big| \ge t\Big) \le 2 \exp\Big(-C_2 N \min\Big(\frac{t^2}{(CR^2d)^2}, \frac{t}{CR^2d}\Big)\Big).$$

Next, we solve for t to achieve the desired confidence $1 - \frac{\delta}{4}$. Set the right-hand side equal to $\frac{\delta}{4}$:

$$2 \exp\left(-C_2 N \min\left(\frac{t^2}{(CR^2d)^2}, \frac{t}{CR^2d}\right)\right) = \frac{\delta}{4}.$$

This equation has two regimes:

- Small t: $\frac{t^2}{(CR^2d)^2} = \frac{1}{C_2N}\log\left(\frac{8}{\delta}\right) \Rightarrow t = CR^2d\sqrt{\frac{\log(8/\delta)}{C_2N}}$.
- Large t: $\frac{t}{CR^2d} = \frac{1}{C_2N}\log\left(\frac{8}{\delta}\right) \Rightarrow t = \frac{CR^2d\log(8/\delta)}{C_2N}$.

To unify both regimes, we define

$$t = CR^2 d\sqrt{\frac{\log(8/\delta)}{C_2 N}} + \frac{CR^2 d\log(8/\delta)}{C_2 N}.$$

Thus, with probability at least $1 - \frac{\delta}{4}$, we conclude

$$\left| \mathrm{Term} \, \mathbf{A} \right| \; \leq \; C R^2 d \sqrt{\frac{\log(8/\delta)}{C_2 N}} + \frac{C R^2 d \log(8/\delta)}{C_2 N}.$$

2. Bounding Term B. Since X is zero-mean, we write

Term B =
$$\|\overline{X}\|^2 - \|\mu\|^2 = \|\overline{X}\|^2$$
.

Applying Lem. 12 and substituting $Y = \bar{X}$ and $R \to R/\sqrt{N}$:

$$P\left(\|\bar{X}\| \ge C_3 \frac{R}{\sqrt{N}}(\sqrt{d} + t)\right) \le \exp(-t^2).$$

Set $e^{-t^2} = \delta/4 \Rightarrow t = \sqrt{\log(4/\delta)}$, substituting t into above inequality:

$$\|\bar{X}\| \le C_3 \frac{R}{\sqrt{N}} (\sqrt{d} + \sqrt{\log(4/\delta)}).$$
 with probability of at least $1 - \frac{\delta}{4}$.

Using $\sqrt{d} + \sqrt{\log(4/\delta)} \le \sqrt{2(d + \log(4/\delta))}$ (by Cauchy-Schwarz), we simplify:

$$\|\bar{X}\| \leq C_3 R \sqrt{\frac{2(d + \log(4/\delta))}{N}}. \quad \text{with probability of at least } 1 - \frac{\delta}{4}$$

Then immediately with probability at least $1 - \frac{\delta}{4}$.

$$\left| \text{Term B} \right| \le ||\bar{X}||^2 \le 2C_3^2 R^2 \frac{d + \log(\frac{4}{\delta})}{N}.$$

3. Combining the bounds. With probability at least $1 - \frac{\delta}{2}$ (by a union bound), both Term A and

Term B satisfy their respective bounds. Hence

$$\begin{split} \left| \operatorname{tr}(\widehat{\Sigma}) - \operatorname{tr}(\Sigma) \right| \; &\leq \; \frac{N}{N-1} \left(\left| \operatorname{Term} \mathbf{A} \right| + \left| \operatorname{Term} \mathbf{B} \right| \right) \\ &\leq \; 2 \left[C R^2 d \sqrt{\frac{\log(8/\delta)}{C_2 N}} + \frac{C R^2 d \log(8/\delta)}{C_2 N} + 2 C_3^2 R^2 \frac{d + \log(\frac{4}{\delta})}{N} . \right]. \end{split}$$

By selecting $\tilde{C}=2\max\{\frac{C}{\sqrt{C_2}},\frac{C}{C_2},2C_3^2\},$ we have

$$\left| \operatorname{tr}(\widehat{\Sigma}) - \operatorname{tr}(\Sigma) \right| \leq \tilde{C} R^2 \left(d\sqrt{\frac{\log(8/\delta)}{N}} + \frac{d + \log(\frac{8}{\delta})}{N} \right).$$

Choosing δ small enough or letting N grow large makes the $\frac{1}{N}$ term less significant, so the main deviation is typically on the order of

$$O\left(d\sqrt{\frac{\log(8/\delta)}{N}}\right)$$
.

This completes the proof.

Lemma 14 (Sub-Gaussian Radius Bound). Let $X \in \mathbb{R}^d$ be a mean-zero R-sub-Gaussian random vector with covariance matrix $\Sigma \in \mathbb{R}^{d \times d}$. For any t > 0, with probability at least $1 - \frac{\delta}{2}$,

$$||X|| \le \sqrt{\operatorname{trace}(\Sigma) + C_1 R^2 d\left(\log(2/\delta) + \sqrt{\log(2/\delta)}\right)}$$

Proof. Since X is R-sub-Gaussian, by Lem. 10, denote $Y = X^2$ and $\mathbb{E}[Y] = \mu$, for any t > 0:

$$P\left(Y \ge \mu + C_1 R^2 d(t + \sqrt{t})\right) \le e^{-t}.$$

By identifying $\mu = \operatorname{trace}(\Sigma)$, we arrive at

$$p(\|X\|^2 \ge \operatorname{trace}(\Sigma) + C_1 R^2 d(t + \sqrt{t})) \le e^{-t}$$

Putting $t = \log(\frac{2}{\delta})$ yields

$$p\Big(\|X\| \ \leq \ \sqrt{\operatorname{trace}(\Sigma) + C_1 R^2 d\left(\log(2/\delta) + \sqrt{\log(2/\delta)}\right)}\Big) \ \geq \ 1 - \frac{\delta}{2}.$$

This completes the proof.

Lemma 15 (Bounded Radius of Sub-Gaussian Variables). Let $Z \in \mathbb{R}^d$ be an R-sub-Gaussian random variable, and $\{z_i\}_{i=1}^N$ are i.i.d copies of Z. Denote $\bar{z} = \frac{1}{N} \sum_{i=1}^N z_i$ and

$$Z_c = [\boldsymbol{z}_1 - \bar{\boldsymbol{z}}, \quad \boldsymbol{z}_2 - \bar{\boldsymbol{z}}, \quad \cdots, \quad \boldsymbol{z}_N - \bar{\boldsymbol{z}}]$$

as the centered embedding matrix. Its singular values are $\sigma_1, \ \sigma_2, \ \dots, \ \sigma_d \ge 0$. For any $0 < \delta < 1$, with probability at least $1 - \delta$:

$$||Z - \mathbb{E}[Z]|| \leq \frac{\sum_{i=1}^{d} \sigma_i(Z_c)}{\sqrt{(N-1)}} + \sqrt{C_1 R^2 d\left(\log(2/\delta) + \sqrt{\log(2/\delta)}\right) + \tilde{C} R^2 \left(d\sqrt{\frac{\log(8/\delta)}{N}} + \frac{d + \log(\frac{8}{\delta})}{N}\right)}.$$

Proof. By Lem. 14, we have:

$$p\Big(\|Z - \mathbb{E}[Z]\| \sqrt{\operatorname{trace}(\Sigma) + C_1 R^2 d\left(\log(2/\delta) + \sqrt{\log(2/\delta)}\right)}\Big) \geq 1 - \frac{\delta}{2}.$$

Denote $\hat{\Sigma} = \frac{1}{N-1} Z_c^T Z_c$ as the centered representation matrix. By Lem. 13, with probability of at least $1 - \delta$ (with the union bound), we have:

$$||Z - \mathbb{E}[Z]|| \leq \sqrt{\operatorname{trace}(\widehat{\Sigma}) + C_1 R^2 d\left(\log(2/\delta) + \sqrt{\log(2/\delta)}\right) + \tilde{C} R^2 \left(d\sqrt{\frac{\log(8/\delta)}{N}} + \frac{d + \log(\frac{8}{\delta})}{N}\right)}.$$

With $\sqrt{a+b} \le \sqrt{a} + \sqrt{b}$, rearrange the terms:

$$\begin{split} \|Z - \mathbb{E}[Z]\| &\leq \sqrt{\operatorname{trace}(\widehat{\Sigma})} \\ &+ \sqrt{C_1 R^2 d \left(\log(2/\delta) + \sqrt{\log(2/\delta)}\right) + \tilde{C} R^2 \left(d\sqrt{\frac{\log(8/\delta)}{N}} + \frac{d + \log(\frac{8}{\delta})}{N}\right)}. \end{split}$$

Since

$$\operatorname{trace}(\widehat{\Sigma}) = \frac{1}{N-1} \sum_{i=1}^{d} \sigma_i(Z_c)^2 \le \frac{1}{N-1} \left(\sum_{i=1}^{d} \sigma_i(Z_c) \right)^2.$$

Thus, with probability at least $1 - \delta$, we have:

$$||Z - \mathbb{E}[Z]|| \leq \frac{\sum_{i=1}^{d} \sigma_i(Z_c)}{\sqrt{(N-1)}} + \sqrt{C_1 R^2 d\left(\log(2/\delta) + \sqrt{\log(2/\delta)}\right) + \tilde{C} R^2 \left(d\sqrt{\frac{\log(8/\delta)}{N}} + \frac{d + \log(\frac{8}{\delta})}{N}\right)}.$$

This completes the proof.

A.2.3 Proof of the Main Theorem

Theorem 16 (Formal version of Thm 2). Let $Z^j=\{Z:y(Z)=j\}$ be the examples in j-th class with $|Z^j|=N_j$. Assume that for all $j\in[K]$, $Z^j\in\mathbb{R}^d$ is an R-sub-Gaussian random variable, and $\{z_i^j\}_{i=1}^{N_j}$ are i.i.d copies of Z^j . Denote $\bar{z}^j=\frac{1}{N_j}\sum_{i=1}^{N_j}z_i^j$ and $Z_c^j=[z_1^j-\bar{z}^j,\ z_2^j-\bar{z}^j,\ \cdots,\ z_N^j-\bar{z}^j]$ as the centered embedding matrix for Z^j . Then, for any $\delta>0$:

$$Err_{\mathcal{D}}(f_{\theta}) \leq \delta + \tilde{P}(\delta).$$

With $\sigma_i(\cdot)$ represents the *i*-th singular value

$$\tilde{P}(\delta) = \mathbb{P}_{\boldsymbol{z}} \left(\frac{\sum_{i=1}^{d} \sigma_i(Z_c^{y(\boldsymbol{z})})}{\sqrt{(N_{y(\boldsymbol{z})} - 1)}} + C_{\delta}^{y(\boldsymbol{z})} > \min_{k \in [K] \setminus y(\boldsymbol{z})} ||\boldsymbol{z} - \boldsymbol{\mu}_k|| \right).$$

Here,

$$C_{\delta}^{y(z)} = \sqrt{C_1 R^2 d\left(\log(2/\delta) + \sqrt{\log(2/\delta)}\right) + \tilde{C} R^2 \left(d\sqrt{\frac{\log(8/\delta)}{N_{y(z)}}} + \frac{d + \log(\frac{8}{\delta})}{N_{y(z)}}\right)}$$

is a positive class-irrelevant bias term and C_1, \tilde{C} are positive constants.

Proof. Using a NN classifier, a representation z could be correctly classified if:

$$||z-\mu_{y(z)}|| \leq ||z-\mu_k||$$

holds for all $k \in [K] \setminus y(z)$.

Using the result of Lem. 14, for any class j and $\delta > 0$, the expected distance of the intra-class distance is bounded with probability at least $1 - \delta$:

$$\mathbb{E}_{\boldsymbol{z}:y(\boldsymbol{z})=j}[\|\boldsymbol{z}-\boldsymbol{\mu}_{y(\boldsymbol{z})}\|] \leq \frac{\sum_{i=1}^{d} \sigma_{i}^{j}(Z_{c})}{\sqrt{N_{j}-1}} + \sqrt{C_{1}R^{2}d\left(\log(2/\delta) + \sqrt{\log(2/\delta)}\right) + \tilde{C}R^{2}\left(d\sqrt{\frac{\log(8/\delta)}{N_{j}}} + \frac{d + \log(\frac{8}{\delta})}{N_{j}}\right)}.$$

For the sake of simplicity, we define

$$C_{\delta}^{j} = \sqrt{C_{1}R^{2}d\left(\log(2/\delta) + \sqrt{\log(2/\delta)}\right) + \tilde{C}R^{2}\left(d\sqrt{\frac{\log(8/\delta)}{N_{j}}} + \frac{d + \log(\frac{8}{\delta})}{N_{j}}\right)}.$$

For the representations in the j-th class, we separate it into two parts: the main part Z_m^j in which all examples lie in the radius, and the outside part $Z_o^j = Z^j \backslash Z_m^j$. From the concentration property, we have $Z^j = Z_m^j \cup Z_o^j$ and $P_{{\boldsymbol z} \in Z^j}({\boldsymbol z} \in Z_m^j) \ge 1 - \delta$.

For the main part, the accuracy of the NN classifier on the j-th class could be calculated by:

$$P_{\boldsymbol{z} \in Z_m^j} \left(\frac{\sum_{i=1}^d \sigma_i(Z_c^j)}{\sqrt{(N_i - 1)}} + C_\delta^j \leq \min_{k \in [K] \setminus j} ||\boldsymbol{z} - \boldsymbol{\mu}_k|| \right).$$

Thus, the error rate of the j-th class should be at least (assuming all the representations in the outside part are misclassified):

$$\operatorname{Err}_{\mathcal{D}}^{j}(f_{\boldsymbol{\theta}}) \leq \delta + P_{\boldsymbol{z} \in Z^{j}} \left(\frac{\sum_{i=1}^{d} \sigma_{i}(Z_{c}^{j})}{\sqrt{(N_{j}-1)}} + C_{\delta}^{j} > \min_{k \in [K] \setminus j} ||\boldsymbol{z} - \boldsymbol{\mu}_{k}|| \right).$$

Rearrange the terms and take an expectation on j yields the result.

A.3 Proof of Corollary 5

We first introduce some useful lemmas for the proof.

Lemma 17 (Sub-Gaussian Norm Concentration [65])). For R-sub-Gaussian vectors $X \in \mathbb{R}^d$:

$$\mathbb{P}\left(\|X\| \ge R\sqrt{d} + t\right) \le 2\exp\left(-\frac{t^2}{2R^2}\right).$$

Lemma 18 (Intra-class Concentration (Adapted from [66])). For any R-sub-Gaussian class j with N_j samples:

$$\frac{1}{\sqrt{N_j - 1}} \sum_{i=1}^d \sigma_i(Z_c^j) \le R\sqrt{d} \left(1 + \sqrt{\frac{\log(8/\delta)}{N_j}} + \frac{\log(8/\delta)}{N_j} \right)$$

holds with probability $\geq 1 - \delta/4$.

Lemma 19 (Dimensional Scaling of the Concentration Term). The concentration term C^j_{δ} admits the dimensional scaling:

$$C^j_\delta \leq R \sqrt{\frac{\sqrt{Cd\log(2/\delta)}}{\text{Sub-Gaussian term}}} + \underbrace{C'd}_{\text{Covariance term}} + \mathcal{O}\left(R\sqrt{\frac{d}{N_j}}\right).$$

For $d \ge \log(8/\delta)$, this simplifies to $C_{\delta}^{j} \le \sqrt{3}R\sqrt{d}$ with probability $\ge 1 - \delta/4$.

Corollary 20 (Formal Version of Corollary 5). *Under Thm 2's assumptions with the condition* $\min_{k \neq y(z)} \|\mu_{y(z)} - \mu_k\| > \sqrt{d}R \left(2 + \sqrt{\frac{\log(8/\delta)}{N_i}} + \sqrt{3}\right)$, for any $\delta > 0$:

$$Err_{\mathcal{D}}(f_{\theta}) \leq \delta + 2K \exp\left(-\tilde{C}_{\delta} \cdot d\right),$$

where $\tilde{C}_{\delta} = \frac{\sqrt{\frac{\log(8/\delta)}{N_j}} + \sqrt{3}}{2} > 0$ is a constant.

Proof. By the triangle inequality

$$\|m{z} - m{\mu}_k\| \ge \|m{\mu}_{y(m{z})} - m{\mu}_k\| - \|m{z} - m{\mu}_{y(m{z})}\|.$$

z would be correctly classified when:

$$\|z - \mu_{y(z)}\| \le \min_{k \ne j} \|\mu_{y(z)} - \mu_k\| - \frac{\sum_{i=1}^d \sigma_i(Z_c^{y(z)})}{\sqrt{(N_{y(z)} - 1)}} - C_\delta^{y(z)}.$$

Using Lemma 18 and Lemma 19, and denote $\Delta = \min_{k \neq y(z)} \|\boldsymbol{\mu}_{y(z)} - \boldsymbol{\mu}_k\| / \sqrt{d}$, we have:

$$\|\boldsymbol{z} - \boldsymbol{\mu}_{y(\boldsymbol{z})}\| \le \sqrt{d} \left(\Delta - R \left(1 + \sqrt{\frac{\log(8/\delta)}{N_j}} + \sqrt{3} \right) \right).$$

For the sake of simplicity, denote $\tilde{R} = R \left(1 + \sqrt{\frac{\log(8/\delta)}{N_j}} + \sqrt{3} \right)$. By the concentration of sub-Gaussian norm (17), we know:

$$\mathbb{P}\left(\|\boldsymbol{z} - \boldsymbol{\mu}_{y(\boldsymbol{z})}\| \ge R\sqrt{d} + t\right) \le 2\exp\left(-\frac{t^2}{2R^2}\right).$$

Selecting $t = (\Delta - \tilde{R} - R)\sqrt{d}$ gives:

$$\mathbb{P}\left(\|\boldsymbol{z} - \boldsymbol{\mu}_{y(\boldsymbol{z})}\| \ge \sqrt{d}(\Delta - \tilde{R})\right) \le 2 \exp\left(-d\frac{\sqrt{\frac{\log(8/\delta)}{N_j}} + \sqrt{3}}{2}\right).$$

Denote $\tilde{C}_{\delta} = \frac{\sqrt{\frac{\log(8/\delta)}{N_j}} + \sqrt{3}}{2}$ and apply the union bound across all classes yields:

$$\tilde{P}(\delta) \le 2K \exp\left(-\tilde{C}_{\delta} \cdot d\right).$$

Immediately,

$$\operatorname{Err}_{\mathcal{D}}(f_{\boldsymbol{\theta}}) \leq \delta + 2K \exp\left(-\tilde{C}_{\delta} \cdot d\right)$$

which completes the proof.

A.4 Analysis of the Effect of k

In the previous analysis, we assumed the pseudo-label to be accurate for simplicity. In this subsection, we analyze the error introduced by $k \neq C$. From the original error decomposition:

$$\operatorname{Err}(k) \leq \delta + \tilde{P}_C(\delta) = \delta + \tilde{P}_k(\delta) + \underbrace{(\tilde{P}_C(\delta) - \tilde{P}_k(\delta))}_{\Delta(k)}.$$

It is easy to tell that $\operatorname{Err}(k) \leq \delta + \tilde{P}_k(\delta) \Longrightarrow \operatorname{Err}(k) \leq \delta + \tilde{P}_C(\delta)$ when $\Delta(k) \geq 0$. In the next theorem, we model the relationship between $\Delta(k)$ and k.

Theorem 21 (Error Bound with Clustering Deviation). For any $\delta > 0$, the error bound satisfies:

$$Err \leq \delta + \tilde{P}_k(\delta) + \Delta(k)$$

where $\Delta(k) := \tilde{P}_C(\delta) - \tilde{P}_k(\delta)$ exhibits the following properties:

- $\Delta(k) \geq 0$ when k > C.
- As k decreases from C to 1, $\Delta(k)$ first increases to a positive peak, then decreases to negative values.

Proof. Intuitively, the proof starts by k=C and discusses the change of $\Delta(k)$ with two cases: 1) Increasing k leads to the splitting of original clusters, and 2) Decreasing k merges two adjacent clusters.

Case 1: k > C (Over-clustering) When k > C, since over clustering converts some false positives into true positives without affecting negative predictions, $\Delta(k) \ge 0$ naturally holds.

Case 2: k < C (Under-clustering) Denote the class center and radius as $\mu_j := \mathbb{E}[z|z \in \mathcal{C}_j]$, $\mathcal{R}_j := \sup_{z \in \mathcal{C}_j} \|z - \mu_j\|$, respectively. The distance between class centers is then defined as $d_{ij} := \|\mu_i - \mu_j\|$. When merging \mathcal{C}_1 and \mathcal{C}_2 into \mathcal{C}^* , define the merged center and radius as μ_* and \mathcal{R}_* , respectively.

The impact of merging depends on the separation between classes. Let C_1 and C_2 be two classes with separation ratio $\rho = \frac{d_{12}}{R_1 + R_2}$. We distinguish two regimes:

- Overlapping Merging ($\rho \ll 1$): Classes are poorly separated, with d_{12} small relative to their radius
- Separated Merging ($\rho \geq 1$): Classes are distinct, with d_{12} comparable to or larger than radius.

When merging two overlapping classes ($\rho \ll 1$), define the center and radius of the merged cluster $\mathcal{C}^* = \mathcal{C}_1 \cup \mathcal{C}_2$ as: $\mu_* = \frac{N_1 \mu_1 + N_2 \mu_2}{N_1 + N_2}$ and $\mathcal{R}_* \leq \max(\mathcal{R}_1, \mathcal{R}_2) + \frac{\min(N_1, N_2)}{N_1 + N_2} d_{12}$, respectively. For $z \notin \mathcal{C}^*$, the minimal distance to other classes improves due to the merged center's shift:

$$D_{\min}^{\pmb{z}}(k) \geq D_{\min}^{\pmb{z}}(C) + \underbrace{\| \pmb{\mu}_* - \pmb{\mu}_{\text{proj}} \|}_{\text{Gain from center shift}} \,,$$

where μ_{proj} is the nearest original center. This increases the margin $D_{\min}^{z}(k) - \mathcal{R}_{*}$ for non-merged classes, reducing $\tilde{P}_{k}(\delta)$.

For $z \in \mathcal{C}^*$, the radius \mathcal{R}_* remains comparable to original radius since d_{12} is small. The dominant effect is the elimination of misclassification between \mathcal{C}_1 and \mathcal{C}_2 . Thus, $\tilde{P}_k(\delta) < \tilde{P}_C(\delta)$, resulting in $\Delta(k) > 0$.

When merging well-separated classes ($\rho \ge 1$), the merged radius $\mathcal{R}_* = \max(\mathcal{R}_1, \mathcal{R}_2) + d_{12}$ becomes significantly larger. For $z \in \mathcal{C}^*$:

$$D_{\min}^{\boldsymbol{z}}(k) - \mathcal{R}_* \le \|\boldsymbol{z} - \boldsymbol{\mu}_*\| - \mathcal{R}_* \le \mathcal{R}_1 + \frac{N_2}{N_1 + N_2} d_{12} - d_{12} \ll 0,$$

increasing $\tilde{P}_k(\delta)$. For $z \notin \mathcal{C}^*$, the minimal distance may decrease slightly, but the dominant effect is the inflated \mathcal{R}_* , leading to $\tilde{P}_k(\delta) > \tilde{P}_C(\delta)$ and $\Delta(k) < 0$.

From above analysis, we see that as k decreases from C to 1, the trajectory of $\Delta(k)$ follows:

- Initial Mergers: Overlapping classes are merged first (since k-means prioritizes reducing within-cluster variance). This reduces $\tilde{P}_k(\delta)$, causing $\Delta(k) > 0$.
- Late Merges: Remaining classes are better separated, and merging them inflates \mathcal{R}_* significantly. $\tilde{P}_k(\delta)$ increases to 1, making $\Delta(k) < 0$.

Since $\Delta(k)$ could possibly be smaller than 0 when k < C, in practice, we suggest taking k equal to or slightly larger than the real number of clusters to ensure the tightest bound.

31

B Additional Related Works

Self-supervised Learning Approaches Beyond Images. In recent years, self-supervised learning has been applied to multiple modalities, including video [46, 45, 63], point clouds [81, 55], time-series data [79, 80], and cell images [18, 54]. While SSL has achieved strong results in these areas, examining whether similar fine-grained performance degradation occurs during training is a valuable direction for future work.

Studies on Class Separability. In this work, our theoretical analysis uncovers the connection between class separability and downstream performance. We clarify that this connection is not unique to our study; it has been widely explored in the machine learning community [22, 6, 23]. In the context of self-supervised learning, several studies have examined this relationship through the lens of alignment and uniformity in contrastive learning [67], or through coding rate reduction [77].

Although we do not introduce the concept of class separability, our work makes non-trivial contributions by bridging the downstream performance with measurable factors. Furthermore, we propose practical methods for evaluating the quality of dense representations.

Supervised Transferability Estimation. Due to the high computational cost of transfer learning, a large number of studies have emerged to estimate downstream performance without fine-tuning [33, 35, 76, 20, 19, 7, 1, 56, 42]. Early approaches focus on approximating the posterior distribution of target datasets [64, 48, 43]. Many works also leverage the energy score [25] or the class separability [75, 51] as the metric. While effective in supervised settings, these methods require labeled data, limiting their applicability to self-supervised scenarios.

Concurrent Studies on SSL Degradation. We also acknowledge a concurrent study, DINO v3 [57], which investigates degradation in self-supervised learning. In this paper, we show that the SDD phenomenon is widespread across sixteen state-of-the-art methods, datasets, and tasks, and we propose a theoretically grounded metric for performance estimation, model selection, and regularization. In contrast, DINO v3 focuses on degradation within the iBOT/DINO v2 family and introduces gram-matrix distillation to address it. The findings in DINO v3 strongly support the scalability of the SDD phenomenon. Their gram-matrix loss selects an early model with hand-crafted iteration steps as the teacher for correlation distillation, which could be improved by integrating our DSE-based model selection. Thus, the techniques in DINO v3 and those in this paper are likely complementary, and DINO v3 provides valuable scaling evidence for the SDD phenomenon that we do not include here due to resource limitations.

We also note that Wen et al. [71] study performance degradation during training from the perspective of insufficient semantic concentration. Their semantic concentration framework effectively mitigates the degradation by improving intra-class compactness.

C Detailed Methodology for DSE-regularized Online Optimization

As mentioned earlier, the DSE metric provides a lower bound on dense performance. Since all operations involved in computing DSE are differentiable, directly optimizing DSE can potentially address the SDD problem effectively. In practice, we include DSE explicitly as a regularizer in the training process:

$$\mathcal{L} = \mathcal{L}_{original} - \beta \cdot \text{DSE}.$$

Although calculating DSE itself is model-agnostic, integrating it into the training procedure requires certain model-specific adjustments. Specifically, all baseline methods in this study use a Joint-Embedding Self-Supervised Learning (JE-SSL) framework, which employs a siamese network with two global views of size 224×224 during pretraining. Thus, we use the student model (the encoder) to extract dense representations from these global views, and then compute the DSE based on these representations. For approaches involving masked modeling, we additionally feed the unmasked views through the student model to obtain dense representations.

To illustrate how DSE can be integrated into the training process, we take the training procedure of DINO [10] as an example. We provide the code for the DSE regularizer class in the Supplementary Material. This allows DSE regularization to be easily integrated into any JE-SSL framework by extracting dense representations from the student model and including the DSE regularizer in the loss function.

Algorithm 2 An example PyTorch pseudocode of DINO with DSE-regularized training.

D Detailed Settings

D.1 Pretraining

For pretraining, we reimplement all methods based on their original settings, but we disable automatic mixed precision (AMP) in I-JEPA [3] to prevent training instability. All models are trained for 800 epochs on ImageNet-1k [39], except for SwAV [9], VICReg [4], VICRegL [5], and DenseCL [70]. We observe model collapse when training SwAV for more epochs, and take the default settings for DenseCL, VICReg and VICRegL. The full pretraining hyperparameters are listed in Tab. 4.

Method Warm-up Epochs Architecture Learning Rate Optimizer Epochs Batch size Image size MoCo v3 [16] ViT-Small-16 AdamW 2×224^2 DenseCL [70] ResNet-50 0.03 SGD 200 2×224^2 BYOL [26] ResNet-50 0.2 LARS 10 800 1024 2×224^2 0.5 SimSiam [14] ResNet-50 800 2×224^{2} SGD 10 512 EsViT [41] Swin-Tiny-7 AdamW 800 2×224^2 5e-4 ResNet-50 0.5 10 MEC [44] SGD 2×224^2 800 Barlow Twins [78] ResNet-50 0.2 LARS 10 2048 2×224^2 2×224^{2} 0.2 VICReg [4] ResNet-50 LARS 10 1000 2048 0.2 2×224^2 ResNet-50 LARS 300 2048 VICRegL [5] 10 ResNet-50 SwAV [9] $2\times224^2+6\times96^2$ 0.6 SGD DINO [10] ViT-Small-16 AdamW 10 800 1024 $2 \times 224^2 + 10 \times 96^2$ 5e-4 $2 \times 224^{\circ} + 10 \times 96^{\circ}$ $2 \times 224^{\circ} + 10 \times 96^{\circ}$ $2 \times 224^{\circ} + 10 \times 96^{\circ}$ 10 iBOT [83] ViT-Small-16 5e-4 AdamW 800 1024 800 1024 Mugs [84] ViT-Small-16 8e-4 AdamW 10 0.5 SGD 1024 2×224^2 ResNet-50 800 ReSA [72] AdamW 40 224^{2} MAE [2 ViT-Small-16 ViT-Base-16 AdamW 800 2048 224^{2} I-JEPA [3]

Table 4: The hyperparameters used for pretraining.

For readers' convenience, we provide a brief introduction to these methods as follows:

- MoCo v3 [16]: Improves contrastive learning stability and performance by combining MoCo [28] with ViTs.
- **DenseCL** [70]: Enhances self-supervised learning by applying contrastive loss [49] at the dense feature level, enabling better dense-level representation learning.
- **BYOL**² [26]: Explores non-contrastive SSL with an asymmetric design including an extra prediction head and a stop-gradient mechanism to avoid collapse.
- SimSiam [14]: Proposes a simple siamese network for self-supervised learning without negative samples, relying on stop-gradient and predictor mechanisms to avoid collapsing solutions.
- EsViT [41]: Enhances self-supervised ViT training by combining masked image modeling with contrastive learning for improved efficiency and scalability.
- MEC [44]: Introduces the idea of maximum entropy encoding that explicitly optimizes on the structure of the representations, leading to generalizable representations.
- **Barlow Twins** [78]: Uses an invariance term and a redundancy-reduction term to optimize representation structure, effectively preventing collapse.
- VICReg [4]: Builds on Barlow Twins by adding a variance regularization loss to keep the representation space well spread.
- VICRegL [5]: Extends VICReg to local features, learning strong dense representations while maintaining image-level quality; it also shows a trade-off between coarse and fine-grained performance.
- SwAV [9]: Combines online clustering with SSL, using swapped prediction between cluster assignments to learn meaningful representations without pairwise comparisons. Assigning representations to clusters naturally enhances the class separability.
- **DINO** [10]: Leverages self-distillation with ViTs and a teacher-student framework to achieve strong instance-level representations.
- **iBOT** [83]: Integrates masked image modeling with self-distillation in features space of ViTs, enabling joint learning of global and local visual representations.

²For BYOL, we use the pytorch implementation from https://github.com/sthalles/PyTorch-BYOL

- **Mugs** [84]: Proposes a multi-granularity discriminative framework that performs well on both instance-level and dense prediction tasks.
- **ReSA** [72]: Improves clustering-based SSL by learning cluster assignments directly from encoder features, removing the need for prototypes.
- MAE [29]: Introduces a simple and scalable masked autoencoder framework for self-supervised learning, reconstructing masked patches from visible ones.
- I-JEPA [3]: Proposes a joint-embedding predictive architecture for self-supervised learning, focusing on predicting representations of masked regions in an abstract latent space.

D.2 Evaluation

Linear Semantic Segmentation. To assess dense representation quality, we adopt the standard linear evaluation protocol standard in SSL [50, 86, 83, 84]. We evaluate four benchmarks: COCO-Stuff27 [8], PASCAL VOC [21], ADE20k [82], and Cityscapes [17]. We remove projectors, and train only a lightweight classifier on dense features. For all models, we use the last layer's patch embeddings. Input images are resized to 336×336 (except 896×896 for Cityscapes to preserve detail) and the classification head is trained on 100,000 images. The head is optimized using a batch size of 256 (64 for Cityscapes due to GPU memory limitation) and a learning rate of $0.01 \times \sqrt{\text{batch size}/256}$ using an Adam [38] optimizer.

In the linear probing setup, the backbone is kept fixed, and only a small classifier is trained using the dense features. In the linear transfer learning setup, the backbone is fine-tuned with a lower learning rate.

Semi-supervised Video Object Segmentation. We evaluate the semi-supervised video object segmentation on the DAVIS 2017 dataset [52] following the [10]. During testing, the labels of the first frame are given, and we propagate predictions to subsequent frames via k-NN similarity matching of dense features.

k-NN Image Classification. (for the motivation) Following the standard k-NN setting in DINO [10], we apply k-NN classification using class tokens (or averaged patch tokens if no class token exists) mapped to ImageNet-1k labels.

D.3 Kendall's τ Coefficient

We measure the correlation between our metric and downstream performance using Kendall's τ coefficient, which is a standard metric for transferability analysis [76, 25, 75]:

$$\tau = \frac{2}{N(N-1)} \sum_{1 \leq i < j \leq N} \mathrm{sign}(M_i - M_j) \mathrm{sign}(P_i - P_j).$$

Here, N is the number of checkpoints, M_i , P_i denote the value of metrics and downstream performance calculated using i-th checkpoint, respectively. When $\tau=1$, the proposed metric is perfectly aligned with the downstream performance, and $\tau=-1$ represents that the proposed metric is inversely correlated with the downstream performance.

D.4 Details in DSE Calculation

For metric calculation, we use 2048 randomly sampled images from the pertaining dataset (ImageNet-lk), and no extra data is introduced for performance estimation. As standard testing protocol, the images are first resized to 256×256 , and then center cropped into 224×224 and normalized. For pseudo-label generation, the k-means clustering is applied. To improve robustness, we compute M_{intra} by averaging the results obtained by B=1 and B=8. We set k=3 and k=24 for k=1 and k=10 and k=12. Sensitivity analyses of these parameters are discussed in later sections.

D.5 Computational Resources

All pretraining are conducted on 8 NVIDIA A100 GPUs, the evaluation is done on 8 NVIDIA 4090 GPUs, and the DSE metric is computed on a single NVIDIA 4090 GPU. The results in Tab. 3 are obtained by averaging the time cost of 5 runs on the 4090 GPU.

E Additional Experiment Results of the SDD Phenomenon

E.1 The SDD Phenomenon Exists Across Datasets

To comprehensively analyze the SDD phenomenon, we visualize training curves for the state-of-the-art methods on the PASCAL VOC, COCO-Stuff, ADE20k, and Cityscapes datasets in Fig. 8, Fig. 9, Fig. 10, and Fig. 11, respectively. The SDD phenomenon persists consistently across all datasets, with training curves exhibiting similar degradation patterns. These results confirm that SDD is dataset-agnostic and closely linked to dense representation quality.

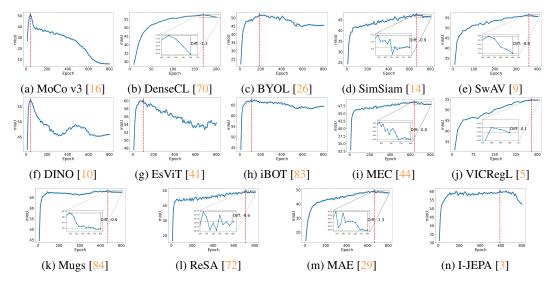


Figure 8: The SDD phenomenon on PASCAL VOC.

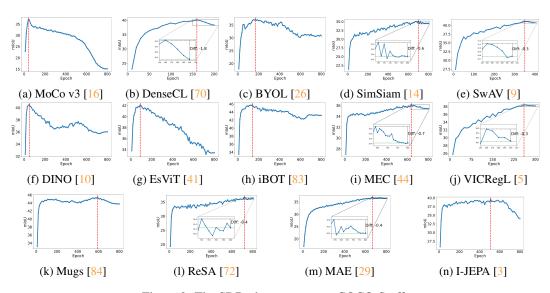


Figure 9: The SDD phenomenon on COCO-Stuff.

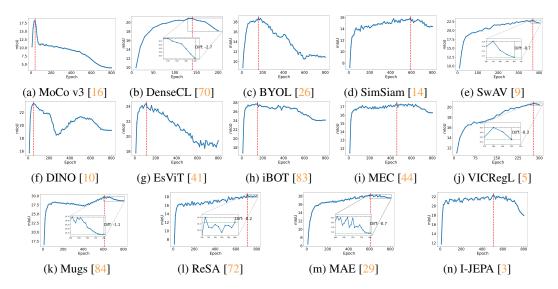


Figure 10: The SDD phenomenon on ADE20k.

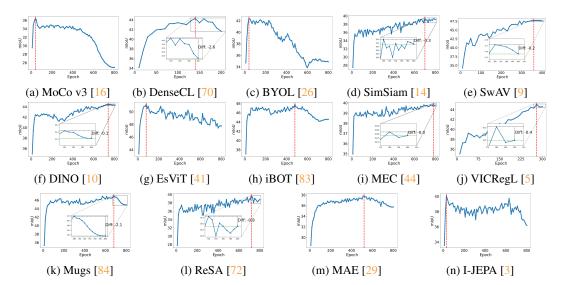


Figure 11: The SDD phenomenon on Cityscapes.

E.2 The SDD Phenomenon Exists Under Varying Evaluation Protocols

E.2.1 The SDD Phenomenon Exists when Backbone is Not Frozen.

To examine how changes to the backbone during fine-tuning influence downstream performance, we evaluated checkpoints where the backbone was unfrozen and present the results in Tab. 5. The findings reveal a pattern similar to that seen in the linear probing setting, with an average decrease of 2.9% in mIoU. While fine-tuning the backbone notably improves downstream performance, the decline becomes less pronounced as the backbone continues to adapt during training. Nevertheless, a similar trend to the fixed-backbone setting remains, further supporting the existence of the SDD phenomenon.

Table 5: A performance gap between the best and the last models is present across all datasets at	ıd
methods when backbone is not frozen.	

Method Type	Method	Architecture	rehitecture COCO-Stuff		PASCAL VOC				ADE20k		
wichiod Type	Method	Architecture	Best	Last	Diff	Best	Last	Diff	Best	Last	Diff
Contrastive	MoCo v3 [16]	ViT-Small-16	39.6	35.4	-4.2	57.3	18.1	-39.2	21.0	6.1	-14.9
Contrastive	DenseCL [70]	ResNet-50	43.2	42.4	-0.8	62.3	61.8	-0.5	25.0	24.8	-0.2
	MEC [44]	ResNet-50	41.4	41.2	-0.2	57.8	57.6	-0.2	22.0	22.0	0.0
	SimSiam [14]	ResNet-50	42.7	42.6	-0.1	60.0	59.9	-0.1	22.5	22.5	0.0
Non-Contrastive	SwAV [9]	ResNet-50	38.8	38.8	0.0	55.3	55.1	-0.2	20.1	20.0	-0.1
Non-Contrastive	DINO [10]	ViT-Small-16	42.6	41.0	-1.6	62.0	60.6	-1.4	25.6	25.5	-0.1
	EsViT [41]	Swin-Tiny-7	42.9	39.5	-3.4	62.5	56.0	-6.5	25.4	22.2	-3.2
	iBOT [83]	ViT-Small-16	46.9	44.9	-2.0	69.7	68.2	-1.5	29.1	28.1	-1.0
Masked Modeling	MAE [29]	ViT-Small-16	39.1	39.0	-0.1	55.3	54.9	-0.4	19.9	19.6	-0.3
wiasked wiodeiling	I-JEPA [3]	ViT-Base-16	46.1	44.8	-1.3	70.0	67.5	-2.5	29.5	28.1	-1.4

E.2.2 The SDD Phenomenon Exists in Varying Evaluation Hyperparameters.

To investigate whether SDD stems from specific evaluation settings, we test different learning rates and fine-tuning durations for downstream tasks. As shown in Fig. 12 and Fig. 13, performance degradation trends remain consistent across hyperparameter configurations. This demonstrates that SDD is not an artifact of specific evaluation hyperparameters but reflects a general correlation with representation quality.

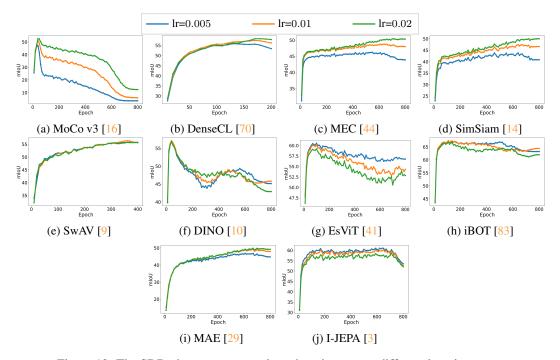


Figure 12: The SDD phenomenon consistently exists across different learning rates.

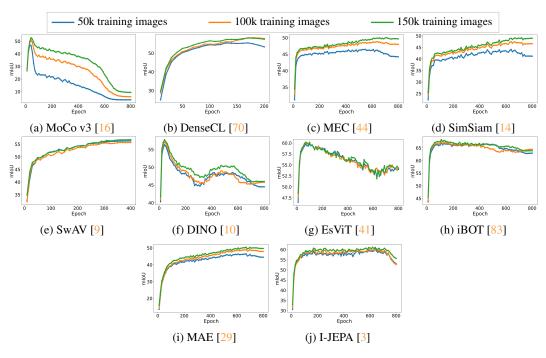


Figure 13: The SDD phenomenon consistently exists across different fine-tuning iterations.

E.3 The SDD Phenomenon is not Caused by Dataset Overfitting

To examine whether SDD is caused by dataset overfitting, we train DINO [10] and evaluate its performance on the same COCO-Stuff dataset. All settings are kept identical to those used for training on ImageNet, except for the training dataset.

As shown in Tab. 6, the model exhibits a similar performance degradation as when pretrained on ImageNet. This results in a large gap between the best and final checkpoints, suggesting that SDD is not due to dataset overfitting.

Table 6: The SDD phenomenon when training and evaluating on the same COCO dataset.

Method	Best	Last	Diff
DINO [10]	36.3	32.3	-4.0

E.4 The SDD Phenomenon Exists in Varying Downstream Tasks

We evaluate semi-supervised video object segmentation on DAVIS-2017 [52] using the protocol from [10]. Due to the limited performance of non-ViT architectures, we focus on ViT models. Results in Tab. 7 reveal persistent performance gaps between optimal and the last checkpoints, confirming that SDD generalizes to diverse dense downstream tasks (More results on depth estimation are provided in Appendix F.1).

Table 7: SDD phenomenon of Video object segmentation task on DAVIS-2017 dataset.

Method Type	Method Architecture		$\mathcal{J}\&\mathcal{F}$			$\mathcal{J}_{ ext{mean}}$				$\mathcal{F}_{ ext{mean}}$	
Memou Type	Method Architecture	Best	Last	Diff	Best	Last	Diff	Best	Last	Diff	
Contrastive	MoCo v3 [16]	ViT-Small-16	61.5	60.7	-0.8	59.6	59.0	-0.6	63.3	62.5	-0.8
Non-Contrastive	DINO [10] iBOT [83]	ViT-Small-16 ViT-Small-16		61.8 61.6	-0.1 -0.8		59.8 60.3	-0.4 -0.8	63.8 64.3	63.7 62.9	-0.1 -1.4
Masked Modeling	MAE [29] I-JEPA [3]	ViT-Small-16 ViT-Base-16	49.2 59.6	41.2 54.5			39.7 53.6	-9.0 -5.1	49.7 60.6	42.8 55.4	-6.9 -5.2

F Additional Experiment Results of the DSE Metric

F.1 DSE Metric is a Precise Estimator for Depth Estimation Task

We further examine whether SDD is unique to semantic segmentation or affects other dense tasks. First, We conduct an additional experiments on depth estimation. Similar to our linear probing setting, we freeze the model and train a linear layer to predict the depth value on NYU-depth v2 dataset [47]. As shown in Tab. 8, the SDD phenomenon persists in depth estimation task. The results validates that SDD is a general phenomenon that not specific to segmentation task.

Next, to see if the proposed DSE metric works well on depth estimation task, we also compute the Kendall's τ coefficient between the DSE metric and depth estimation performance. As shown in Tab. 9, metric positively correlates with the depth estimation performance, and DSE-based model selection consistently improves the model performance. While these results demonstrate the effectiveness of DSE across different downstream tasks, we note that the Kendall's τ coefficient is relatively lower compared with the segmentation task. The reasons are two-fold: 1) Our DSE metric is derived from the class-relevant performance, the class-seperability may not be fully useful for depth estimation task. 2) As a regression task, the RMSE curve quakes, making it hard to predict. We would like to further imporve the DSE metric to strengthen its capability on depth estimation task.

Table 8: The SDD phenomenon in depth estimation task. We report the RMSE metric (lower is better) on NYU-depth v2 dataset.

Method	Best ↓	Last ↓	Difference
MoCo v3 [16]	0.638	1.589	0.951
DenseCL [70]	0.547	0.559	0.012
SimSiam [14]	0.597	0.608	0.011
SwAV [9]	0.705	0.725	0.020
DINO [10]	0.515	0.553	0.038
MAE [29]	0.680	0.775	0.095
I-JEPA [3]	0.460	0.487	0.027

Table 9: Left: Kendall's τ coefficient between DSE and RMSE on depth estimation. Right: DSE-based model selection results. We report RMSE metric with the improvement compared to the last epoch.

Method	Kendall's $ au$	Method	RMSE ↓
MoCo v3 [16]	0.452	MoCo v3 [16]	0.638 (-0.951)
DenseCL [70]	0.752	DenseCL [70]	0.547 (-0.012)
SimSiam [14]	0.659	SimSiam [14]	0.598 (-0.010)
SwAV [9]	0.247	SwAV [9]	0.714 (-0.009)
DINO [10]	0.071	DINO [10]	0.532 (-0.021)
MAE [29]	0.107	MAE [29]	0.709 (-0.066)
I-JEPA [3]	0.190	I-JEPA [3]	0.479 (-0.008)

F.2 DSE Precisely Indicates Performance Trends Across Unseen Datasets

To comprehensively analyze the relationship between DSE and downstream performance, we present full results across all datasets and methods. These results confirm that DSE consistently identifies performance trends and peaks, establishing it as a reliable metric for model selection.

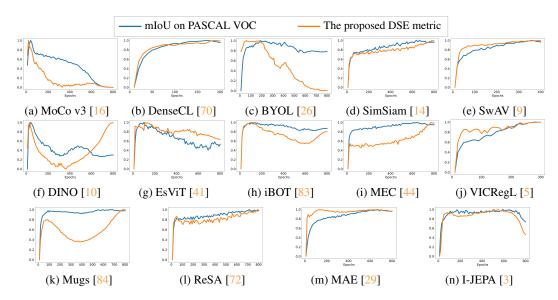


Figure 14: The proposed DSE metric precisely predicts the downstream performance on the PASCAL VOC dataset.

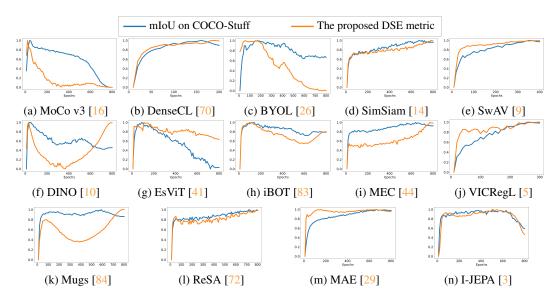


Figure 15: The proposed DSE metric precisely predicts the downstream performance on the COCO-Stuff dataset.

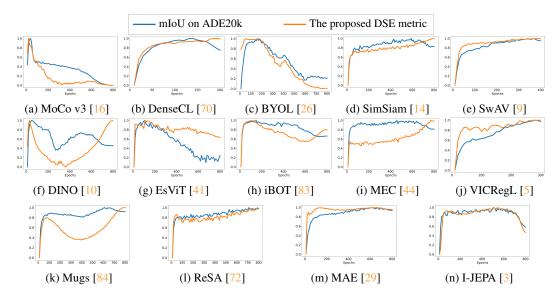


Figure 16: The proposed DSE metric precisely predicts the downstream performance on the ADE20k dataset.

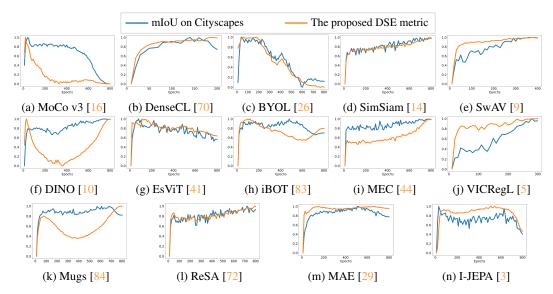


Figure 17: The proposed DSE metric precisely predicts the downstream performance on the Cityscapes dataset.

F.3 Ablation Study of Different Components.

As shown in Tab. 10, both the class-separability and dimensionality measures exhibit a positive correlation with downstream performance, supporting our theoretical findings. Combining these components further enhances Kendall's τ coefficient. Notably, when only M_{dim} is used (Line 2 in Tab. 10), it is equivalent to adapting the RankMe metric [24] to dense representations. These results highlight DSE's superior ability to assess dense representation quality.

F.4 Sensitivity Analysis on Number of Images and Clusters

While DSE itself does not explicitly depend on hyperparameters, pseudo-label generation introduces unavoidable parameters. We analyze their sensitivity to assess whether DSE's accuracy is contingent on specific choices.

Table 10: Ablation study of different components of DSE. We report the average Kendall's τ coefficient across different methods.

$M_{inter} - M_{intra}$	M_{dim}	COCO	VOC	ADE	City	Avg
√		0.45	0.42	0.33	0.37	0.39
	✓	0.25	0.26	0.22	0.23	0.24
\checkmark	✓	0.45 0.25 0.58	0.60	0.56	0.49	0.57

Effect of the number of images. By default, we only use a small amount of data (2048 images, $\sim 0.16\%$ of the training dataset). This is the minimum batch size B' to meet the requirement of $B' \gg d$ in order to obtain B' independent representations. In this part, we test whether this limited sample size suffices for robust DSE estimation. Fig. 18 shows no significant change in DSE with larger datasets, confirming that 2,048 images suffice for accurate performance estimation.

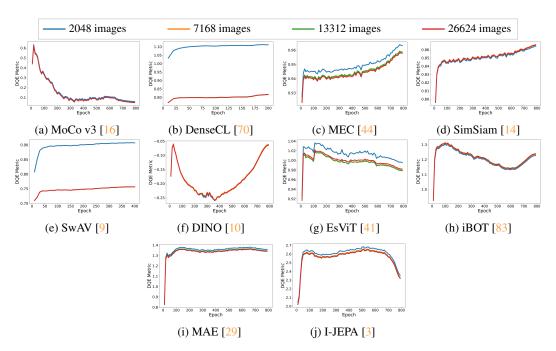


Figure 18: Sensitivity analysis of different numbers of images used for the DSE calculation.

Effect of the number of clusters. In this part, we examine the impact of the choice of k. The number of clusters is a key parameter for estimating representation quality because the actual number of classes in an image is unknown without labels. As a result, the estimated class-related metrics may be biased if the number of clusters does not match the true number of classes. Based on the earlier theoretical analysis, we recommend choosing k slightly larger than the expected number of classes. Since ImageNet is curated, we assume that the average number of classes per image (including the background as one class) is about 2 or 3. Therefore, we explore the effect of setting k between 3 and 5. Fig. 19 shows that while absolute DSE values vary with k, performance trends remain consistent, preserving DSE's utility of predicting the relative performance between models.

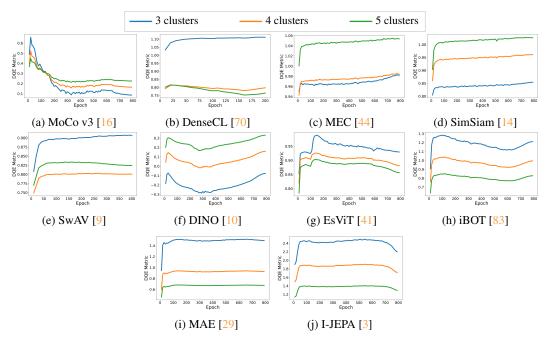


Figure 19: Sensitivity analysis of different numbers of clusters k.

F.5 Analysis of Bias Introduced by Data Distribution Shift and Label Estimation Error

When estimating downstream performance, two types of bias are introduced: 1) Data distribution shift—DSE is computed on the training dataset (i.e., ImageNet), while the target performance is evaluated on testing sets (e.g., COCO, VOC). 2) Label estimation error—DSE relies on pseudo-labels generated by the *k*-means algorithm, which may introduce bias during label assignment.

We conduct two groups of experiments to analyze their impact: 1) DSE calculated on the training set, which serves as the default setting for metric computation, and 2) DSE calculated on the testing set using ground truth labels. As shown in Fig. 21, Fig. 20, and Fig. 22 although systematic errors are present, these biases act as shifting factors rather than altering the overall trend of M_{inter} , M_{intra} , and M_{dim} . Consequently, the estimated DSE metric remains reliable for comparing model performance.

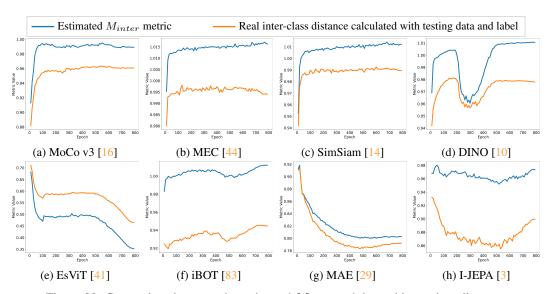


Figure 20: Comparison between the estimated M_{inter} and the real inter-class distance.

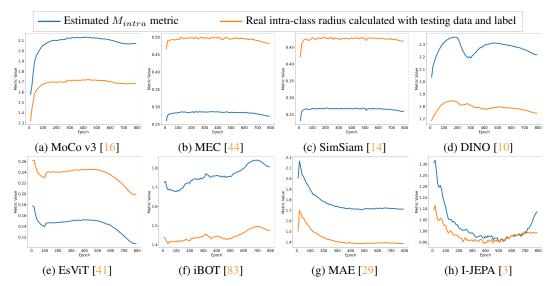


Figure 21: Comparison between the estimated M_{intra} and the real intra-class radius.

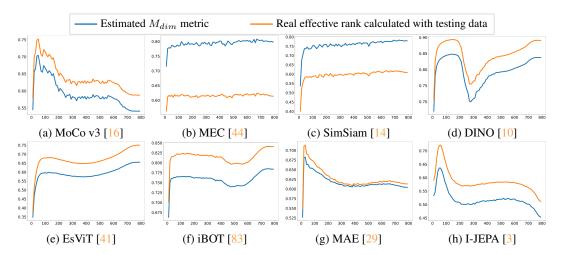


Figure 22: Comparison between the estimated M_{dim} and the real effective rank.

F.6 Understanding the Reason behind SDD Phenomenon

DSE not only measures downstream performance but also helps explain the causes of SDD. The behavior of different components of DSE offers insights into SDD. For instance, Fig. 23 shows that DINO's performance drop at 300 epochs coincides with a slower reduction in intra-class distance compared to inter-class distance, which decreases class separability. In contrast, MoCo v3's degradation is associated with a collapse in the dimensionality of dense features, making downstream separation more difficult. These findings offer a deeper understanding of the training dynamics in SSL methods and provide useful guidance for improving algorithm design.

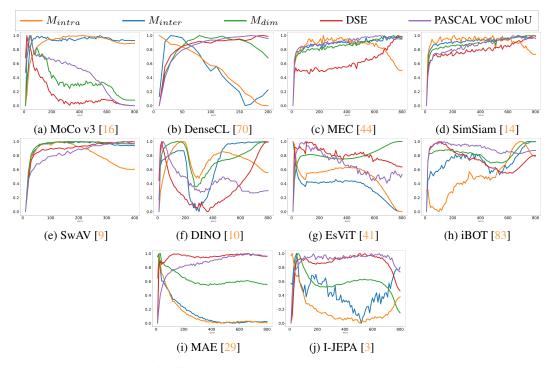


Figure 23: Visualization of different components of DSE metric, which reveals the reason behind SDD phenomenon.

F.7 Extending the DSE Metric toward Image-level Performance Estimation

Technically, our analysis should hold for both dense and image-level settings, as image-level learning can be viewed as a special case where an image is only split into one patch. To investigate how DSE performs at the image level, we conducted the following experiments. DSE originally accepts input tensors shaped as (Num_images, Num_patch_tokens, Dimensionality). For image-level tasks, we first extract the class token (or averaged dense tokens), obtaining a tensor of shape (Num_images, Dimensionality). Then, we resize this tensor to (Num_images / 200, 200, Dimensionality) and calculate the DSE metric. For evaluation, we computed Kendall's τ coefficient between the adapted DSE and ImageNet k-NN performance, comparing it with the state-of-the-art unsupervised transferability estimation method RankMe. The results are shown in the following table:

Table 11: Kendall's τ coefficient between metrics and ImageNet k-NN performance.

Method	DINO	MEC	EsViT	iBOT	I-JEPA	Average
RankMe DSE(Ours)	0.57 0.87	0.81 0.90	0.71 0.92	0.93 0.86	0.90 0.73	0.79 0.86

F.8 Visualizations

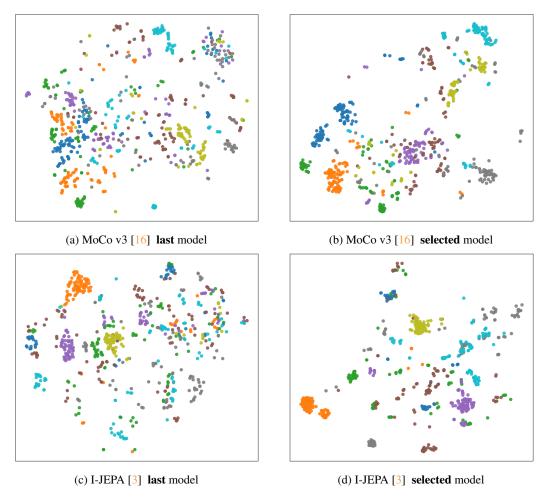


Figure 24: The selected model achieves better intra-class alignment and intra-class separatebility compared with the last model.

Furthermore, we visualize the t-SNE plot and representation similarity matrix of the last and the selected model. The inputs are sorted by class labels thus an ideal similarity matrix should be block diagonal. As spotted in Fig. 24 and Fig. 25, the selected models have better intra-class alignment and larger inter-class distance, which means that it would be easier for the downstream classifier to separate the representations. The results well support our theoretical and empirical conclusions.

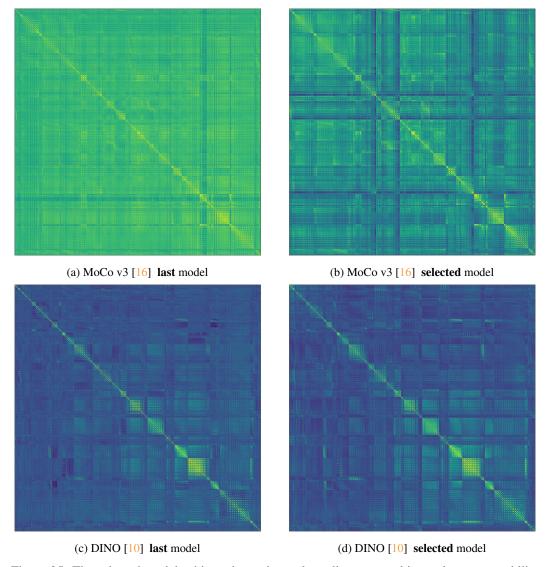


Figure 25: The selected model achieves better intra-class alignment and intra-class separatebility compared with the last model.

F.9 Ablation Study of DSE Regularization

To evaluate the effectiveness of different components of DSE, we conduct an additional ablation study in Tab. 12. We draw three main conclusions from these results:

- Different components of DSE separately contribute to improving class separability and effective dimensionality. Consequently, this improves downstream performance, which aligns with our theoretical analysis of the DSE metric.
- The performance improvement mainly comes from addressing model-specific degradation. For example, DINO's performance improves primarily due to enhanced class separability, whereas resolving dimensional collapse notably boosts dense-task performance in MoCo v3. These observations align with our model-specific analysis presented in Sec. 5.2.
- When both factors are used together, class separability often dominates optimization. Thus, introducing a hyperparameter to balance these factors may further enhance performance.

Table 12: Ablation studies of different components of DSE regularization. We report $M_{inter}-M_{intra}$, M_{dim} , and mIoU on VOC dataset.

Method	$\mid M_{inter} - M_{intra}$	M_{dim}	VOC mIoU
DINO [10]	-1.221	0.863	56.6
DINO [10]+ M_{dim}	-1.210 (+0.011)	0.884 (+0.021)	56.9 (+0.3)
DINO [10]+ M_{cls}	-1.103 (+0.118)	0.851 (-0.012)	57.4 (+0.8)
DINO [10]+ M_{dim} + M_{cls} (DSE)	-1.115 (+0.106)	0.865 (+0.002)	57.8 (+1.2)
MoCo v3 [16]	-0.958	0.744	49.1
MoCo v3 [16]+ M_{dim}	-1.021 (-0.063)	0.892 (+0.148)	52.3 (+3.2)
MoCo v3 [16]+ M_{cls}		collapsed	
MoCo v3 [16]+ M_{dim} + M_{cls} (DSE)	-0.867 (+0.091)	0.752 (+0.008)	52.1 (+3.0)

G Additional Discussions

Limitations. This paper's theoretical analysis mainly focuses on the linear probing setting and does not account for potential distributional shifts during transfer learning. When the backbone is fine-tuned for more iterations on data that differs significantly from the pre-training distribution, the estimated performance may become inaccurate.

Future Works Several prior studies on supervised transferability estimation have addressed related issues [33, 35, 76, 20, 19, 7, 1, 56, 42]. Extending the proposed DSE metric to such scenarios would be an interesting direction for future research.

It would also be valuable to investigate model-specific causes of dense degradation. For example, understanding why dimensional collapse occurs in MoCo v3 or what leads to separability degradation in DINO could provide deeper insights.

Extensions. Although the proposed DSE metric is primarily designed to address the SDD phenomenon, and we introduce two methods to reduce its negative impact, its applications are not limited to this context. For example, image-level tasks can be viewed as a special case of dense tasks, where the number of patches is one. In such cases, our theoretical analysis and the DSE metric can be applied directly.