

CLUE: FINE-GRAINED SELF-SUPERVISED LEARNING WITH MULTI-LEVEL REGULARIZATION

000
001
002
003
004
005 **Anonymous authors**
006 Paper under double-blind review
007
008
009
010

ABSTRACT

011 Self-supervised learning (SSL) has achieved strong results on coarse-grained tasks
012 but often struggles with fine-grained recognition, where categories differ only by
013 subtle local cues. For strong downstream transfer, features must form compact
014 within-class clusters with large inter-class margins at the fine level. However,
015 standard SSL losses either over-separate visually similar subcategories by treating
016 all non-positives as equally negative, or overlook part-based evidence and thus
017 merge them under coarse prototypes. We propose a multi-level regularization
018 framework that improves clustering across granularities. At the *global level*, a soft
019 variant of InfoNCE reduces false negatives and enhances class separation; at the
020 *part level*, clustering on local descriptors preserves subtle intra-class distinctions,
021 at the *instance level*, semantic descriptions from vision–language models provide
022 attribute-level anchors. Together, these components yield representations with bal-
023 anced clustering across granularities. Experiments on multiple fine-grained datasets
024 show consistent improvements in both classification and retrieval, validating the
025 effectiveness of our approach for fine-grained SSL.
026

1 INTRODUCTION

027 Self-supervised learning (SSL) has achieved remarkable success in learning visual representations
028 without human annotations, enabling models to exploit large-scale unlabeled data. Recent ad-
029 vances—spanning contrastive methods (Chen et al., 2020a; He et al., 2020) and non-contrastive
030 paradigms (Grill et al., 2020; Bardes et al., 2021; Oquab et al., 2023; Siméoni et al., 2025) have deliv-
031 ered strong performance on downstream tasks including image classification, object detection, and
032 semantic segmentation. However, SSL still underperforms in fine-grained visual recognition (FGVR),
033 where the objective is to distinguish visually similar subcategories (e.g., bird species or car models).
034 Such tasks impose stricter requirements on the discriminative ability of learned representations, as
035 subtle local differences must be captured reliably.
036

037 Recent studies have revealed that SSL representations tend to exhibit clustering behavior, where
038 learned features are naturally grouped into semantic categories. By decomposing SSL objectives
039 into an invariance term and a regularization term, (Ben-Shaul et al., 2023) demonstrate that while
040 the invariance term saturates early (e.g., in VICReg (Bardes et al., 2021)), the regularization term
041 continues to shape the geometry of feature space and is primarily responsible for the emergence
042 of semantic clustering. Such clustering structures yield well-formed feature arrangements that are
043 beneficial for downstream transfer, making the design of effective regularization particularly critical.
044

045 Although a number of SSL methods have been explored for fine-grained visual recognition (FGVR),
046 their performance remains limited. FGVR demands discrimination between visually similar sub-
047 categories (Shu et al., 2023; Wang et al., 2024), which requires preserving subtle and localized
048 cues throughout pre-training. We argue that the bottleneck lies less in invariance learning and more
049 in *granularity mismatch* in regularization: current objectives shape feature geometry mainly at a
050 coarse level, but lack explicit guidance at fine levels. As a consequence (Fig. 1), two failure modes
051 frequently arise: **Over-dispersion (left)**. Fine categories scatter excessively around coarse clusters:
052 intra-class coherence is weak at the fine level despite clear coarse-level separation. **Over-collapse**
053 (**middle**). Fine categories collapse toward their coarse centers: fine-level distinctions vanish even
054 though coarse clusters are well formed. **Ideal structure (right)**. Coarse categories are well separated
055 while fine-category centers remain distinct with balanced spacing. In Sec. 3, these phenomena are

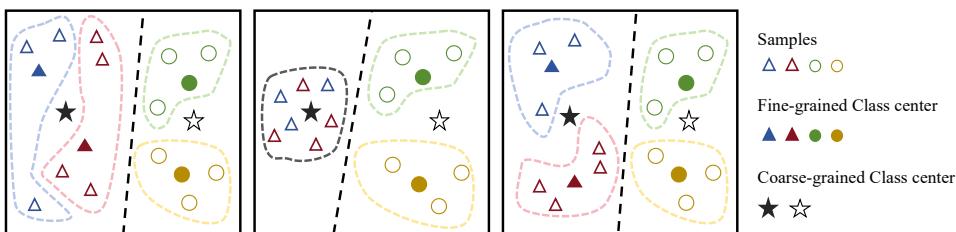


Figure 1: Illustration of granularity-related failure modes in self-supervised learning. **Left:** Over-dispersion, where coarse categories are separated but fine-grained categories scatter too widely. **Middle:** Over-collapse, where fine-grained categories collapse into coarse centers, losing intra-coarse distinctions. **Right:** Ideal structure, where coarse categories remain separated and fine-grained categories form distinct, balanced clusters. In this visualization, each *shape* represents a coarse category, each *color* within a shape denotes a fine category, solid symbols mark fine-category centers, and stars mark coarse-category centers.

quantified by geometry-aware metrics (e.g., coherence and dispersion at both coarse and fine levels) and empirical results show that multi-granularity regularization *enhances the retention of fine-grained cues*, which correlates with improved recognition performance.

Motivated by these observations, and building on the analysis of (Ben-Shaul et al., 2023) which links well-structured clustering with strong downstream transfer, we design **CLUE**, a framework that explicitly strengthens fine-grained capabilities through **CLUStEring**-aware regularization. Concretely, our approach integrates three complementary components: **(1) Class-level regularization.** We employ a soft contrastive loss that addresses the limitations of InfoNCE (Weng et al., 2025). Rather than treating all non-positives as equally negative, the soft assignment reweights pairwise relations, reducing false negatives and yielding more coherent class-level structures. **(2) Part-level regularization.** To capture subtle local differences, we extend the clustering objective to part-aware representations derived from intermediate feature maps. Applying soft assignment at this level encourages the model to disentangle fine-grained subcategories within the same coarse class, mitigating over-collapse. **(3) Instance-level regularization.** Finally, we incorporate textual guidance from vision-language models (VLMs). By aligning image features with automatically generated textual embeddings, the model is anchored to diverse semantic directions, which helps prevent instance-level collapse and further enhances fine-grained discrimination. Together, these three components form a multi-level regularization framework that balances coarse- and fine-grained structures, enabling SSL models to achieve consistently stronger performance on fine-grained recognition benchmarks.

We evaluate our approach on widely used fine-grained benchmarks, including CIFAR-100 (Krizhevsky et al., 2009), Stanford Cars (Krause et al., 2013), CUB-200 (Wah et al., 2011), and FGVC Aircraft (Maji et al., 2013). Across these datasets, our method consistently surpasses strong SSL baselines, yielding notable gains in top-1 accuracy and retrieval performance. In addition, we demonstrate that the proposed components help alleviate the granularity-related failure modes discussed above.

2 RELATED WORK

2.1 FINE-GRAINED SELF-SUPERVISED LEARNING

Self-supervised learning (SSL) has achieved remarkable progress in visual representation learning. Contrastive methods such as SimCLR (Chen et al., 2020a) and MoCo (He et al., 2020) learn by enforcing consistency between augmented views of the same image while repelling negatives, effectively shaping instance-discriminative features. Later approaches including Barlow Twins (Zbontar et al., 2021) and VICReg (Bardes et al., 2021) reduce redundancy and decorrelate features without explicit negatives, while reconstruction-based paradigms such as MAE (He et al., 2022) and BEiT (Bao et al., 2021) leverage masked image modeling to capture global structures. These methods yield strong general-purpose representations and have become standard baselines in SSL.

108 However, applying these frameworks directly to fine-grained recognition reveals critical limitations.
 109 Instance-level objectives primarily emphasize global alignment and overlook subtle local cues, which
 110 are essential for distinguishing visually similar subcategories (Cole et al., 2022). Contrastive methods
 111 can suffer from over-dispersion by aggressively separating semantically related samples, while
 112 redundancy-reduction or reconstruction-based methods may underexploit discriminative part-level
 113 information, leading to over-collapse within coarse categories. Consequently, although SSL methods
 114 transfer well to generic tasks, their ability to capture fine-grained semantics remains limited. Our
 115 work addresses this gap by introducing explicit regularization across multiple granularities to preserve
 116 both coarse- and fine-level structures.

117 2.2 HIERARCHICAL AND PART-AWARE REPRESENTATION LEARNING

120 A number of self-supervised methods for FGVR explicitly model hierarchy or parts. Prototype-
 121 based approaches (Tan et al., 2025) construct semantic prototypes and refine them stage-wise,
 122 while Particle (Saha & Maji, 2023) discovers object parts and applies contrastive learning over
 123 the discovered regions. HIRL (Xu et al., 2022) learns hierarchical image representations and S-
 124 JEA (Manová et al., 2023) stacks joint-embedding branches to capture multi-scale invariances;
 125 meanwhile, ViT-based SSL methods such as DINOv2 and CMD (Oquab et al., 2023; Bi et al., 2025)
 126 show that multi-/local-crop strategies can implicitly enhance fine-grained modeling. Overall, these
 127 works tend to operationalize hierarchy through architectural components (e.g., prototype heads, part
 128 branches, stacked embeddings) or carefully designed cropping schemes.

129 CLUE is related to these methods in that it also incorporates a lightweight part-extraction module,
 130 but its focus is different. Rather than designing a specific hierarchical architecture, CLUE starts
 131 from a *granularity-aware clustering* perspective and treats global features, part descriptors, and
 132 attribute-level (text) anchors as three coupled views that jointly shape the regularization term of
 133 the SSL objective. The part module in CLUE is used as a vehicle to extend a soft clustering loss
 134 to the local level, with the same assignment structure shared across global and part features, while
 135 the VLM-guided term introduces additional attribute-level anchors. This shifts the emphasis from
 136 building an explicit hierarchy of predictors to explicitly controlling how clusters form and separate
 137 across granularities, and makes CLUE complementary to prior hierarchical SSL approaches (Tan
 138 et al., 2025; Saha & Maji, 2023; Xu et al., 2022; Manová et al., 2023).

139 2.3 VISION-LANGUAGE MODELS AND SEMANTIC GUIDANCE

140 Large-scale vision–language models (VLMs) such as CLIP (Radford et al., 2021) and ALIGN (Jia
 141 et al., 2021) have demonstrated remarkable transferability by aligning images and text in a joint
 142 embedding space through large-scale contrastive pretraining on image–text pairs. These models
 143 enable zero-shot and open-vocabulary recognition and have inspired follow-up work (Mu et al., 2022;
 144 Gu et al., 2021) that integrates linguistic cues into visual representation learning. However, prior
 145 efforts have primarily employed VLMs for global or class-level supervision, while their potential to
 146 enhance fine-grained, part-level representations remains underexplored. In this work, we exploit VLM-
 147 generated semantic cues as external anchors to guide self-supervised learning, thereby improving
 148 fine-grained recognition especially under limited annotation.

149 3 GRANULARITY-AWARE COLLAPSE IN SSL

150 3.1 PRELIMINARIES

151 **CDNV (Class-Distance Normalized Variance).** Following Galanti et al. (2021), Class-Distance
 152 Normalized Variance (CDNV) quantifies *feature-variability collapse*, i.e., the extent to which samples
 153 from the same class are compressed into a narrow region of the feature space. Let $f : \mathbb{R}^d \rightarrow \mathbb{R}^p$
 154 denote the representation function, and let $S_1, \dots, S_C \subset \mathbb{R}^d$ be disjoint sets of samples belonging to
 155 different classes. For any pair of classes (S_i, S_j) , CDNV is defined as

$$156 \quad 161 \quad V_f(S_i, S_j) = \frac{\text{Var}_f(S_i) + \text{Var}_f(S_j)}{2 \|\mu_f(S_i) - \mu_f(S_j)\|_2^2}, \quad (1)$$

162 where $\mu_f(S)$ is the class centroid in the feature space and $\text{Var}_f(S) = \mathbb{E}_{x \in S} \|f(x) - \mu_f(S)\|_2^2$ is the
 163 within-class variance. The overall CDNV is obtained by averaging over all class pairs:
 164

$$165 \quad \text{CDNV} = \text{Avg}_{i \neq j} V_f(S_i, S_j). \quad (2)$$

167 Intuitively, CDNV compares intra-class variance with inter-class centroid distance. A lower CDNV
 168 implies that samples within a class form tight, compact clusters that are still well separated from
 169 other classes. Hence, CDNV is a useful diagnostic for whether learned features balance compactness
 170 and separability across categories.
 171

172 **NCC (Nearest Class-Center Separability).** Normalized Class Confusion (NCC) evaluates how
 173 well class centroids can serve as decision boundaries in the learned feature space. Formally, the
 174 nearest class-center (NCC) classifier is defined as
 175

$$176 \quad h(x) = \arg \min_{c \in [C]} \|f(x) - \mu_f(S_c)\|_2, \quad (3)$$

178 where $\mu_f(S_c)$ denotes the centroid of class c in the feature space. The NCC accuracy is then computed
 179 by applying $h(\cdot)$ to all samples. High NCC accuracy implies that embeddings are naturally organized
 180 around their class centroids, i.e., samples lie close to the correct centroid and far from others. This
 181 centroid-like geometry is a hallmark of SSL representations and provides an interpretable measure of
 182 class separability.
 183

184 **Granularity Variants.** To capture hierarchical structure, we extend CDNV and NCC to multiple
 185 granularities, yielding richer insights than standard fine-level metrics. Consider a dataset with C_{fine}
 186 fine classes grouped into C_{coarse} coarse classes via a mapping $\pi : [C_{\text{fine}}] \rightarrow [C_{\text{coarse}}]$. Let S_i denote
 187 the sample set of fine class i , and $S_A = \bigcup_{i: \pi(i)=A} S_i$ the union of fine classes belonging to coarse
 188 class A . We define three variants:
 189

- 190 (1) CDNV_A . For each coarse class A , we compute CDNV over the fine classes $\{S_i : \pi(i) = A\}$.
 191 This reflects the compactness and separability of fine categories within A . A lower CDNV_A
 192 means that fine classes in the same coarse group are tightly clustered around their centroids
 193 yet remain distinguishable. In contrast to the global CDNV_{all} , which aggregates over all fine
 194 classes, CDNV_A focuses specifically on intra-coarse compactness and separability.
- 195 (2) NCC_{fine} . This metric computes NCC accuracy using fine-class labels and centroids $\mu_f(S_i)$
 196 across all samples. Higher accuracy indicates that embeddings align closely with their fine-class
 197 centroids, demonstrating stronger fine-level discriminability.
- 198 (3) $\text{NCC}_{\text{coarse}}$. This variant computes NCC accuracy using coarse-class labels and centroids
 199 $\mu_f(S_A)$. A higher score suggests that fine-class samples are well organized around their
 200 coarse-class centroids, showing that the coarse semantic structure is well preserved in the
 201 representation space.

203 In summary, CDNV_A captures intra-coarse compactness and separability among fine classes within
 204 each coarse group, whereas CDNV_{all} reflects global relations across all fine classes. Similarly,
 205 NCC_{fine} and $\text{NCC}_{\text{coarse}}$ evaluate alignment with fine- and coarse-level centroids, respectively. Taken
 206 together, these metrics provide a complementary suite for quantifying clustering quality across
 207 semantic granularities.
 208

209 3.2 OBSERVATION

211 CIFAR-100 provides a hierarchical structure with 100 fine classes organized into 20 coarse categories
 212 (five per coarse group). We pretrain a ResNet-50 on this dataset and evaluate the learned represen-
 213 tations using the granularity-aware metrics introduced in Sec. 3.1. As shown in Fig. 2, the three
 214 geometric patterns illustrated in Fig. 1, over-dispersion, over-collapse, and the ideal structure, indeed
 215 manifest empirically. For reference, the dataset-level mean NCC_{fine} on the training split is 59.16%;
 unless otherwise specified, all reported results are computed on the training set.

216 **Over-dispersion (e.g., coarse ID 14).** Here $\text{NCC}_{\text{coarse}}$ is clearly above the
 217 dataset mean, yet CDNV_A for its five
 218 fine classes is markedly elevated and
 219 NCC_{fine} (34.28%) is depressed. This
 220 indicates that fine classes scatter widely
 221 despite good coarse-level separation,
 222 matching the left pattern in Fig. 1.
 223

224 **Over-collapse (e.g., coarse ID 17).** In
 225 this case CDNV_A is notably low and
 226 $\text{NCC}_{\text{coarse}}$ remains high, but NCC_{fine}
 227 (39.04%) stays below the mean. Fine
 228 categories therefore contract toward the
 229 coarse centroid and lose intra-coarse dis-
 230 tinctions, corresponding to the middle
 231 pattern in Fig. 1.

232 **Ideal structure (e.g., coarse ID 9).**
 233 Here both NCC_{fine} (73.52%) and
 234 $\text{NCC}_{\text{coarse}}$ are high, while CDNV_A re-
 235 mains relatively low. Fine clusters are
 236 compact and separable, nested within
 237 well-separated coarse groups, consistent
 238 with the right pattern in Fig. 1.
 239

240 These observations reveal that standard
 241 SSL can break down at the fine level in
 242 two opposite ways: *over-dispersion* or *over-collapse*, even when coarse structure appears satisfactory.
 243 The root cause lies in how current objectives regulate geometry. On the one hand, contrastive
 244 losses typically treat all non-positives as equally negative, which pushes apart visually similar (but
 245 semantically related) samples and leads to over-dispersion. On the other hand, existing methods
 246 underexploit *local* image cues that encode subtle distinctions, so fine classes may collapse toward
 247 their coarse centroid. Without labels, the definition of “equivalence” remains ambiguous, making
 248 this trade-off inherently challenging. Motivated by this analysis, we seek to preserve fine-grained
 249 structure while maintaining coarse-level separation. In Sec. 4, we present a multi-level regularization
 250 framework that integrates: (i) a *global-level* soft-alignment term to temper the effect of hard negatives,
 251 (ii) a *part-level* term that leverages local descriptors to preserve intra-coarse distinctions, and (iii)
 252 *semantic* cues from vision–language models to anchor instance-level uniqueness.
 253

4 METHOD

4.1 SOFT-INFONCE AS A STRONGER BASELINE REGULARIZATION

254 Following (Ben-Shaul et al., 2023), the SSL objective can be decomposed into an invariance term and
 255 a regularization term. While the invariance term quickly saturates in early training, the regularization
 256 term continues to decrease and plays the dominant role in shaping semantic clustering. An effective
 257 regularization should not only prevent representational collapse but also promote meaningful class
 258 separation. However, the standard InfoNCE loss treats all non-positive pairs as equally negative,
 259 ignoring the underlying fine-grained relations. This indiscriminate repulsion often leads to either
 260 over-dispersion or collapse toward coarse-level centers.
 261

262 To mitigate this limitation, we adopt ReSA (Weng et al., 2025) as our baseline. ReSA replaces the
 263 hard one-to-one target distribution in InfoNCE with a soft assignment matrix, thereby alleviating the
 264 uniform repulsion among non-positive samples. Formally, given a batch of $2m$ augmented samples,
 265 let $Z, Z' \in \mathbb{R}^{d \times m}$ denote the embeddings, and define the similarity matrix $S_Z = Z^\top Z'$. Standard
 266 InfoNCE minimizes the cross-entropy loss with the identity matrix I as the target distribution:
 267

$$L_{\text{InfoNCE}} = -\frac{1}{2m} \sum_{i,j} I_{ij} \log D(S_Z)_{ij} + I_{ji} \log D(S_Z^\top)_{ji}, \quad (4)$$

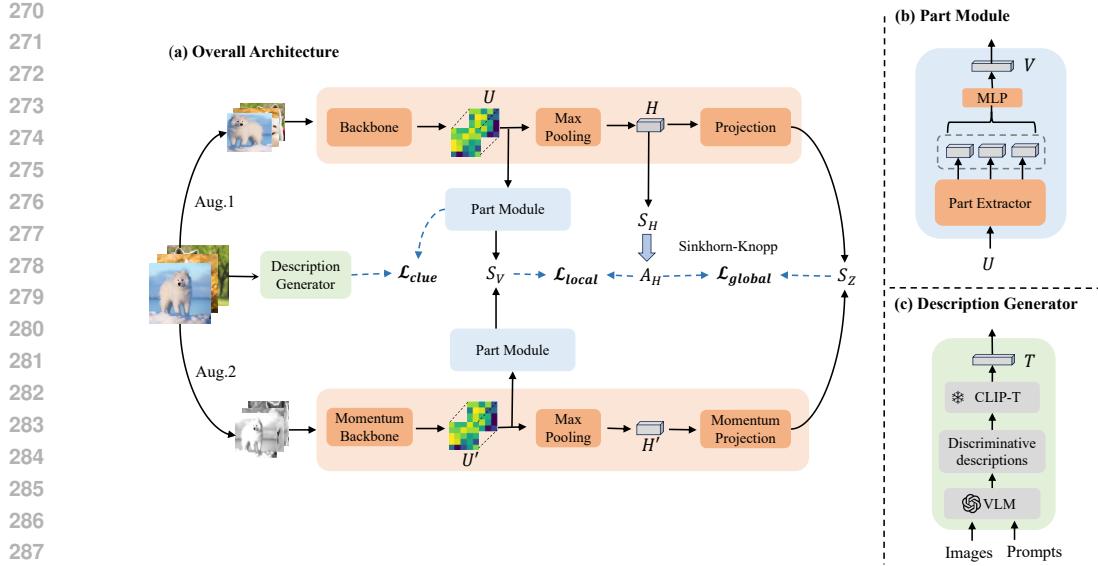


Figure 3: Pipeline of the proposed clustering-guided self-supervised learning framework (**CLUE**). (a) Overall architecture: the Sinkhorn–Knopp algorithm is applied to the similarity matrix S_H to produce soft assignments A_H , which guide both global and local contrastive losses. (b) Part-assignment module: semantic part features are extracted from activation maps and used to capture fine-grained distinctions. (c) Description generator: a vision–language model (VLM) produces discriminative textual descriptions, which are embedded by CLIP-T into the text feature space to provide semantic guidance for representation learning.

where $D(\cdot)$ denotes the row-wise softmax normalization with temperature τ . ReSA introduces an online clustering step that produces a doubly-stochastic assignment matrix $A \in \mathbb{R}^{m \times m}$ via the Sinkhorn–Knopp algorithm. The resulting soft target distribution leads to the following soft contrastive loss:

$$L_{\text{soft-InfoNCE}} = -\frac{1}{2m} \sum_{i,j} A_{ij} \log D(S_Z)_{ij} + A_{ji} \log D(S_Z^\top)_{ji}. \quad (5)$$

When $A = I$, this reduces to standard InfoNCE. The doubly-stochastic constraint enforces balanced assignments and prevents trivial collapse to a single mode. In practice, A is computed from the encoder’s intermediate features (rather than the projection head), which have been shown to provide more stable clustering (Weng et al., 2025).

Compared with InfoNCE, the gradients of $L_{\text{soft-InfoNCE}}$ reweight pairwise relations: similar samples receive larger positive weights, which reduces false negatives and preserves fine-grained structures, while truly dissimilar samples remain repelled. Thus, soft-InfoNCE offers a stronger and more principled form of semantic regularization, and we use it as the foundation for our multi-level extension in Sec. 4.

4.2 EXTENDING THE CLUSTERING LOSS TO PART-LEVEL

While the soft-InfoNCE in Sec. 4.1 mitigates over-dispersion at the *global* level, relying only on a holistic feature vector h is insufficient for distinguishing instances within the same coarse class. To explicitly capture fine-grained cues, we extend the clustering loss to operate on *part-level* features.

Part-aware representation. Given a convolutional feature map $F \in \mathbb{R}^{C \times H \times W}$, we adopt a VLAD-like residual aggregation (Arandjelovic et al., 2016) to obtain K part descriptors: $P = \{p_1, p_2, \dots, p_K\}$, $p_k = \sum_u \alpha_{uk} (f_u - c_k)$, where f_u is the local descriptor at spatial position u , c_k is the k -th part centroid, and α_{uk} is the normalized assignment weight across parts. The resulting descriptors are flattened and concatenated into a part-aware vector $v = [p_1, p_2, \dots, p_K]$.

324 **Part-level soft contrastive loss.** We replace the similarity matrix in Eq. equation 4 with part-aware
 325 similarities $S_V = V^\top V'$, where V stacks part-aware vectors from a batch. The same soft assignment
 326 matrix A is reused to construct the part-level loss:

$$328 \quad L_{\text{local}} = -\frac{1}{2m} \sum_{i,j} A_{ij} \log D(S_V)_{ij} + A_{ji} \log D(S_V^\top)_{ji}. \quad (6)$$

330 This formulation mirrors Eq. equation 5, but in the part-aware coordinate system. Rather than acting
 331 as an auxiliary objective, it imposes stronger regularization directly on discriminative local subspaces.
 332

333 **Geometric effect.** VLAD-style residuals encourage local descriptors to cluster around multiple
 334 part centroids, naturally forming multi-modal fine-grained subclusters within each coarse class. The
 335 gradient of L_{part} , proportional to $D(S_V) - A$, enforces consistency between these subclusters and
 336 the global soft assignment A . As a result, fine-grained categories are pushed apart while maintaining
 337 intra-coarse coherence, thereby mitigating over-collapse and preserving subtle distinctions that would
 338 otherwise be absorbed by coarse-class centers.

339 4.3 VLM-DRIVEN GUIDANCE FOR INSTANCE INFORMATION

341 Vision-language models (VLMs) excel at broad recognition and reasoning but often lack the gran-
 342 ularity needed for fine-grained categorization (Peng et al., 2024; Jing et al., 2024). Even when a
 343 VLM confuses closely related species at the *class* level, it typically produces *reliable region-level*
 344 *descriptions* (e.g., attributes, textures, part shapes) that domain experts rely on to separate fine
 345 categories (Zhao et al., 2025). Thus, although class predictions may be unreliable, the accompa-
 346 nying textual cues provide *stable fine-grained semantics* that can serve as anchors for representation
 347 learning. Building on this observation (Fig. 4), we regard a pre-trained VLM as an external expert
 348 and incorporate its descriptions as semantic priors, following the spirit of recent VLM-assisted
 349 approaches (Bang et al., 2024; El Banani et al., 2023; Shrivastava et al., 2021).

350 Complementary to the global and part-level
 351 objectives, text guidance addresses both over-
 352 dispersion and over-collapse. Images sharing
 353 fine-grained attributes are drawn toward the
 354 same textual anchors, which reduces variance
 355 along attribute dimensions and lowers CDNV
 356 without collapsing to zero. Meanwhile, within a
 357 common coarse category, samples with distinct
 358 attribute signatures are attracted to different
 359 anchors, effectively partitioning a coarse cluster
 360 into multiple fine-level centers.

361 To generate such anchors, we design a simple
 362 prompting strategy that (a) produces a global
 363 summary of the image, (b) enumerates salient
 364 parts and attributes, and (c) condenses these into
 365 a concise description. After light de-duplication,
 366 each description is encoded by a frozen CLIP
 367 text encoder to yield t_i . This semantic prior sta-
 368 bilizes early clustering and injects fine-grained
 369 supervision without labels. Given an image embedding v_i and its corresponding text embedding t_i ,
 we align them using a temperature-scaled cross-entropy:

$$370 \quad \mathcal{L}_{\text{text}} = -\frac{1}{m} \sum_{i=1}^m \log \frac{\exp(\langle v_i, t_i \rangle / \tau_t)}{\sum_{j=1}^m \exp(\langle v_i, t_j \rangle / \tau_t)}. \quad (7)$$

373 Without loss of generality, we combine the three objectives into a single loss:

$$374 \quad \mathcal{L}_{\text{total}} = \alpha \mathcal{L}_{\text{global}} + \beta \mathcal{L}_{\text{local}} + \gamma \mathcal{L}_{\text{text}}, \quad (8)$$

376 where $\mathcal{L}_{\text{global}}$ is the soft-InfoNCE loss from Sec. 4.1, $\mathcal{L}_{\text{local}}$ is the part-aware loss from Sec. 4.2,
 377 and $\mathcal{L}_{\text{text}}$ is the VLM-guided alignment from Sec. 4.3. The coefficients α, β, γ balance the relative
 strength of each term; unless otherwise specified, we set them all to 1.

378 Table 1: Fine-grained classification accuracy (%) and retrieval performance (Rank-1 / Rank-5, %) on
 379 CUB200, Stanford Cars, and FGVC-Aircraft. Use a ResNet-50 backbone unless otherwise specified.
 380

381 Method	382 Classification			383 Retrieval					
	384 CUB200	385 Cars	386 Aircraft	387 CUB200	388 Rank1	389 Rank5	390 Cars	391 Rank1	392 Aircraft
383 SimSiam (Chen & He, 2021)	384 46.75	385 45.72	386 38.52	387 16.24	388 –	389 12.45	390 –	391 18.49	392 –
384 MoCo v2 (Chen et al., 2020b)	385 63.98	386 62.02	387 51.13	388 39.72	389 67.14	390 30.51	391 56.15	392 30.02	393 52.87
385 LEWEL (Huang et al., 2022)	386 64.59	387 62.91	388 51.90	389 39.91	390 –	391 32.36	392 –	393 31.09	394 –
386 Contrastive Crop (Peng et al., 2022)	387 64.23	388 63.29	389 52.04	390 39.84	391 –	392 32.71	393 –	394 30.37	395 –
387 SAM-SSL-Bilinear (Shu et al., 2022)	388 64.94	389 62.85	390 52.83	391 40.08	392 –	393 33.19	394 –	395 30.52	396 –
388 MAE (He et al., 2022)	389 38.92	390 43.30	391 55.72	392 12.45	393 28.49	394 13.04	395 27.97	396 31.75	397 57.72
389 BEiT (Bao et al., 2021)	390 25.70	391 36.48	392 46.62	393 5.62	394 15.03	395 11.45	396 25.87	397 24.49	398 31.02
390 Barlow Twins (Zbontar et al., 2021)	391 33.45	392 31.91	393 34.77	394 15.24	395 38.35	396 11.99	397 30.17	398 16.32	399 35.55
391 VICReg (Bardes et al., 2021)	392 37.78	393 30.80	394 36.00	395 17.02	396 42.85	397 12.31	398 30.52	399 14.43	400 37.02
392 LCR (Shu et al., 2023)	393 65.24	394 63.96	395 53.22	396 41.26	397 –	398 34.74	399 –	400 31.55	401 –
393 LDF (Wang et al., 2024)	394 66.17	395 65.60	396 55.28	397 42.06	398 69.59	399 35.81	400 61.94	401 33.27	402 56.80
394 PAPN (Tan et al., 2025)	395 69.93	396 67.48	397 60.13	398 45.39	399 72.81	400 35.98	401 59.94	402 35.13	403 58.75
395 ReSA (Weng et al., 2025)	396 65.82	397 64.76	398 56.70	399 42.53	400 71.31	401 34.92	402 60.46	403 34.64	404 58.84
396 CLUE (Ours)	397 69.62	398 72.66	399 58.59	400 48.53	401 74.71	402 43.45	403 69.49	404 40.66	405 63.56
397 EsViT (Swin-T) (Li et al., 2021)	398 70.54	399 59.12	400 55.18	401 43.48	402 73.08	403 31.95	404 58.40	405 27.06	406 53.02
398 LoDisc (ViT-B) (Shi et al., 2025)	399 73.23	400 69.72	401 62.17	402 45.89	403 72.75	404 41.55	405 67.24	406 41.49	407 68.59
399 CLUE (ViT-B) (Ours)	400 77.83	401 75.67	402 62.76	403 61.93	404 83.91	405 52.05	406 78.14	407 40.26	408 65.11

397 5 EXPERIMENTS

398 In this section, we evaluate the performance of the proposed method on three fine-grained image
 399 datasets: Caltech UCSD-Birds (CUB200) (Wah et al., 2011), Stanford Cars (Cars) (Krause et al.,
 400 2013), FGVC-Aircraft (Aircraft) (Maji et al., 2013).

401 5.1 SETTINGS

402 5.1.1 IMPLEMENTATION DETAILS

403 For all experiments, we adopt ResNet-50 (He et al., 2016) as the backbone and a standard three-layer
 404 MLP projector. Following common practice (He et al., 2020; Grill et al., 2020; Caron et al., 2020),
 405 we use a momentum encoder with coefficient 0.999. The number of part centroids in the part-level
 406 module is set to $K = 3$, which provides a good balance between capturing discriminative parts
 407 and computational efficiency. All models are optimized with SGD under a cosine learning-rate
 408 schedule. Our training protocol is aligned with prior work (Shu et al., 2023; Wang et al., 2024) for
 409 fair comparison; further implementation details are provided in the Appendix. All experiments are
 410 conducted on four NVIDIA RTX 3090 GPUs.

411 5.1.2 EVALUATION PROTOCOLS

412 We adopt two complementary evaluation settings: **linear probing** and **image retrieval**. Linear
 413 probing is a widely used protocol in self-supervised learning (SSL). After pretraining, the backbone
 414 is frozen and a linear classifier is trained on top of the learned representations. The classification
 415 accuracy of this linear head serves as a direct measure of the discriminative quality of the features.
 416 Image retrieval evaluates how well the representations capture semantic similarity. For each query
 417 image, we perform nearest-neighbor search in the feature space and retrieve images from the gallery.
 418 Performance is measured by rank- k accuracy, i.e., whether a correct match appears among the top- k
 419 retrieved images. Unless otherwise specified, we report top-1 accuracy for linear probing and rank-1 /
 420 rank-5 accuracy for retrieval.

421 5.2 MAIN RESULTS

422 Table 1 reports a comprehensive comparison with prior methods on three fine-grained benchmarks.
 423 Our method consistently outperforms the baselines in both classification and retrieval. In particular, it
 424 surpasses LDF by **+3.45**, **+7.06**, and **+3.31** points in classification accuracy on CUB200, Cars, and
 425 Aircraft, respectively. Similar gains are observed in retrieval, where our approach achieves higher
 426 rank-1 and rank-5 accuracy across all datasets. These results highlight the robustness of our framework

432
 433 Table 2: Classification accuracy (%) and retrieval performance (Rank-1 / Rank-5, %) for different
 434 ablation configurations of the proposed CLUE framework. The baseline (#1) uses standard contrastive
 435 learning with InfoNCE loss. Each row corresponds to incrementally adding key components: global
 436 loss (soft clustering alignment), local loss (part-aware contrastive learning), and text loss (VLM-
 437 guided semantic cues).

438 ID	439 Losses			440 Classification			441 Retrieval								
	442 L_{global}	443 L_{local}	444 L_{text}	445 CUB200	446 Cars	447 Aircraft	448 CUB200	449 Rank1	450 Rank5	451 Cars	452 Rank1	453 Rank5	454 Aircraft	455 Rank1	456 Rank5
#1				62.29	60.20	51.13	37.74	57.07	31.74	56.05	31.74	52.07			
#2			✓	64.63	64.46	52.29	42.06	67.43	35.75	60.69	35.97	59.77			
#3	✓			65.82	64.76	56.70	42.53	71.31	34.92	60.46	34.64	58.84			
#4	✓	✓		66.95	65.58	57.76	43.56	72.21	35.27	62.16	36.42	60.24			
#5	✓	✓	✓	69.62	72.66	58.59	48.53	74.71	43.45	69.49	40.66	63.56			

444 Table 3: Fine-grained classification accuracy (%) and retrieval performance (Rank-1 / Rank-5, %) on
 445 CUB200 and FGVC-Aircraft for different number of clusters.

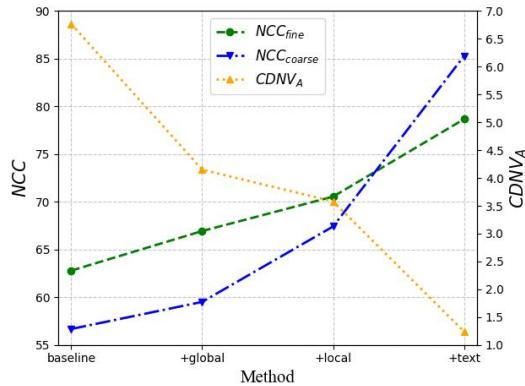
447 Clusters	448 Classification			449 Retrieval								
	450 CUB200	451 Cars	452 Aircraft	453 CUB200	454 Rank1	455 Rank5	456 Cars	457 Rank1	458 Rank5	459 Aircraft	460 Rank1	461 Rank5
2	64.60	63.20	55.93	42.46	71.47	34.27	59.16	33.74	56.07			
3	65.82	64.76	56.70	42.53	71.31	34.92	60.46	34.64	58.84			
4	66.10	65.91	56.76	43.48	71.52	34.57	59.11	33.31	57.28			
8	66.95	65.58	57.03	43.66	71.31	33.44	56.95	36.42	59.54			

454 for both non-rigid objects (CUB200) and rigid categories (Cars, Aircraft), demonstrating its ability
 455 to capture subtle inter-class differences. Moreover, when instantiated with a ViT-Base backbone,
 456 our method further improves top-1 accuracy and consistently outperforms Transformer-based SSL
 457 approaches such as EsViT Li et al. (2021) and LoDisc Shi et al. (2025) in both classification and
 458 retrieval, confirming that the proposed multi-level regularization is architecture-agnostic and remains
 459 effective on modern vision Transformers.

460 5.3 ABLATION STUDY

462 **Effect of Key Modules.** We perform ablations
 463 on CUB200, Stanford Cars, and FGVC-Aircraft
 464 to isolate the contribution of each component
 465 (Table 2). The baseline (#1) relies on standard
 466 InfoNCE contrastive learning (Eq. equation 4).
 467 Replacing it with global soft clustering align-
 468 ment (#3; soft-InfoNCE, Eq. equation 5) leads
 469 to clear gains in both classification and retrieval,
 470 highlighting the benefit of adapting to batch-
 471 level semantic structure. Adding the part-level
 472 loss (#4; Sec. 4.2) provides further improve-
 473 ments by leveraging local descriptors to cap-
 474 ture fine-grained cues within coarse categories.
 475 Additionally, incorporating the VLM-guided ob-
 476 jective (#5; Sec. 4.3) achieves the strongest re-
 477 sults, showing that external semantic anchors of-
 478 ffer complementary supervision. Taken together,
 479 these results demonstrate that global soft align-
 480 ment, part-level discrimination, and VLM guid-
 481 ance contribute in a complementary and additive manner, producing the most effective fine-grained
 482 representations. Finally, when training with only the text loss (#2), we observe a noticeable degra-
 483 dation in linear probing performance, suggesting that VLM guidance alone is insufficient to learn strong
 484 fine-grained visual representations.

485 **Number of Clusters** In the part-assignment module, deep descriptors are grouped into K clusters.
 486 As shown in Table 3, setting $K = 2$ typically separates foreground from background but fails to
 487 capture richer part-level semantics, leading to suboptimal performance. Increasing K to 4 yields clear



488 Figure 5: Clustering behavior on CIFAR-100.

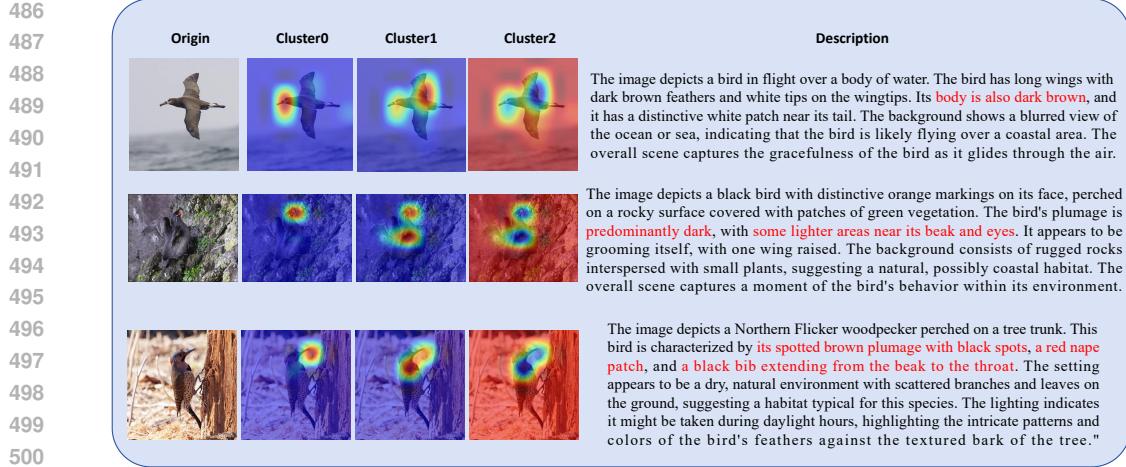


Figure 6: Visualization of the learned part clusters on fine-grained bird images. Each row shows the original image (left) and the response maps of three clusters (Cluster 0–2). Warmer colors indicate stronger activation of the corresponding cluster at that location. Cluster 0 consistently focuses on discriminative regions such as the head and beak, Cluster 1 attends to the torso/wing area, while Cluster 2 mainly captures background regions, illustrating that the part extractor discovers semantically meaningful parts.

improvements, indicating that moderate granularity helps model fine-grained cues. However, further enlarging K does not consistently provide additional benefits and can even degrade results due to overfitting. Overall, the effect of cluster count tends to plateau beyond a moderate value. Balancing discriminative power and computational cost, we fix $K = 3$ for all experiments in this paper.

Clustering Behavior of Module-Specific Clusters We evaluate clustering dynamics on CIFAR-100 using three granularity-aware metrics: (1) the average $\overline{\text{CDNV}_A}$ over the 20 coarse categories as an indicator of cluster compactness, (2) the average $\overline{\text{NCC}_{\text{fine}}}$ over the 100 fine classes to measure fine-grained discrimination, and (3) the average $\overline{\text{NCC}_{\text{coarse}}}$ over the 20 coarse categories to reflect super-class separability. An ideal outcome corresponds to a smaller $\overline{\text{CDNV}_A}$ together with higher $\overline{\text{NCC}_{\text{fine}}}$, while maintaining $\overline{\text{NCC}_{\text{coarse}}}$. As shown in Figure 5, our module effectively suppresses cluster collapse and enhances fine-grained discrimination, without sacrificing coarse-level separability. This confirms its role in shaping balanced feature geometry across semantic granularities.

The effect of the Part Extractor Our part extractor groups spatial features into clusters that correspond to semantically coherent regions in the image. To illustrate its behavior, we visualize several examples in Fig. 6. Even though no bounding-box annotations are used during training, the module automatically discovers meaningful parts: for instance, one cluster (Cluster 0) consistently focuses on the head and beak region, while another (Cluster 2) concentrates on the background. Similar patterns are observed across different categories, which aligns well with the common intuition in fine-grained recognition that stable, reusable parts are crucial for discrimination. In the visualizations, higher response (“hotter” colors) indicates stronger focus from the corresponding cluster on that region.

6 CONCLUSION

We studied the problem of fine-grained recognition in self-supervised learning, where standard objectives often suffer from over-dispersion or over-collapse of fine categories. To address this, we proposed a multi-level regularization framework that integrates soft-InfoNCE, part-aware learning, and VLM-guided alignment to shape feature geometry across granularities. Experiments on CUB200, Stanford Cars, and FGVC-Aircraft demonstrated consistent improvements in both classification and retrieval, with ablations confirming the complementary effect of each module. Our results highlight the importance of granularity-aware regularization for learning discriminative and transferable representations without labels.

540 ETHICS STATEMENT
541542 This work does not involve human subjects, personally identifiable information, or sensitive data. All
543 datasets used are publicly available, properly cited, and comply with their respective licenses. The
544 proposed methodology is designed for fine-grained self-supervised learning research and does not
545 introduce foreseeable risks of misuse or harmful societal impact. We have adhered to the ICLR Code
546 of Ethics throughout the research and preparation of this paper.
547548 REPRODUCIBILITY STATEMENT
549550 To facilitate reproducibility of CLUE, we provide comprehensive dataset descriptions and full experi-
551 mental details in in Appendix A, The supplementary materials include the complete source code,
552 training scripts, and step-by-step instructions, covering data preprocessing, model configurations, and
553 hyperparameters, enabling independent verification of all reported results.
554555 REFERENCES
556557 Relja Arandjelovic, Petr Gronat, Akihiko Torii, Tomas Pajdla, and Josef Sivic. Netvlad: Cnn
558 architecture for weakly supervised place recognition. In *Proceedings of the IEEE conference on*
559 *computer vision and pattern recognition*, pp. 5297–5307, 2016.560 Jihwan Bang, Sumyeong Ahn, and Jae-Gil Lee. Active prompt learning in vision language models.
561 In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.
562 27004–27014, 2024.564 Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. Beit: Bert pre-training of image transformers.
565 *arXiv preprint arXiv:2106.08254*, 2021.
566567 Adrien Bardes, Jean Ponce, and Yann LeCun. Vicreg: Variance-invariance-covariance regularization
568 for self-supervised learning. *arXiv preprint arXiv:2105.04906*, 2021.569 Ido Ben-Shaul, Ravid Shwartz-Ziv, Tomer Galanti, Shai Dekel, and Yann LeCun. Reverse engineering
570 self-supervised learning. *Advances in Neural Information Processing Systems*, 36:58324–58345,
571 2023.573 Qi Bi, Wei Ji, Jingjun Yi, Haolan Zhan, and Gui-Song Xia. Cross-level multi-instance distillation for
574 self-supervised fine-grained visual categorization. *IEEE Transactions on Image Processing*, 2025.576 Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin.
577 Unsupervised learning of visual features by contrasting cluster assignments. *Advances in neural*
578 *information processing systems*, 33:9912–9924, 2020.579 Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and
580 Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the*
581 *IEEE/CVF international conference on computer vision*, pp. 9650–9660, 2021.583 Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for
584 contrastive learning of visual representations. In *International conference on machine learning*, pp.
585 1597–1607. PMLR, 2020a.586 Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of*
587 *the IEEE/CVF conference on computer vision and pattern recognition*, pp. 15750–15758, 2021.589 Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum
590 contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020b.592 Elijah Cole, Xuan Yang, Kimberly Wilber, Oisin Mac Aodha, and Serge Belongie. When does
593 contrastive visual representation learning work? In *Proceedings of the IEEE/CVF conference on*
594 *computer vision and pattern recognition*, pp. 14755–14764, 2022.

594 Mohamed El Banani, Karan Desai, and Justin Johnson. Learning visual representations via language-
 595 guided sampling. In *Proceedings of the ieee/cvf conference on computer vision and pattern*
 596 *recognition*, pp. 19208–19220, 2023.

597

598 Tomer Galanti, András György, and Marcus Hutter. On the role of neural collapse in transfer learning.
 599 *arXiv preprint arXiv:2112.15121*, 2021.

600

601 Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena
 602 Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar,
 603 et al. Bootstrap your own latent-a new approach to self-supervised learning. *Advances in neural*
 604 *information processing systems*, 33:21271–21284, 2020.

605

606 Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Zero-shot detection via vision and language
 607 knowledge distillation. *arXiv preprint arXiv:2104.13921*, 2(3):4, 2021.

608

609 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image
 610 recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*,
 611 pp. 770–778, 2016.

612

613 Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for
 614 unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on*
 615 *computer vision and pattern recognition*, pp. 9729–9738, 2020.

616

617 Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked
 618 autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer*
 619 *vision and pattern recognition*, pp. 16000–16009, 2022.

620

621 Lang Huang, Shan You, Mingkai Zheng, Fei Wang, Chen Qian, and Toshihiko Yamasaki. Learning
 622 where to learn in cross-view self-supervised learning. In *Proceedings of the IEEE/CVF conference*
 623 *on computer vision and pattern recognition*, pp. 14451–14460, 2022.

624

625 Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung,
 626 Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with
 627 noisy text supervision. In *International conference on machine learning*, pp. 4904–4916. PMLR,
 2021.

628

629 Dong Jing, Xiaolong He, Yutian Luo, Nanyi Fei, Wei Wei, Huiwen Zhao, Zhiwu Lu, et al. Fineclip:
 630 Self-distilled region-based clip for better fine-grained understanding. *Advances in Neural Information*
 631 *Processing Systems*, 37:27896–27918, 2024.

632

633 Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron
 634 Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in neural*
 635 *information processing systems*, 33:18661–18673, 2020.

636

637 Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained
 638 categorization. In *Proceedings of the IEEE international conference on computer vision workshops*,
 639 pp. 554–561, 2013.

640

641 Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.

642

643 Chunyuan Li, Jianwei Yang, Pengchuan Zhang, Mei Gao, Bin Xiao, Xiyang Dai, Lu Yuan, and
 644 Jianfeng Gao. Efficient self-supervised vision transformers for representation learning. *arXiv*
 645 *preprint arXiv:2106.09785*, 2021.

646

647 Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained
 648 visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013.

649

650 Alžběta Manová, Aiden Durrant, and Georgios Leontidis. S-jea: Stacked joint embedding archi-
 651 tectures for self-supervised visual representation learning. *arXiv preprint arXiv:2305.11701*,
 652 2023.

653

654 Norman Mu, Alexander Kirillov, David Wagner, and Saining Xie. Slip: Self-supervision meets
 655 language-image pre-training. In *European conference on computer vision*, pp. 529–544. Springer,
 656 2022.

648 Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov,
 649 Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaeldin El-Nouby, et al. Dinov2: Learning
 650 robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023.

651

652 Wujian Peng, Sicheng Xie, Zuyao You, Shiyi Lan, and Zuxuan Wu. Synthesize diagnose and optimize:
 653 Towards fine-grained vision-language understanding. In *Proceedings of the IEEE/CVF Conference
 654 on Computer Vision and Pattern Recognition*, pp. 13279–13288, 2024.

655

656 Xiangyu Peng, Kai Wang, Zheng Zhu, Mang Wang, and Yang You. Crafting better contrastive views
 657 for siamese representation learning. In *Proceedings of the IEEE/CVF conference on computer
 658 vision and pattern recognition*, pp. 16031–16040, 2022.

659

660 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
 661 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual
 662 models from natural language supervision. In *International conference on machine learning*, pp.
 663 8748–8763. PMLR, 2021.

664

665 Oindrila Saha and Subhransu Maji. Particle: Part discovery and contrastive learning for fine-grained
 666 recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp.
 667 167–176, 2023.

668

669 Jialu Shi, Zhiqiang Wei, Jie Nie, and Lei Huang. Lodisc: Learning global-local discriminative
 670 features for self-supervised fine-grained visual recognition. *IEEE Transactions on Circuits and
 671 Systems for Video Technology*, 2025.

672

673 Aman Shrivastava, Ramprasaath R Selvaraju, Nikhil Naik, and Vicente Ordonez. Clip-lite: In-
 674 formation efficient visual representation learning with language supervision. *arXiv preprint
 675 arXiv:2112.07133*, 2021.

676

677 Yangyang Shu, Baosheng Yu, Haiming Xu, and Lingqiao Liu. Improving fine-grained visual
 678 recognition in low data regimes via self-boosting attention mechanism. In *European Conference
 679 on Computer Vision*, pp. 449–465. Springer, 2022.

680

681 Yangyang Shu, Anton Van den Hengel, and Lingqiao Liu. Learning common rationale to improve
 682 self-supervised representation for fine-grained visual recognition problems. In *Proceedings of the
 683 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11392–11401, 2023.

684

685 Oriane Siméoni, Huy V Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose,
 686 Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, et al. Dinov3. *arXiv
 687 preprint arXiv:2508.10104*, 2025.

688

689 Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *Advances
 690 in neural information processing systems*, 30, 2017.

691

692 Baofeng Tan, Xiu-Shen Wei, and Lin Zhao. Prototype-based contrastive learning with stage-wise pro-
 693 gressive augmentation for self-supervised fine-grained learning. In *Proceedings of the IEEE/CVF
 694 International Conference on Computer Vision*, pp. 4125–4134, 2025.

695

696 Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd
 697 birds-200-2011 dataset. 2011.

698

699 Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through align-
 700 ment and uniformity on the hypersphere. In *International conference on machine learning*, pp.
 701 9929–9939. PMLR, 2020.

702

703 Zihu Wang, Lingqiao Liu, Scott Ricardo Figueroa Weston, Samuel Tian, and Peng Li. On learning
 704 discriminative features from synthesized data for self-supervised fine-grained visual recognition.
 705 In *European Conference on Computer Vision*, pp. 101–117. Springer, 2024.

706

707 Xi Weng, Jianing An, Xudong Ma, Binhang Qi, Jie Luo, Xi Yang, Jin Song Dong, and Lei Huang.
 708 Clustering properties of self-supervised learning. *arXiv preprint arXiv:2501.18452*, 2025.

702 Minghao Xu, Yuanfan Guo, Xuanyu Zhu, Jiawen Li, Zhenbang Sun, Jian Tang, Yi Xu, and Bingbing
703 Ni. Hirl: A general framework for hierarchical image representation learning. *arXiv preprint*
704 *arXiv:2205.13159*, 2022.

705 Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised
706 learning via redundancy reduction. In *International conference on machine learning*, pp. 12310–
707 12320. PMLR, 2021.

709 Xin-yang Zhao, Jian Jin, Yang-yang Li, and Yazhou Yao. Twofold debiasing enhances fine-grained
710 learning with coarse labels. In *Proceedings of the AAAI Conference on Artificial Intelligence*,
711 volume 39, pp. 22831–22839, 2025.

712

713

714

715

716

717

718

719

720

721

722

723

724

725

726

727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

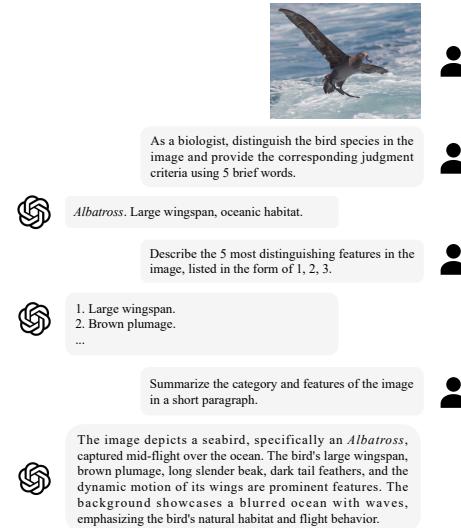
756 **A IMPLEMENTATION DETAILS**
757758 **A.1 DATASETS**
759760 We conducted experiments on widely used fine-grained recognition datasets to validate the effectiveness
761 of our algorithm. Specifically, CUB-200-2011 consists of 11,788 images spanning 200 bird
762 species, with 5,994 images allocated for training and 5,794 for testing. The Stanford Cars dataset
763 comprises 16,185 images across 196 categories, divided into 8,144 training images and 8,041 testing
764 images. FGVC-Aircraft includes 10,000 images covering 100 categories, with 6,667 images used for
765 training and 3,333 reserved for testing.
766767 For additional evaluation, we also consider four commonly used recognition benchmarks. The Oxford
768 102 Flowers dataset contains 8,189 images from 102 flower categories, following the standard split
769 into training, validation, and test sets. The Oxford-IIIT Pet dataset includes 7,349 images of 37 breeds
770 of cats and dogs, with roughly half of the images used for training and the remainder for testing.
771 Food-101 consists of 101,000 images from 101 food categories, with 750 training images and 250
772 test images per class. Caltech-256 contains 30,607 images spanning 256 object categories, where we
773 follow the conventional protocol and sample a fixed number of images per class for training while
774 using the remaining images for testing.
775776 **A.2 SETTINGS**
777778 For all our experiments, we use ResNet-50 (He et al.,
779 2016) as the backbone, initialized with ImageNet-1K
780 pre-trained weights, and employ a standard three-
781 layer MLP as the projector. **This initialization is kept**
782 **consistent with prior fine-grained SSL works such as**
783 **LCR**(Shu et al., 2023) and PAPN(Tan et al., 2025);
784 **we use the official TensorFlow ResNet-50 checkpoint**
785 **trained in a supervised manner on ImageNet-1K only.**
786787 Following common practice (He et al., 2020; Grill
788 et al., 2020; Caron et al., 2020), we adopt a momen-
789 tum encoder; **in our main experiments the momentum**
790 **coefficient is fixed to 0.999.** To ensure a fair com-
791 **parison, this choice is aligned with LCR**(Shu et al.,
792 2023) and LDF(Wang et al., 2024), which also use a
793 momentum value of 0.999. For reference, other self-
794 supervised methods such as ReSA(Weng et al., 2025)
795 and DINO(Caron et al., 2021) decrease the effective
796 update rate by gradually increasing the momentum
797 **from 0.996 to 1.0 over training.** Since clustering in
798 our framework is performed online, there is no need
799 for a memory bank to store negatives as in MoCo.
800 Instead, to ensure sufficient diversity within each
801 batch, we set the batch size to 512, which is still
802 much smaller than the queue size commonly used in
803 MoCo-style methods. The number of prototypes for part-level
804 clustering is set to $N = 3$.
805806 Based on the above settings, **the only tuned hyperparameters are the optimizer-related ones, including**
807 **the learning rate, weight decay, and the number of warm-up epochs, which are adjusted according to**
808 **the specific encoder architecture and dataset; all other settings are kept fixed.** We follow the linear
809 scaling rule, setting the learning rate as $lr = lr_{base} \times \text{batch_size}/256$. After the warm-up phase, the
learning rate decays following a cosine schedule.810 For the MAE baseline (He et al., 2022), we follow the official protocol, employing a ViT-Small
811 backbone with a masking ratio of 75% and pre-training for 100 epochs to ensure a fair comparison.
812 **ViT-Small is chosen because its model size is comparable to that of ResNet-50. Unless otherwise**
813 **specified, all methods share the same data preprocessing pipeline and optimization hyperparameters.**
814

Figure 7: An example of generating a discriminative image description through a chain-of-thought dialogue with a vision-language model.

810	811	812	813	814	815	816	817	818	819	820	821	822	823	824	825	826	827
Images	Descriptions																
	The image depicts a bird swimming in water. The bird has a predominantly brown plumage with lighter patches on its head and neck . It possesses a long, pointed beak typical of seabirds adapted for catching prey. Its wings are partially submerged , indicating it is actively swimming or diving. The water around the bird shows ripples and splashes, suggesting movement. The background consists of a calm body of water with gentle waves, providing a natural aquatic habitat for the bird.																
	The image depicts a vibrant red cardinal perched on a lichen-covered branch. The bird's striking plumage is characterized by its bright red feathers, black mask around the beak, and a prominent crest atop its head . The background is a soft green gradient, providing a natural contrast that highlights the cardinal's vivid coloration. The branch, adorned with patches of white lichen, adds texture to the scene, emphasizing the bird's delicate perch. This composition captures the essence of nature, showcasing the cardinal's beauty against a serene backdrop.																
	The image depicts a bird perched on a tree branch against a clear blue sky. The bird has a predominantly gray head with a black patch around its neck and a speckled pattern on its body . Its tail feathers are reddish-brown , adding a striking contrast to its overall appearance. The background is slightly blurred, highlighting the bird as the focal point. The branches surrounding the bird have sparse leaves, suggesting it might be autumn or early spring. The lighting indicates that the photo was taken during daylight hours, possibly under direct sunlight.																

828 Figure 8: Visualization of discriminative descriptions identified by a vision-language model (VLM)
 829 on the CUB-200 dataset. The highlighted descriptions in red denote the most discriminative textual
 830 cues associated with the corresponding bird images.

831
 832
 833 For a fair comparison with LoDisc (ViT-B) (Shi et al., 2025), we adopt a ViT-Base backbone in our
 834 ViT-based experiments, matching the model size used in LoDisc. Likewise, both ViT-Small and
 835 ViT-Base are initialized from the official TensorFlow checkpoints pre-trained in a supervised manner
 836 on ImageNet-1K.

837 All experiments are conducted on 4 NVIDIA RTX 3090 GPUs.
 838

839 B GENERATING DESCRIPTION WITH VLM

840 We explored using Vision-Language Models (VLMs) to generate discriminative descriptions. Inspired
 841 by the chain-of-thought approach, we adopted a multi-turn dialogue process for description generation.
 842 As illustrated in Figure 7, we first assign a specific role for each dataset, and then prompt the model
 843 from both global and local perspectives. Finally, a comprehensive textual description is generated for
 844 each image, ensuring appropriate length and detail.

845 After obtaining the generated feature descriptions, we utilize the text encoder of CLIP to project the
 846 textual information into the feature space. These projected text features serve as anchors, aiding the
 847 model's learning process.

848 Figure 8 presents example visualizations of discriminative descriptions generated by GLM-4V-flash
 849 on the CUB-200 dataset. Although GLM-4V-flash is not currently the most advanced vision-language
 850 model, it is nonetheless capable of effectively identifying and articulating distinctive image features.

851 C MORE EXPERIMENTS

852 To further validate the generality of CLUE, we additionally evaluate on three fine-grained benchmarks
 853 (Flowers, Pets, Food-101) and a more generic recognition benchmark (Caltech-256). For all these
 854 datasets, we strictly follow the training protocol described in the main text: the same ResNet-50
 855 backbone initialized from the ImageNet-1K supervised checkpoint, identical data augmentations,
 856 batch size, number of epochs, and optimizer settings. The only change is the underlying dataset. For
 857 LCR, we use the official implementation and re-train under this unified setup to obtain the numbers
 858 reported in Tables 4 and 5.

864 Table 4: Fine-grained classification accuracy (%) and retrieval performance (Rank-1 / Rank-5, %) on
 865 Flowers, Pets, and Food-101.

Method	Classification			Retrieval								
	Flower	Pet	Food	Flower	Rank1	Rank5	Pet	Rank1	Rank5	Food	Rank1	Rank5
LCR (Shu et al., 2023)	85.24	90.73	73.34	90.26	97.15	89.78	97.77	52.62	74.70			
ReSA (Weng et al., 2025)	84.17	91.82	72.91	93.89	97.92	89.24	97.87	56.08	75.99			
CLUE (Ours)	86.28	91.63	83.37	95.51	98.06	90.68	97.87	67.14	83.89			

872 Table 5: Fine-grained classification accuracy (%) and retrieval performance (Rank-1 / Rank-5, %) on
 873 Caltech-256.

Method	Classification		Retrieval	
	Rank1	Rank5	Rank1	Rank5
LCR (Shu et al., 2023)	74.44		75.82	87.85
ReSA (Weng et al., 2025)	84.07		78.36	88.04
CLUE (Ours)	86.39		80.22	90.90

881 The results show that CLUE consistently matches or outperforms strong baselines across all four
 882 datasets. On the fine-grained benchmarks in Table 4, CLUE achieves the best classification accuracy
 883 on Flowers (86.28%, +1.0 over LCR) and Food-101 (83.37%, more than +10 points over both LCR
 884 and ReSA), and also yields higher rank-1 retrieval scores on all three datasets, e.g., +14.5 points over
 885 LCR on Food-101. On Pets, CLUE attains competitive classification performance (91.63%) while
 886 still improving rank-1 retrieval over both LCR and ReSA. On the more generic Caltech-256 dataset
 887 (Table 5), CLUE significantly surpasses LCR (86.39% vs. 74.44% in classification, and 80.22% vs.
 888 75.82% in rank-1 retrieval) and also improves over ReSA. These results indicate that the proposed
 889 multi-level regularization not only benefits fine-grained recognition but also transfers well to broader
 890 object recognition scenarios.

D GEOMETRIC IDENTITIES AND OPTIMIZATION VIEW OF CLUE

894 It is widely acknowledged that providing a rigorous theoretical explanation of how regularization
 895 terms induce clustered representations is challenging; (Ben-Shaul et al., 2023) have emphasized this
 896 difficulty in their discussions of representation learning. Inspired by prior geometric analyses of
 897 contrastive and prototype-based methods (Wang & Isola, 2020; Khosla et al., 2020; Snell et al., 2017),
 898 we take a modest step in this direction. In this appendix, we present a simple geometric view of how
 899 the proposed multi-level regularization shapes the feature space. We show that our objectives tend to
 900 (i) reduce intra-class variance and (ii) enlarge inter-class separation, thereby improving the CNDV
 901 metric used in the main text.

902 Throughout, we assume that feature vectors are ℓ_2 -normalized, i.e., $\|z_i\|_2 = 1$ for all i , which is
 903 standard in contrastive learning.

D.1 INTRA-CLASS VARIANCE AND SIMILARITY

906 Let C be a set of indices (e.g., a class or a cluster) with $|C| = n$ and centroid

$$\mu = \frac{1}{n} \sum_{i \in C} z_i.$$

911 Define the intra-class variance

$$\text{Var}_{\text{in}}(C) = \frac{1}{n} \sum_{i \in C} \|z_i - \mu\|_2^2.$$

915 **Lemma 1** (Variance–similarity identity). *For unit-norm vectors, the intra-class variance is*

$$\text{Var}_{\text{in}}(C) = 1 - \text{Sim}_{\text{in}}(C), \quad \text{Sim}_{\text{in}}(C) = \frac{1}{n^2} \sum_{i,j \in C} \langle z_i, z_j \rangle. \quad (9)$$

918 *Proof.* Expanding the variance,
 919

$$920 \quad \text{Var}_{\text{in}}(C) = \frac{1}{n} \sum_{i \in C} (\|z_i\|_2^2 - 2\langle z_i, \mu \rangle + \|\mu\|_2^2).$$

923 Using $\|z_i\|_2^2 = 1$ and $\frac{1}{n} \sum_{i \in C} \langle z_i, \mu \rangle = \|\mu\|_2^2$ gives $\text{Var}_{\text{in}}(C) = 1 - \|\mu\|_2^2$. Finally,
 924

$$925 \quad \|\mu\|_2^2 = \left\langle \frac{1}{n} \sum_i z_i, \frac{1}{n} \sum_j z_j \right\rangle = \frac{1}{n^2} \sum_{i,j \in C} \langle z_i, z_j \rangle = \text{Sim}_{\text{in}}(C).$$

□

930 Thus, on the unit sphere, maximizing intra-class similarity $\text{Sim}_{\text{in}}(C)$ is exactly equivalent to mini-
 931 mizing intra-class variance $\text{Var}_{\text{in}}(C)$.
 932

933 D.2 CLASS SEPARATION AND INTER-CLASS SIMILARITY

935 Consider two disjoint sets C and C' with $|C| = n$, $|C'| = n'$ and centroids

$$936 \quad \mu = \frac{1}{n} \sum_{i \in C} z_i, \quad \mu' = \frac{1}{n'} \sum_{j \in C'} z_j.$$

939 Define the inter-class similarity
 940

$$941 \quad \text{Sim}_{\text{out}}(C, C') = \frac{1}{nn'} \sum_{i \in C, j \in C'} \langle z_i, z_j \rangle.$$

944 **Lemma 2** (Centroid distance decomposition). *The squared distance between centroids satisfies*

$$945 \quad \|\mu - \mu'\|_2^2 = \text{Sim}_{\text{in}}(C) + \text{Sim}_{\text{in}}(C') - 2 \text{Sim}_{\text{out}}(C, C'). \quad (10)$$

947 *Proof.* We have
 948

$$949 \quad \|\mu - \mu'\|_2^2 = \|\mu\|_2^2 + \|\mu'\|_2^2 - 2\langle \mu, \mu' \rangle.$$

950 Using $\|\mu\|_2^2 = \frac{1}{n^2} \sum_{i,j \in C} \langle z_i, z_j \rangle = \text{Sim}_{\text{in}}(C)$, similarly $\|\mu'\|_2^2 = \text{Sim}_{\text{in}}(C')$, and
 951

$$952 \quad \langle \mu, \mu' \rangle = \frac{1}{nn'} \sum_{i \in C, j \in C'} \langle z_i, z_j \rangle = \text{Sim}_{\text{out}}(C, C'),$$

954 we obtain equation 10. □
 955

956 This shows that increasing intra-class similarities $\text{Sim}_{\text{in}}(C)$, $\text{Sim}_{\text{in}}(C')$ and decreasing inter-class
 957 similarity $\text{Sim}_{\text{out}}(C, C')$ both enlarge the centroid distance.
 958

959 D.3 CONNECTION TO CNDV

961 For two sets C and C' with centroids μ, μ' , the pairwise CNDV metric (Eq. (1) in the main text) is
 962

$$963 \quad \text{CNDV}(C, C') = \frac{\text{Var}_{\text{in}}(C) + \text{Var}_{\text{in}}(C')}{2 \|\mu - \mu'\|_2^2}. \quad (11)$$

965 Assuming $\mu \neq \mu'$, Lemma 1 and Lemma 2 give

$$966 \quad \text{CNDV}(C, C') = \frac{(1 - \text{Sim}_{\text{in}}(C)) + (1 - \text{Sim}_{\text{in}}(C'))}{2(\text{Sim}_{\text{in}}(C) + \text{Sim}_{\text{in}}(C') - 2 \text{Sim}_{\text{out}}(C, C'))}. \quad (12)$$

971 Thus, strictly increasing $\text{Sim}_{\text{in}}(C)$, $\text{Sim}_{\text{in}}(C')$ and strictly decreasing $\text{Sim}_{\text{out}}(C, C')$ decreases the
 972 numerator and increases the denominator, and therefore strictly reduces $\text{CNDV}(C, C')$.

972 D.4 OPTIMIZATION DYNAMICS OF SOFT-INFO NCE
973974 We now show how the Soft-InfoNCE loss used in CLUE acts on the pairwise similarities in a way
975 that matches the CDNV analysis above.976 Let $S \in \mathbb{R}^{m \times m}$ be a similarity matrix and define a row-wise softmax (for simplicity, without explicit
977 temperature) by

978
$$D_{ij} = \frac{\exp(S_{ij})}{\sum_{k=1}^m \exp(S_{ik})}.$$

980 Let $A \in \mathbb{R}^{m \times m}$ be a row-stochastic target matrix ($\sum_j A_{ij} = 1$ for all i), and consider

982
$$L(S, A) = -\frac{1}{m} \sum_{i=1}^m \sum_{j=1}^m A_{ij} \log D_{ij}. \quad (13)$$

985 **Lemma 3** (Gradient of row-softmax cross-entropy). *For the loss $L(S, A)$ in equation 13,*

987
$$\frac{\partial L}{\partial S_{ij}} = \frac{1}{m} (D_{ij} - A_{ij}), \quad i, j = 1, \dots, m. \quad (14)$$

990 *Proof.* Writing $L = \frac{1}{m} \sum_i L_i$ with $L_i = -\sum_j A_{ij} \log D_{ij}$ and using

992
$$\log D_{ij} = S_{ij} - \log \left(\sum_k \exp(S_{ik}) \right),$$

994 a standard calculation shows $\frac{\partial}{\partial S_{ij}} \log D_{ik} = \delta_{jk} - D_{ij}$, hence $\frac{\partial L_i}{\partial S_{ij}} = D_{ij} - A_{ij}$ and $\frac{\partial L}{\partial S_{ij}} = \frac{1}{m} (D_{ij} - A_{ij})$. \square 997 **Corollary 4** (Effect on pairwise similarities). *A gradient descent step $S_{ij} \leftarrow S_{ij} - \eta \partial L / \partial S_{ij}$ with
998 $\eta > 0$ has:*1000

- if $A_{ij} > D_{ij}$, then $\partial L / \partial S_{ij} < 0$ and S_{ij} increases (the pair (i, j) is pulled closer);
- if $A_{ij} < D_{ij}$, then $\partial L / \partial S_{ij} > 0$ and S_{ij} decreases (the pair (i, j) is pushed apart).

1003 In CLUE, A is obtained from a Sinkhorn-based clustering and assigns larger mass to semantically
1004 related examples. Therefore the global and local regularization terms:1005

- increase similarities for pairs that should be close (increasing Sim_{in} and reducing the CDNV
1006 numerator);
- decrease similarities for unrelated pairs (reducing Sim_{out} and increasing the CDNV denominator).

1010 By Lemma 1 and Lemma 2, this means that the global and local losses tend to reduce intra-class
1011 variance and enlarge inter-class distances, thus decreasing CDNV.1013 D.5 PREVENTION OF TRIVIAL COLLAPSE VIA TEXT GUIDANCE
10141015 Finally, we show that a CLIP-style text alignment loss rules out the trivial solution where all image
1016 features collapse to the same constant vector.1017 Let v_i be (normalized) image features and t_i the corresponding text features. Consider the CLIP-style
1018 loss

1019
$$L_{\text{text}} = -\frac{1}{m} \sum_{i=1}^m \log \frac{\exp(\langle v_i, t_i \rangle / \tau)}{\sum_{j=1}^m \exp(\langle v_i, t_j \rangle / \tau)}, \quad (15)$$

1022 with temperature $\tau > 0$, and let

1023
$$p_{ij} = \frac{\exp(\langle v_i, t_j \rangle / \tau)}{\sum_{k=1}^m \exp(\langle v_i, t_k \rangle / \tau)}$$

1024 be the softmax probabilities.

1026
1027 **Lemma 5** (Constant mapping is not stationary). *Assume that the text features $\{t_i\}_{i=1}^m$ are not all
1028 equal. Then the mapping $v_1 = \dots = v_m = c$ (for any constant c) is not a stationary point of L_{text} .*

1029 *Proof.* The gradient w.r.t. v_i is

$$1031 \quad \frac{\partial L_{\text{text}}}{\partial v_i} = \frac{1}{m\tau} \left(\sum_{j=1}^m p_{ij} t_j - t_i \right).$$

1034 If $v_1 = \dots = v_m = c$, then p_{ij} does not depend on i , say $p_{ij} = p_j$, and

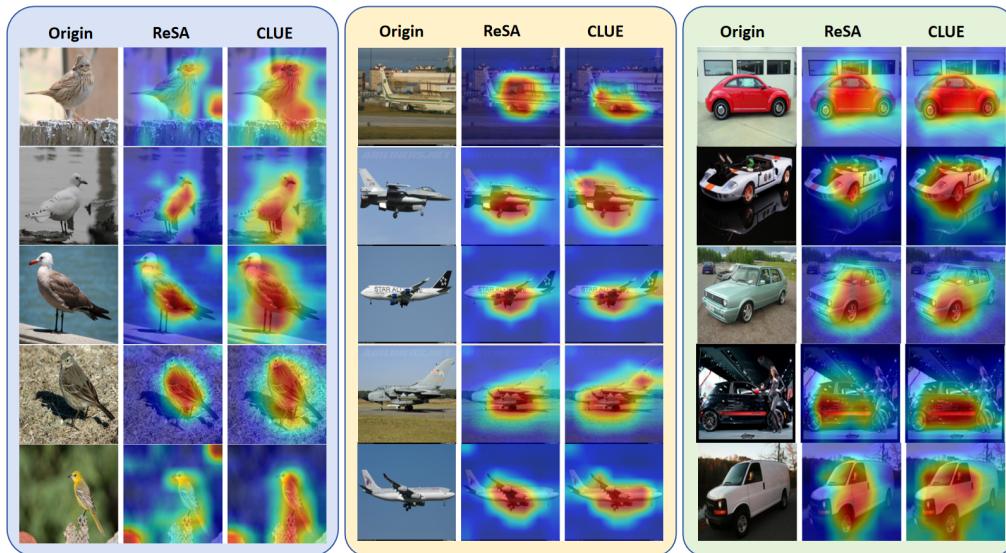
$$1036 \quad \frac{\partial L_{\text{text}}}{\partial v_i} = \frac{1}{m\tau} \left(\sum_j p_j t_j - t_i \right).$$

1038 Stationarity would require $\sum_j p_j t_j = t_i$ for all i , i.e., all t_i equal the same convex combination of
1039 $\{t_j\}$, which is impossible if the $\{t_i\}$ are not all identical. Hence at least one gradient is nonzero. \square

1041 Therefore, the text-based regularization term explicitly rules out the trivial “all features are identical”
1042 solution and helps prevent extreme instance-level collapse, complementing the clustering-based
1043 global and local regularization discussed above.

1045 E VISUALIZATION

1046 For a clearer illustration of the model’s effectiveness, we conduct Grad-CAM visualizations on the
1047 evaluation datasets, highlighting the regions most relevant to fine-grained discrimination 9.



1069 Figure 9: Grad-CAM visualizations on fine-grained benchmarks. Compared with baseline SSL,
1070 our CLUE model attends more accurately to discriminative regions (e.g., textures, shapes, or parts),
1071 which facilitates precise recognition and highlights its advantage in fine-grained categorization.

1073 F THE USE OF LARGE LANGUAGE MODELS(LLMS)

1076 Large Language Models (LLMs) were used only as language and formatting assistants during
1077 manuscript preparation. Specifically, LLMs were employed to (i) polish grammar and improve
1078 fluency, (ii) standardize terminology, tense, and voice, (iii) suggest alternative phrasings for clarity
1079 and concision, and (iv) provide suggestions for table layouts and LaTeX typesetting (e.g., caption
style, column alignment, and cross-referencing).

1080 Within our proposed framework, LLMs are included solely as supportive modules, functioning as
1081 tools to facilitate the architecture rather than as core research contributions.

1082 LLMs did not participate in designing experiments, analyzing data, deriving theoretical results, or
1083 drawing conclusions. All technical ideas, methods, proofs, experimental protocols, and findings
1084 are authored, validated, and interpreted by the authors. All LLM-assisted edits were reviewed and
1085 approved by the authors to ensure accuracy and faithfulness to the intended meaning.

1086

1087

1088

1089

1090

1091

1092

1093

1094

1095

1096

1097

1098

1099

1100

1101

1102

1103

1104

1105

1106

1107

1108

1109

1110

1111

1112

1113

1114

1115

1116

1117

1118

1119

1120

1121

1122

1123

1124

1125

1126

1127

1128

1129

1130

1131

1132

1133