

Let the Void Be Void: Robust Open-Set Semi-Supervised Learning via Selective Non-Alignment

You Rim Choi¹, Subeom Park¹, Seojun Heo¹, Eunchung Noh², Hyung-Sin Kim¹

¹Graduate School of Data Science, Seoul National University, Seoul, South Korea

²Samsung Electronics

yrchoi@snu.ac.kr, sbpark7@snu.ac.kr, seodalzzz@snu.ac.kr, eunchung.noh@samsung.com, hyungkim@snu.ac.kr

Abstract

Open-set semi-supervised learning (OSSL) leverages unlabeled data containing both in-distribution (ID) and unknown out-of-distribution (OOD) samples, aiming simultaneously to improve closed-set accuracy and detect novel OOD instances. Existing methods either discard valuable information from uncertain samples or force-align every unlabeled sample into one or a few synthetic “catch-all” representations, resulting in geometric collapse and overconfidence on only seen OODs. To address the limitations, we introduce *selective non-alignment*, adding a novel “skip” operator into conventional pull and push operations of contrastive learning. Our framework, *SkipAlign*, selectively skips alignment (pulling) for low-confidence unlabeled samples, retaining only gentle repulsion against ID prototypes. This approach transforms uncertain samples into a pure repulsion signal, resulting in tighter ID clusters and naturally dispersed OOD features. Extensive experiments demonstrate that *SkipAlign* significantly outperforms state-of-the-art methods in detecting unseen OOD data without sacrificing ID classification accuracy.

Introduction

Semi-supervised learning (SSL) improves performance by leveraging unlabeled data, typically assuming that *all* unlabeled samples belong strictly to one of the known in-distribution (ID) classes (Sohn et al. 2020). However, practical unlabeled datasets inevitably form **open sets**, containing unknown out-of-distribution (OOD) samples. To address this limitation, Open-Set Semi-Supervised Learning (OSSL) learns from open-set unlabeled data to simultaneously improve ID classification and reliably detect OOD samples, including OOD classes present in the unlabeled pool (*seen* OOD) and novel OOD classes that appear only at test time (*unseen* OOD).

However, generalizable OOD detection via OSSL remains challenging. This is primarily because **the true OOD space is inherently vast and heterogeneous, whereas training provides only limited, potentially misleading glimpses of OOD samples**. Recent state-of-the-art methods take opposite stances toward OOD: aggressive or conservative. On the aggressive side, SCOMatch (Wang et al. 2024) collapses the diverse OOD universe into a single synthetic

class, causing *geometric collapse and overconfidence on seen OOD samples*. Similarly, ProSub (Wallin et al. 2024b) creates a broad ID subspace spanning all ID clusters, aggressively pulling or pushing every unlabeled sample toward or away from this subspace even when pseudo-labels are *noisy*. Consequently, unseen OOD samples that *fall within the broad ID subspace* but outside any specific ID cluster tend to be misclassified as ID. On the conservative side, SSB (Fan et al. 2023) trains per-class one-vs-all (OVA) classifiers but discards low-confidence unlabeled data, *losing valuable geometric cues* that could improve feature space discrimination between ID and novel OOD.

To address these limitations, we conceptualize the embedding space as a universe consisting of compact **galaxies** (ID classes) surrounded by a vast **interstellar void** (OOD). Guided by this analogy, we propose *Selective Non-Alignment (SNA)*, a novel contrastive learning principle that adds a third operator—“**skip**”—alongside traditional pull-push pair. SNA pursues two objectives: (1) preserve the natural heterogeneity of OOD samples, and (2) carefully guide uncertain unlabeled samples to minimize the risk of amplifying ID/OOD pseudo-labeling errors.

Specifically, an unlabeled sample is pulled into an ID galaxy only when it passes a **stringent confidence test**, significantly reducing the risk of mistakenly aligning an OOD sample. Otherwise, unlike previous prototype-based contrastive methods that attract every sample toward a prototype (Li et al. 2020), SNA **nullifies the pull term** and applies only a *non-selective angular update* computed from a softmax-weighted average of all ID prototype directions. This ensures that points closer to ID galaxies experience stronger repulsion, preventing premature attachment to any galaxy. Over time, genuine ID samples in the void naturally gravitate toward their correct galaxies as the model matures, whereas genuine OODs remain securely dispersed. This produces a distinctive geometric structure for robust detection of unseen OODs: **dense ID galaxies clearly separated by a diffuse void**. Figure 1 illustrates the concept of SNA.

We instantiate SNA in the *SkipAlign* framework, integrating it alongside an ID classifier and an OVA OOD detector sharing a unified backbone. To ensure representational independence and flexibility, *SkipAlign* employs an additional non-linear projection layer before applying SNA. During training, SNA synergistically interacts with the ID classifier

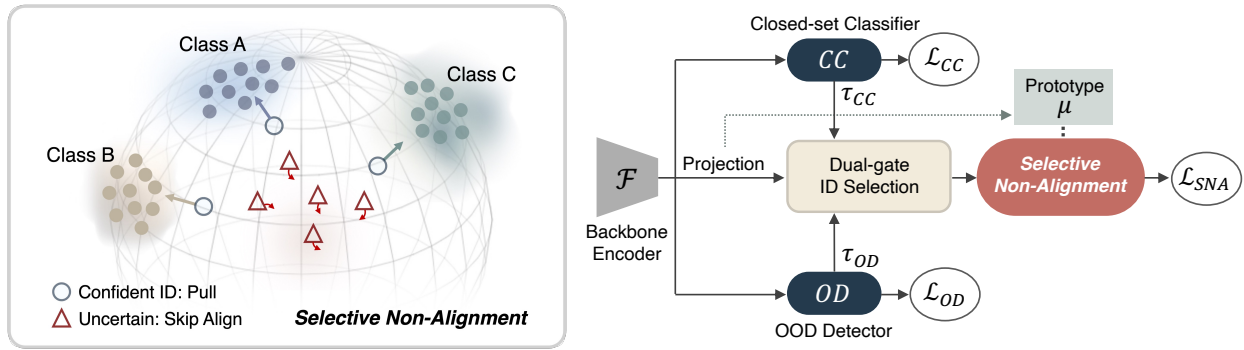


Figure 1: Overview of *SkipAlign*. Left: Illustration of the Selective Non-Alignment (SNA) concept. Confident ID samples are attracted toward their corresponding class prototypes, while uncertain or OOD samples are softly repelled and remain unaligned. Right: The overall framework of *SkipAlign*. A shared encoder feeds into a closed-set classifier ($CC(\cdot)$) and an OOD detector ($OD(\cdot)$), which jointly perform dual-gate ID selection. Projection embeddings are used to form class prototypes, enabling SNA to refine representation learning through prototype-based alignment and repulsion.

and OOD detector in two ways: (1) For the stringent confidence criterion of SNA, *SkipAlign* utilizes a dual-gate mechanism, requiring simultaneous high-confidence agreement from both the ID classifier and OOD detector, and (2) SNA progressively refines the shared backbone to enhance performance on both downstream tasks. Furthermore, *SkipAlign* progressively refines prototypes by integrating confidently pseudo-labeled unlabeled samples into class centroids, stabilizing training.

Our contributions can be summarized as follows:

- **Selective non-alignment principle:** We introduce *skip* as an essential third contrastive operator, effectively preventing OOD overfitting by selectively suppressing alignment.
- **Tailored framework integration:** *SkipAlign* seamlessly integrates SNA with a dedicated projection layer, dual-gate confidence checks, and progressive prototype refinement to robustly leverage both labeled and unlabeled data.
- **Extensive empirical validation:** Extensive evaluations on CIFAR-10/100, ImageNet-30, and TinyImageNet, and diverse unseen OOD datasets demonstrate *SkipAlign*'s superior generalization, improving Overall-OOD AUC by +3.1 points on average (max +7.1) over the strongest baseline, while maintaining comparable closed-set accuracy.

Related Work

Open-set Semi-Supervised Learning (OSSL)

Early approaches to OSSL identify and exclude OOD samples, thus minimizing their negative impact on ID classifiers (Yu et al. 2020; Guo et al. 2020; Saito, Kim, and Saenko 2021). Recent methods, however, have shifted toward actively incorporating OOD samples. To leverage representations from OOD data, auxiliary self-supervised tasks, such as orientation prediction (Huang et al. 2021) or feature consistency enforcement (Wallin et al. 2024a), have been explored. However, these general-purpose tasks do not directly target OOD detection, resulting in limited gains.

More recent, targeted OSSL approaches handle OOD samples either aggressively or conservatively. On the

aggressive side, IOMatch (Li et al. 2023) and SCOMatch (Wang et al. 2024) introduce OSSL-specific modules using a $(K + 1)$ -way classifier, assigning a single additional class for all OOD samples. While SCOMatch improves over IOMatch by enhancing pseudo-label reliability, collapsing diverse OOD samples into one homogeneous class disrupts the embedding space geometry and induces overfitting to seen OOD samples, thereby hindering generalization to novel OOD data. Similarly, ProSub (Wallin et al. 2024b) defines an ID subspace spanning class prototypes and actively repels OOD samples from the subspace based on angular distances, establishing clear boundaries for known OODs. Despite being less rigid than the single-class strategy, ProSub's subspace assumption can inadvertently classify unseen OOD samples located within the ID subspace but outside specific ID prototypes as ID, leading to detection errors. Furthermore, these aggressive methods consistently align unlabeled samples to prototypes or classes regardless of pseudo-label noise, compromising performance due to incorrect alignments.

On the conservative side, SSB (Fan et al. 2023) employs one-vs-all (OVA) binary classifiers for each ID class—partitioning samples into “ONE” (target ID class) and “ALL” (other classes) categories—with the OOD region defined as the intersection of all “ALL” regions. During training, the “ONE” class exclusively uses labeled samples, while the “ALL” class additionally includes only a limited subset of highly confident OOD samples. Although this conservative approach reduces the risk of erroneous alignment, discarding low-confidence unlabeled data sacrifices valuable geometric cues essential for distinguishing ID from novel OOD.

Therefore, achieving robust and generalizable OOD detection through OSSL remains an open challenge, demanding (1) a tailored characterization of the OOD space and (2) a nuanced approach to handling uncertain unlabeled samples.

SkipAlign

We propose *SkipAlign* to improve generalizable OOD detection in OSSL by addressing the key limitations of existing

methods. Our core innovation for OSSL is **Selective Non-Alignment (SNA)**, characterized by two complementary strategies. First, SNA forms compact, individual ID clusters (no subspace spanning), selectively pulling only high-confidence ID samples toward their respective class prototypes. Second, for all other unlabeled samples—including confidently identified OOD and uncertain samples—SNA completely removes the pull operation, applying instead a non-selective angular repulsion. This has two benefits: (1) By refraining from pulling confident OOD samples to specific regions, *SkipAlign* preserves the inherent diversity and dispersion of the OOD interstellar void; (2) Mildly pushing uncertain samples from all prototypes provides robust contrastive learning that effectively mitigates sensitivity to pseudo-label noise.

Overview of *SkipAlign*

For K -way classification, let $X = \{(x_i, y_i)\}_{i=1}^B$ be a batch of B labeled samples from \mathcal{D}_l with labels $y_i \in \{1, \dots, K\}$. Let $\mathcal{U} = \{u_i\}_{i=1}^{\gamma B}$ be γB unlabeled samples from \mathcal{D}_u , where γ controls the labeled-to-unlabeled ratio. Unlike SSL, OSSL includes *OOD samples* from unknown classes in \mathcal{D}_u .

As illustrated in Figure 1, *SkipAlign* consists of a shared backbone and three task-specific heads. These heads share representations and mutually reinforce learning—particularly the SNA module aiding both classification and OOD detection—while MLP projection heads are inserted to reduce harmful gradient interference, preserving constructive information flow across tasks. Each head is optimized for a specific objective:

- **Closed-set (ID) classifier (CC)** performs standard K -way classification over ID classes, outputting a softmax probability vector $\mathbf{p} \in \mathbb{R}^K$.
- **OOD detector (OD)** distinguishes ID from OOD samples using an OVA classifier with K binary sub-classifiers. Each detector $OD_k(\cdot)$ outputs a probability pair $(\varphi_k^{ID}, \varphi_k^{OOD})$ indicating the likelihood of being ID or OOD for class k .
- **Selective Non-Alignment (SNA) module**, the core of *SkipAlign*, shapes the feature space to improve both ID classification and OOD detection, facilitating better generalization to unseen OODs. It operates on the embedding $\mathbf{z} \in \mathbb{R}^d$ produced by a dedicated projection head on top of the shared backbone.

Selective Non-Alignment

Unlabeled data in OSSL inevitably include samples from unknown classes, introducing ambiguity in their assignment to ID prototypes. To address this, the SNA module attracts only those samples confidently predicted as ID by both the closed-set classifier and the OOD detector—a mechanism we term **dual-gate ID selection**. For each unlabeled embedding z_i , its confidence mask Φ_i is defined as:

$$\Phi_i = \mathbb{1}(p_{i, \hat{k}} > \tau_{ID} \wedge \varphi_{i, \hat{k}}^{ID} > \eta_{ID}), \quad \hat{k} = \arg \max_k p_{i, k} \quad (1)$$

where τ_{ID} and η_{ID} are thresholds for the closed-set classifier and the OOD detector, respectively.

Let $\text{sim}(a, b) = \frac{a}{\|a\|} \cdot \frac{b}{\|b\|}$ denote the cosine similarity between two vectors. For an unlabeled embedding z_i with predicted class \hat{k} and its prototype $\mu_{\hat{k}} \in \mathbb{R}^d$, the unlabeled SNA loss is:

$$\mathcal{L}_{\text{USNA}}(z_i) = -\Phi_i \frac{\text{sim}(z_i, \mu_{\hat{k}})}{T} + \log \sum_{j=1}^K \exp\left(\frac{\text{sim}(z_i, \mu_j)}{T}\right) \quad (2)$$

where T is a temperature parameter.

The gradient of $\mathcal{L}_{\text{USNA}}$ with respect to z_i is:

$$\nabla_{z_i} \mathcal{L}_{\text{USNA}} = \frac{1}{T \|z_i\|} (I - \hat{z} \hat{z}^\top) \left(\sum_{j=1}^K \alpha_j \hat{\mu}_j - \Phi_i \hat{\mu}_{\hat{k}} \right) \quad (3)$$

where $\hat{z} = z_i / \|z_i\|$, $\hat{\mu}_j = \mu_j / \|\mu_j\|$, and α_j is the softmax probability of z_i belonging to class j ($\frac{\exp(\text{sim}(z_i, \mu_j)/T)}{\sum_{\ell} \exp(\text{sim}(z_i, \mu_\ell)/T)}$). The projection term $(I - \hat{z} \hat{z}^\top)$ removes radial components, ensuring that updates act purely in the angular direction.

This design selectively aligns only confident ID samples ($\Phi_i = 1$) toward their respective class prototypes. In contrast, uncertain or OOD samples ($\Phi_i = 0$) receive **no targeted pull and instead undergo a non-selective angular update** determined by the averaged term $\sum_j \alpha_j \hat{\mu}_j$, preventing unwanted alignment to any specific prototype. Complementing the *CC* and *OD* heads, SNA addresses a key limitation of their objectives (e.g., cross-entropy or pseudo-label consistency), which may unintentionally enlarge feature norms even for incorrectly pseudo-labeled samples. SNA’s stricter dual-gate mechanism classifies these ambiguous samples as uncertain ($\Phi_i = 0$). The corresponding $\mathcal{L}_{\text{USNA}}$ then provides an orthogonal corrective “torque”, redirecting uncertain features away from incorrectly matched prototypes without aligning them elsewhere. Consequently, norm growth occurs exclusively along reliable directions, effectively mitigating the risk of uncertain samples from accumulating large norms in misleading regions. Jointly, SNA, *CC*, and *OD*, yield a more discriminative and robust feature space, suitable for challenging open-set scenarios.

For labeled samples, we exploit reliable ground-truth supervision to promote intra-class compactness and implicit inter-class separation. The Instance-wise Alignment loss (\mathcal{L}_{IA}), following the supervised contrastive formulation (Khosla et al. 2020), attracts same-class samples:

$$\mathcal{L}_{\text{IA}}(z_i) = -\frac{1}{|P(i)|} \sum_{p \in P(i)} \log \frac{\exp(\text{sim}(z_i, z_p)/T)}{\sum_{j=1}^B \mathbb{1}_{i \neq j} \exp(z_i \cdot z_j/T)} \quad (4)$$

where $P(i)$ is the set of same-class samples excluding i , and T is the temperature. Prototype-based Alignment loss (\mathcal{L}_{PA}) always pulls each embedding toward its class prototype:

$$\mathcal{L}_{\text{PA}}(z_i) = -\frac{\text{sim}(z_i, \mu_{y_i})}{T} + \log \sum_{j=1}^K \exp\left(\frac{\text{sim}(z_i, \mu_j)}{T}\right) \quad (5)$$

where μ_{y_i} denotes the prototype of class y_i .

The overall SNA loss combines labeled and unlabeled objectives:

$$\mathcal{L}_{\text{SNA}} = \lambda_{\text{USNA}} \mathcal{L}_{\text{USNA}}(\mathcal{U}) + \lambda_{\text{IA}} \mathcal{L}_{\text{IA}}(X) + \lambda_{\text{PA}} \mathcal{L}_{\text{PA}}(X), \quad (6)$$

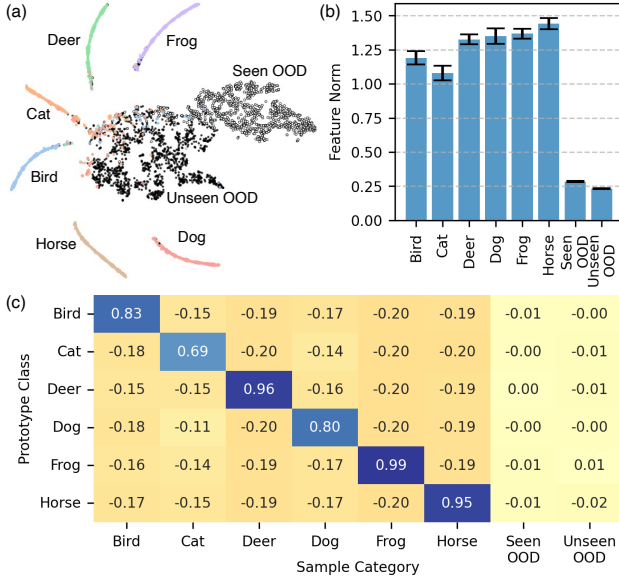


Figure 2: Effect of the SNA module. (a) t-SNE visualization of projection embeddings \mathbf{z} . (b) Average feature norm for each class category. (c) Average cosine similarity between sample categories and ID class prototypes.

where λ_{USNA} , λ_{IA} , λ_{PA} weight the contribution of each term.

Figure 2 provides a visual illustration of the geometric effects induced by our SNA module. This figure is generated under a CIFAR-10 OSSL setting, where six animal classes are treated as inliers with 50 labeled samples per class, four classes as seen OOD, and **six additional datasets as unseen OOD** (see Section Experiments for details). Figure 2(a) shows a t-SNE plot of projection embeddings \mathbf{z} . As intended, ID samples form distinct, compact clusters with smoothly curved manifolds, demonstrating clear alignment with prototype direction. In contrast, both seen and unseen OOD samples remain in the surrounding void without attaching to any ID cluster, reflecting the effect of angular updates from SNA. Some uncertain ID samples, excluded by our dual-gate criterion, lie near the ID clusters but partially intermingle with nearby OOD embeddings.

To further assess discriminability, we analyze feature norms and prototype alignment. Figure 2(b) reveals a clear norm gap between ground-truth ID and OOD samples. This separation reflects how USNA encourages reliable samples to align with prototypes—thereby accumulating norm through the synergistic pull of CC , OD , and SNA alignment—while applying a selective repulsion to misaligned or unlabeled samples. As a result, ID samples exhibit significantly higher norms than both seen and unseen OODs. Figure 2(c) illustrates the average cosine similarity between each sample category and the ID prototypes. ID samples exhibit high similarity with their respective prototypes, whereas both seen and unseen OODs consistently display low similarity to all prototypes. Collectively, SNA, together with CC and OD , naturally shapes a geometric structure consisting of dense, high-norm ID “galaxies” clearly sepa-




Figure 3: Adaptive Prototype Generation. For each class, the labeled prototype μ_l is the mean embedding of labeled samples, while the unlabeled prototype μ_u is obtained from confident ID samples among the unlabeled data. The final prototype μ is computed as a weighted sum of the two.

rated by a sparse, low-norm OOD “interstellar void”, effectively generalizing to unseen OOD distributions.

Prototype Refinement

In SNA, prototypes serve as positive anchors that guide ID embeddings toward class centers. However, relying solely on limited labeled data can lead to suboptimal prototypes. To mitigate this, we dynamically refine prototypes by combining labeled features with reliably pseudo-labeled unlabeled features selected via a stricter dual-gate criterion. Figure 3 illustrates this process, where the final prototype is computed as a weighted average of labeled and selected unlabeled embeddings. The weights adaptively account for sample counts and a hyperparameter controlling the influence of unlabeled data, as detailed in Supplementary Material. This refinement stabilizes prototype representations and improves class separability throughout training.

Objectives

The objective functions used to jointly train the closed-set classifier, the OOD detector, and the SNA module are described below.

Closed-set classifier. We adopt the FixMatch (Sohn et al. 2020) framework to train the closed-set classifier using labeled and unlabeled data. The loss consists of a supervised cross-entropy loss \mathcal{L}_x on labeled samples and a consistency regularization loss \mathcal{L}_u based on unlabeled samples pseudo-labeled via softmax predictions:

$$\mathcal{L}_{CC} = \mathcal{L}_x(X) + \lambda_u \mathcal{L}_u(\mathcal{U}) \quad (7)$$

The balancing weight λ_u controls the relative strength of the unsupervised term.

OOD detector. We follow the loss formulation of OpenMatch (Saito, Kim, and Saenko 2021) and SSB (Fan et al. 2023) to train the OOD detector. The total loss includes the One-vs-All (OVA) loss \mathcal{L}_{od}^{OVA} for labeled data, an entropy minimization loss \mathcal{L}_{od}^{em} to sharpen predictions on unlabeled samples, and a soft open-set consistency regularization loss \mathcal{L}_{od}^{SOCR} that promotes prediction consistency under augmentation. In addition, we incorporate a pseudo-negative loss \mathcal{L}_{od}^{neg} , which leverages confidently identified OOD samples from \mathcal{U} to enhance ID-OOD separation. To further improve robustness, both weak and strong augmentations are applied

Dataset	CIFAR-10 6/4				CIFAR-100 55/45				CIFAR-100 80/20				ImageNet30		TIN	
	25		50		25		50		25		50		5%		5%	
Metric	Acc.	AUC	Acc.	AUC	Acc.	AUC	Acc.	AUC	Acc.	AUC	Acc.	AUC	Acc.	AUC	Acc.	AUC
IOMatch	93.2	83.7	93.3	84.9	69.5	73.6	73.2	76.6	<u>64.5</u>	70.0	68.6	71.6	84.5	71.4	57.9	69.0
	± 0.6	± 0.6	± 0.4	± 1.0	± 0.5	± 0.2	± 0.0	± 2.0	± 0.5	± 1.2	± 0.1	± 1.4	± 0.3	± 2.4	± 0.3	± 1.6
SSB	84.5	75.5	91.0	<u>93.5</u>	69.4	80.8	74.0	83.4	64.3	<u>75.7</u>	69.8	75.3	91.8	82.0	59.9	73.8
	± 3.9	± 4.5	± 1.3	± 0.5	± 0.8	± 0.7	± 0.5	± 1.7	± 0.4	± 2.7	± 0.3	± 1.2	± 0.0	± 0.3	± 0.3	± 1.0
SCOMatch	<u>92.1</u>	72.0	92.4	73.3	<u>71.0</u>	62.2	<u>74.6</u>	60.7	64.1	65.2	69.7	72.5	87.6	61.4	51.3	45.8
	± 0.6	± 1.0	± 0.3	± 0.7	± 0.8	± 0.9	± 0.4	± 0.6	± 0.3	± 3.0	± 0.5	± 0.8	± 0.5	± 1.8	± 0.4	± 4.2
ProSub	88.0	<u>87.8</u>	88.5	88.0	63.6	76.5	69.2	79.2	55.2	74.6	61.1	<u>78.5</u>	<u>92.7</u>	88.5	61.0	<u>75.8</u>
	± 1.0	± 1.1	± 2.1	± 2.5	± 1.0	± 2.2	± 1.4	± 0.9	± 1.1	± 1.9	± 1.9	± 3.4	± 1.0	± 0.7	± 0.2	± 0.3
<i>SkipAlign</i>	92.0	94.9	<u>93.1</u>	96.6	71.6	82.8	75.0	83.9	67.2	80.1	70.4	82.8	92.8	90.7	<u>60.2</u>	77.1
	± 0.4	± 0.7	± 0.5	± 0.2	± 0.4	± 1.6	± 0.2	± 3.2	± 0.1	± 0.7	± 0.5	± 0.6	± 0.4	± 1.0	± 0.4	± 1.3

Table 1: Closed-set accuracy (Acc.) and overall OOD detection AUC across all experimental configurations. Results are reported as mean \pm standard deviation over three runs with a fixed random seed. The best performance is shown in bold, and the second best is underlined.

to these negative samples. The final OOD detector loss is defined as:

$$\mathcal{L}_{OD} = \mathcal{L}_{od}^{OVA}(X) + \lambda_{em} \mathcal{L}_{od}^{em}(\mathcal{U}) + \lambda_{SOCR} \mathcal{L}_{od}^{SOCR}(\mathcal{U}) + \lambda_{neg} \mathcal{L}_{od}^{neg}(\mathcal{U}) \quad (8)$$

Each λ weights the corresponding loss component.

Overall Objective. The final training objective combines the three loss terms:

$$\mathcal{L}_{Total} = \lambda_{CC} \mathcal{L}_{CC} + \lambda_{OD} \mathcal{L}_{OD} + \lambda_{SNA} \mathcal{L}_{SNA} \quad (9)$$

Here, \mathcal{L}_{SNA} refers to the selective alignment loss described above, and λ_{CC} , λ_{OD} , λ_{SNA} control the relative weights of each objective.

Experiments

Experimental Setup

We evaluate *SkipAlign* under standard OSSL protocols, measuring both ID (closed-set) classification accuracy and OOD detection performance using AUC (Area Under the Receiver Operating Characteristic curve). To rigorously assess OOD generalization, we include a **diverse set of unseen OOD datasets** in addition to the seen OODs used during training. For overall detection performance, we report *overall AUC*, defined as the mean AUC across all OOD datasets (seen and unseen), weighting each dataset equally.

Datasets and Backbones. For CIFAR-10 and CIFAR-100 (Krizhevsky, Hinton et al. 2009), we use WideResNet-28-2 (Zagoruyko 2016) and vary the number of ID/OOD classes with 25 or 50 labeled samples per class. For CIFAR-10, six animal classes are selected as ID classes, and the remaining four are used as OOD. In CIFAR-100, we first group classes by their semantic super-classes and then assign them into disjoint ID and OOD subsets. We evaluate two settings with different OOD proportions: one with 55 ID and 45 OOD classes, and another with 80 ID and 20 OOD classes. For each configuration, 25 or 50 labeled samples per ID class are used for training. For ImageNet30 (Hendrycks and Gimpel 2016), a 30-class subset of ImageNet (Deng et al. 2009), we adopt ResNet-18 (He et al. 2016) and use

20 classes as ID and 10 as OOD, with 5% of the ID data labeled (65 samples per class). For TinyImageNet (Le and Yang 2015), we use WideResNet-28-4 on a 100 ID / 100 OOD split, with 10% labeled data per ID class.

Unseen OOD Evaluation. To evaluate the generalization ability to unseen OOD data, we augment the standard seen OOD test with additional unseen OOD sources. For CIFAR-10 and CIFAR-100, we use CIFAR-100 and CIFAR-10, respectively, as unseen OOD datasets, along with CIFAR-100C and CIFAR-10C (Hendrycks and Dietterich 2019), ImageNet (Deng et al. 2009), LSUN (Yu et al. 2015), SVHN (Netzer et al. 2011), and Gaussian noise. For ImageNet30, we include a diverse set of natural and synthetic categories such as Gaussian noise, Stanford Dogs (Khosla et al. 2011), CUB-200 (Welinder et al. 2010), Flowers102 (Nilsback and Zisserman 2008), Caltech-256 (Griffin et al. 2007), Describable Textures Dataset (DTD) (Cimpoi et al. 2014), LSUN, and SUN397 (Xiao et al. 2010). For TinyImageNet (TIN), we adopt unseen OOD datasets including Gaussian noise, SVHN, DTD, LSUN, and SUN397. **Baselines.** We compare *SkipAlign* against recent OSSL methods, including IOMatch, SSB, ProSub, and SCOMatch. All methods are evaluated under the same dataset splits and hardware environment (H100 or RTX 3090), using official or faithfully reproduced implementations.

Main Results

Performances. Table 1 reports closed-set accuracy and overall AUC across all experimental configurations. *SkipAlign* consistently achieves the highest AUC in every configuration, with improvements of **up to 7.1 points** over the strongest baseline. Closed-set accuracy is also the best in nearly all settings, with only marginal drops on CIFAR-10 and TinyImageNet, and shows a maximum gain of 2.7 percentage points (pp) compared to competing methods.

Figure 4 further breaks down AUC performance for each OOD dataset (including both seen and unseen sources). *SkipAlign* demonstrates **consistently strong and balanced detection performance** across nearly all OOD datasets. One exception is LSUN under the CIFAR-100 setting: as the




Figure 4: Per-dataset OOD detection AUC for various OSSL methods across different experimental configurations (CIFAR-10, CIFAR-100 with varying ID/OOD splits and 25 labels, and ImageNet30). *SkipAlign* consistently achieves strong and balanced detection performance on most OOD datasets.




Figure 5: t-SNE visualization of feature embeddings from the OOD detector (CIFAR-10, 6/4-50). Red contours highlight problematic regions in prior methods: *SkipAlign* yields compact, class-specific ID clusters and clear ID-OOO separation, with unseen OODs naturally positioned in non-ID regions or around the OOD boundary.

number of ID classes grows, many scene-centric categories overlap semantically with LSUN, causing feature-level similarity and reduced separability of OOD samples. Interestingly, ProSub appears to perform well on LSUN in the 80-class CIFAR-100 setting, but this stems from poor ID classification—its closed-set accuracy is 12 pp lower than that of *SkipAlign*, indicating that it misclassifies many ID samples rather than truly improving OOD discrimination. Finally, unlike other methods, whose AUC varies substan-

tially between 25- and 50-label settings, *SkipAlign* maintains consistently high OOD detection performance regardless of the number of labeled samples.

t-SNE Analysis of OOD Detector Embeddings. To better understand *SkipAlign*'s performance gains, we visualize and compare how different methods structure embeddings within the OOD detector using t-SNE (Figure 5). ProSub separates seen OODs from ID clusters but leaves unseen OODs nearby, which can be located within the ID subspace and cause mis-detection. Additionally, the cluster for the ID “Bird” class remains fragmented due to early misalignment when relying solely on feature-subspace similarity. SCOMatch defines a synthetic OOD class that appears as a distinct region in the embedding space, yet most unseen OODs fall outside this predefined region and are misclassified as ID. SSB produces mostly compact ID clusters, but the “Dog” class lacks clear boundaries and overlaps with nearby seen or unseen OODs. Unseen OODs tend to gather near the center and intrude into ID regions, reflecting the limitations of OVA classifiers trained with very limited labeled data.

In contrast, *SkipAlign* leverages SNA to better structure the embedding space. The SNA module rotates ID samples toward their corresponding class prototypes, facilitating class alignment, while uncertain or OOD samples are rotated toward directions orthogonal to all prototypes, preventing OOD data from being assigned to any class. Through this mechanism, ID sample embeddings naturally condense around their class prototypes, while OOD samples reside in the voids outside the regions occupied by ID classes.

Ablation

Selective Non-Alignment Loss Variants. We assess the effectiveness of our Unlabeled SNA loss (\mathcal{L}_{USNA} , Eq. 2) by comparing it with two alternatives in Table 2. The first variant, *OVA with Confident ID*, augments OVA classifier training by treating confidently filtered ID samples from unlabeled data as additional positives. This approach significantly lowers both accuracy and AUC compared to ours, suggesting that even small amounts of label noise can destabilize the one-vs-all decision boundaries that are designed to rely on clean labeled data. The second variant, *Cosine*

Loss	80/20-25		55/45-25	
	Acc.	AUC	Acc.	AUC
OVA w/ Conf. ID	64.4	69.7	68.5	69.7
Cosine Similarity	65.9	79.4	71.4	83.4
\mathcal{L}_{USNA} (Ours)	67.2	80.8	71.8	84.1

Table 2: Comparison of Unlabeled SNA loss variants on CIFAR-100 under two dataset configurations.

Exp.	Loss Combination			Acc.	AUC		
	IA	PA	USNA		Seen	Unseen	Overall
1				64.7	88.6	76.2	78.0
2	✓	✓		67.2	91.5	77.7	79.7
3			✓	67.2	89.9	78.3	80.0
4	✓	✓	✓	67.2	91.2	79.1	80.9

Table 3: Ablation on SNA loss combinations on CIFAR-100 80/20-25.

Similarity, only pulls confident samples toward prototypes without repelling uncertain ones (i.e., excluding them from training). On the 55/45-25 setting, its accuracy is close to ours, but on the 80/20-25 setting it drops by about 1.3 pp, and AUC is consistently lower in both settings. This indicates that selective attraction alone (pull only) cannot maintain sufficient separation between ID and OOD embeddings. By contrast, \mathcal{L}_{USNA} achieves the best balance across both settings, improving both accuracy and OOD detection performance. Its selective attraction combined with mild repulsion of uncertain features yields a more structured and discriminative feature space for robust OOD detection.

SNA Loss Combination. Table 3 evaluates the contribution of each loss component within the SNA module: \mathcal{L}_{IA} , \mathcal{L}_{PA} , and \mathcal{L}_{USNA} . Without these losses (Exp. 1), both accuracy and AUC are low due to poor feature structuring. Supervised alignment alone (Exp. 2) improves accuracy and seen OOD detection but fails to leverage unlabeled data, limiting the generalization ability of OOD detection. USNA alone (Exp. 3) slightly improves overall AUC over Exp. 2 by avoiding false alignment of uncertain samples, but lacking supervised alignment degrades ID cluster quality and lowers seen AUC. Combining all three losses (Exp. 4) yields the best performance: labeled losses form compact and well-separated ID clusters, while USNA keeps uncertain samples away from prototypes, enlarging ID–OOD margins. This synergy achieves the highest total AUC, highlighting the benefit of combining supervised alignment with selective non-alignment for robust OOD detection.

Prototype Refinement. We analyze the effect of the unlabeled contribution ratio r_u , a weighting parameter in our adaptive prototype generation that determines how much confidently pseudo-labeled samples contribute relative to labeled data. Figure 6 shows that increasing r_u from 0 (labeled only) to a moderate value (around 0.3–0.5) improves both accuracy and overall OOD AUC, indicating that leveraging reliable unlabeled samples helps refine prototypes and leads




Figure 6: Effect of unlabeled contribution r_u in prototype refinement on CIFAR-100 80/20-25.

η_{ID}	CIFAR10 6/4-25		CIFAR100 55/45-25	
	Acc.	AUC	Acc.	AUC
0.0	91.4	80.0	70.9	82.5
0.3	91.6	89.5	71.5	83.1
0.5	92.1	94.4	71.8	83.2
0.7	92.4	91.4	71.5	85.1
0.9	91.5	89.8	70.8	83.1

Table 4: Effect of the OOD detector threshold η_{ID} in dual-gate ID selection ($\tau_{ID} = 0.99$, temperature 0.5).

to a more stable decision geometry that generalizes beyond seen OODs. However, setting r_u too high introduces residual pseudo-label noise, causing slight performance degradation.

ID Selection Threshold. We analyze the impact of varying the OOD confidence threshold η_{ID} in our dual-gate ID selection, with the classifier threshold fixed at $\tau_{ID} = 0.99$ and softmax temperature 0.5 (Table 4). All nonzero thresholds outperform the classifier-only gating ($\eta_{ID} = 0$), demonstrating the importance of incorporating the OOD detector for reliable ID selection. On both CIFAR-10 and CIFAR-100, performance steadily improves as η_{ID} increases from 0 and remains stable around moderate values (e.g., 0.5–0.7). While overly conservative settings (e.g., $\eta_{ID} = 0.9$) may slightly reduce accuracy by filtering out some useful ID samples, this behavior reflects the binary nature of ID/OOD separation and indicates that our method performs robustly across a wide range of threshold values.

Conclusion

We introduced *SkipAlign*, an OSSSL framework based on Selective Non-Alignment (SNA) to improve generalization beyond seen OOD samples. SNA avoids overfitting by skipping alignment for uncertain samples and applying angular repulsion, encouraging OOD features to remain in a separate void. Combined with dual-gate filtering and prototype refinement, *SkipAlign* effectively learns under label noise. Experiments show strong AUC and generalization gains over prior methods while maintaining competitive accuracy. Although performance is robust across diverse OODs, challenges remain when unseen classes closely resemble ID ones (e.g., LSUN under the CIFAR-100 setting). Future work includes enhancing feature disentanglement for such cases.

Acknowledgments

This work was supported by the Information and Communications Technology Planning and Evaluation (IITP) grant funded by the Korea Ministry of Science and ICT (MSIT) (Grant No. 2710085245); the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) through the Engineering Research Center (ERC) program (Grant No. 2710077975); and by the Mobile eXperience (MX) Business, Samsung Electronics Co., Ltd. Hyung-Sin Kim is the corresponding author.

References

- Cimpoi, M.; Maji, S.; Kokkinos, I.; Mohamed, S.; and Vedaldi, A. 2014. Describing textures in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3606–3613.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Fan, Y.; Kukleva, A.; Dai, D.; and Schiele, B. 2023. SSB: Simple but Strong Baseline for Boosting Performance of Open-Set Semi-Supervised Learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 16068–16078.
- Griffin, G.; Holub, A.; Perona, P.; et al. 2007. Caltech-256 object category dataset. Technical report, Technical Report 7694, California Institute of Technology Pasadena.
- Guo, L.-Z.; Zhang, Z.-Y.; Jiang, Y.; Li, Y.-F.; and Zhou, Z.-H. 2020. Safe deep semi-supervised learning for unseen-class unlabeled data. In *International conference on machine learning*, 3897–3906. PMLR.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hendrycks, D.; and Dietterich, T. 2019. Benchmarking Neural Network Robustness to Common Corruptions and Perturbations. *Proceedings of the International Conference on Learning Representations*.
- Hendrycks, D.; and Gimpel, K. 2016. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *arXiv preprint arXiv:1610.02136*.
- Huang, J.; Fang, C.; Chen, W.; Chai, Z.; Wei, X.; Wei, P.; Lin, L.; and Li, G. 2021. Trash to treasure: Harvesting ood data with cross-modal matching for open-set semi-supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8310–8319.
- Khosla, A.; Jayadevaprakash, N.; Yao, B.; and Fei-Fei, L. 2011. Novel Dataset for Fine-Grained Image Categorization. In *First Workshop on Fine-Grained Visual Categorization, IEEE Conference on Computer Vision and Pattern Recognition*. Colorado Springs, CO.
- Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; and Krishnan, D. 2020. Supervised contrastive learning. *Advances in neural information processing systems*, 33: 18661–18673.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Le, Y.; and Yang, X. 2015. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7): 3.
- Li, J.; Zhou, P.; Xiong, C.; and Hoi, S. C. 2020. Prototypical contrastive learning of unsupervised representations. *arXiv preprint arXiv:2005.04966*.
- Li, Z.; Qi, L.; Shi, Y.; and Gao, Y. 2023. Iomatch: Simplifying open-set semi-supervised learning with joint inliers and outliers utilization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15870–15879.
- Netzer, Y.; Wang, T.; Coates, A.; Bissacco, A.; Wu, B.; Ng, A. Y.; et al. 2011. Reading digits in natural images with unsupervised feature learning. In *NIPS workshop on deep learning and unsupervised feature learning*. Granada.
- Nilsback, M.-E.; and Zisserman, A. 2008. Automated Flower Classification over a Large Number of Classes. In *Indian Conference on Computer Vision, Graphics and Image Processing*.
- Saito, K.; Kim, D.; and Saenko, K. 2021. Openmatch: Open-set semi-supervised learning with open-set consistency regularization. *Advances in Neural Information Processing Systems*, 34: 25956–25967.
- Sohn, K.; Berthelot, D.; Carlini, N.; Zhang, Z.; Zhang, H.; Raffel, C. A.; Cubuk, E. D.; Kurakin, A.; and Li, C.-L. 2020. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems*, 33: 596–608.
- Wallin, E.; Svensson, L.; Kahl, F.; and Hammarstrand, L. 2024a. Improving open-set semi-supervised learning with self-supervision. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2356–2365.
- Wallin, E.; Svensson, L.; Kahl, F.; and Hammarstrand, L. 2024b. ProSub: Probabilistic Open-Set Semi-supervised Learning with Subspace-Based Out-of-Distribution Detection. In *European Conference on Computer Vision*, 129–147. Springer.
- Wang, Z.; Xiang, L.; Huang, L.; Mao, J.; Xiao, L.; and Yamasaki, T. 2024. SCOMatch: Alleviating Overtrusting in Open-set Semi-supervised Learning. In *European Conference on Computer Vision*, 217–233. Springer.
- Welinder, P.; Branson, S.; Mita, T.; Wah, C.; Schroff, F.; Belongie, S.; and Perona, P. 2010. Caltech-UCSD birds 200.
- Xiao, J.; Hays, J.; Ehinger, K. A.; Oliva, A.; and Torralba, A. 2010. Sun database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE computer society conference on computer vision and pattern recognition*, 3485–3492. IEEE.
- Yu, F.; Seff, A.; Zhang, Y.; Song, S.; Funkhouser, T.; and Xiao, J. 2015. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*.
- Yu, Q.; Ikami, D.; Irie, G.; and Aizawa, K. 2020. Multi-task curriculum framework for open-set semi-supervised learning. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XII 16*, 438–454. Springer.

Zagoruyko, S. 2016. Wide residual networks. *arXiv preprint arXiv:1605.07146*.