Two-Stage Angular Alignment for Positive-Unlabeled Learning

Vasileios Sevetlidis[1,2][a], George Pavlidis[1][b] and Antonios Gasteratos[2][c]
[1]*Athena Research Center, Xanthi, Greece*
[2]*Department of Production and Management Engineering, Democritus University of Thrace, Xanthi, Greece*

Keywords: Positive-Unlabeled Learning, Angular Prototypes, von Mises-Fisher Distribution, Curriculum Learning, Interpretable Deep Learning.

Abstract: Positive-Unlabeled (PU) learning addresses the binary classification problem where only positive and unlabeled data are available—a setting common in applications such as medical diagnosis and web mining. We introduce a novel two-stage approach based on angular alignment in feature space, where a learnable prototype vector represents the directional centroid of the positive class. In the first stage, the model aligns labeled positives toward this prototype to promote angular compactness; in the second, it repels overly similar unlabeled instances to refine the decision boundary without prematurely assigning negative labels. Our method employs a directional loss inspired by von Mises–Fisher geometry, a dynamic stage-switching curriculum, and maintains a highly parameter-efficient design. Experiments on CIFAR-10 and SVHN demonstrate strong performance and competitive results compared to state-of-the-art PU learning methods. The approach also yields semantically structured latent spaces, highlighting the value of angular geometry for interpretable and effective representation-based PU learning in visual domains.

1 INTRODUCTION

Learning from partially labeled data is a core challenge in machine learning. In *Positive-Unlabeled* (PU) learning, only some positives are labeled and all remaining samples are unlabeled, a setting common in medical diagnosis (Bekker and Davis, 2020), fault detection (Sevetlidis et al., 2024a; Sevetlidis et al., 2023), and web filtering (Elkan and Noto, 2008) where negatives are costly or ambiguous. Classical PU methods adapt empirical risk minimization to the absence of labeled negatives, often assuming *SCAR* and relying on class-prior estimation or risk-correction terms (Du Plessis et al., 2014; Kiryo et al., 2017); despite strong theory, these risk-based approaches can be brittle when priors are unknown or data are high-dimensional.

We instead propose a geometric PU approach that avoids risk rebalancing and class-prior estimation. Our *two-stage angular alignment* framework learns in a normalized embedding space with a prototype vector μ^+ that represents the positive direction on the unit hypersphere. In stage one, labeled positives are aligned toward μ^+ to induce angular compactness. In stage two, we add a soft repulsion for unlabeled samples that lie within a cosine-similarity margin of μ^+, creating a confidence-aware exclusion zone that refines the boundary without explicit negative supervision. The method is lightweight and interpretable for vision models, requiring no class prior, pseudo-labeling, or labeled negatives. Experiments on CIFAR-10 and SVHN show state-of-the-art or competitive performance against recent PU baselines, and we provide theoretical insights into prototype stability and margin-based separation under directional losses.

The remainder of this paper is organized as follows: Section 2 reviews related work, Section 3 presents the proposed framework, Section 4 analyzes its geometric properties, Section 5 reports results with ablations in Section 6, and the work concludes in Section 7.

2 RELATED WORK

PU learning addresses weak supervision settings where only some positives are labeled and all other samples are unlabeled, often under the SCAR as-

[a] https://orcid.org/0000-0001-9348-8786
[b] https://orcid.org/0000-0002-9909-1584
[c] https://orcid.org/0000-0002-5421-0332

Sevetlidis, V., Pavlidis, G. and Gasteratos, A.
Two-Stage Angular Alignment for Positive-Unlabeled Learning.
DOI: 10.5220/0013964100004067
Paper published under CC license (CC BY-NC-ND 4.0)
In *Proceedings of the 15th International Conference on Pattern Recognition Applications and Methods (ICPRAM 2026)*, pages 225-232
ISBN: 978-989-758-797-9; ISSN: 2184-4313
Proceedings Copyright © 2026 by SCITEPRESS – Science and Technology Publications, Lda.

sumption. Elkan and Noto (Elkan and Noto, 2008) formalized this setting and proposed a posterior correction, motivating risk-based ERM methods such as uPU/nnPU (Kiryo et al., 2017), where nnPU enforces non-negativity to curb overfitting. Despite strong empirical performance, these approaches depend on accurate class-prior estimation and SCAR, which can break in practice (Bekker and Davis, 2020). To relax SCAR, instance-dependent labeling models (e.g., SAR-PU, PUSB (Kato et al., 2019), PGPU (He et al., 2020)) estimate labeling propensity from features, improving robustness when positives are selected non-randomly (Bekker et al., 2019).

A complementary line of work reduces reliance on priors by exploiting representation geometry: Dense-PU (Sevetlidis et al., 2024b) uses norm/density cues, and follow-up work synthesizes dense positives in latent space to model the positive region (Sevetlidis et al., 2024c); tensor-decomposition approaches have also been applied successfully to PU settings (Žunkovič, 2023). Recent methods further incorporate pseudo-labeling and contrastive objectives, e.g., PUUPL (Dorigatti et al., 2024) refines pseudo-labels via uncertainty/consistency, Dist-PU (Zhao et al., 2022) aligns predicted label distributions to the expected prior, and contrastive PU objectives based on clustering have been explored in (Acharya et al., 2022). PU learning has also been adapted to specialized applications such as drug interaction prediction, road black-spot detection, and knowledge graph completion (Zheng et al., 2019; Karamanlis et al., 2023; Tang et al., 2022).

Our approach is representation-based but avoids class-prior estimation, pseudo-labeling, and loss-correction schemes. We instead propose a two-stage angular alignment method on the unit hypersphere, aligning labeled positives to a prototype and repelling nearby unlabeled samples via a cosine-margin criterion. Unlike norm/density-based separation (e.g., Dense-PU), we enforce directional separation directly through cosine similarity, and unlike pseudo-labeling pipelines (e.g., PUUPL, PSPU), we avoid iterative relabeling, yielding a lightweight and more interpretable training dynamic.

3 PROPOSED METHOD

We propose a two-stage angular alignment framework for PU learning that operates in a normalized embedding space, leveraging the geometry of the unit hypersphere. At the core of our approach is a learnable prototype vector $\mu^+ \in \mathbb{S}^{k-1}$, which serves as the angular centroid of the positive class. During training, the model first aligns known positives toward this prototype and subsequently repels overly similar unlabeled instances to induce a decision boundary. This structure-aware formulation promotes class separation without requiring explicit negative supervision.

Unlike classical PU methods that rely on class prior estimation or surrogate risk correction, our framework enforces angular constraints directly in feature space. The prototype μ^+ is updated using an exponential moving average (EMA), which ensures smooth adaptation across training iterations.

Let $\mathcal{X} \subset \mathbb{R}^d$ denote the input space, and let $y \in \{0,1\}$ represent binary labels, where $y = 1$ indicates a positive instance. In the PU setting, we are given a dataset $\mathcal{D}_{\text{train}} = \mathcal{D}_P \cup \mathcal{D}_U$, where $\mathcal{D}_P = \{(\mathbf{x}_i, y_i = 1)\}_{i=1}^{n_P}$ consists of labeled positives, and $\mathcal{D}_U = \{\mathbf{x}_j\}_{j=1}^{n_U}$ contains unlabeled instances that may belong to either class.

Most existing approaches rely on the *Selected Completely At Random* (SCAR) assumption, under which labeled positives are a random subset of all positives. This allows reformulating empirical risk minimization (ERM) using estimable components from $\mathcal{D}_P$ and $\mathcal{D}_U$. However, these risk estimators often require class prior estimation and suffer from instability under distribution shift or data scarcity.

In contrast, our approach directly operates in a normalized embedding space, leveraging cosine geometry to guide training without class priors or labeled negatives.

The first stage deals with positive alignment via angular attraction. During the early training phase ($t < s$), the model encourages positive samples to align with the prototype vector μ^+. Let $f_\theta : \mathcal{X} \to \mathbb{R}^k$ be the encoder, and define the normalized embedding $\mathbf{z}_i = \text{normalize}(f_\theta(\mathbf{x}_i))$. The stage-1 loss is:

$$\mathcal{L}_{\text{stage1}} = -\kappa \cdot \frac{1}{|\mathcal{P}|} \sum_{i \in \mathcal{P}} \mathbf{z}_i^\top \mu^+, \tag{1}$$

where $\kappa > 0$ controls the strength of alignment. This term builds a compact positive cluster in angular space while **ignoring unlabeled data**, avoiding premature supervision on uncertain points.

The second stage performs angular separation via margin-based repulsion. Once the prototype has stabilized (after epoch s), the model enters a contrastive phase. In addition to aligning positives, it repels unlabeled samples whose cosine similarity with μ^+ exceeds a predefined margin m. The stage-2 loss is:

$$\mathcal{L}_{\text{stage2}} = -\kappa \cdot \tfrac{1}{|\mathcal{P}|} \textstyle\sum_{i \in \mathcal{P}} \mathbf{z}_i^\top \mu^+ + \kappa \cdot \tfrac{1}{|\mathcal{U}|} \textstyle\sum_{i \in \mathcal{U}} \max(0, \mathbf{z}_i^\top \mu^+ - m), \tag{2}$$

The margin $m \in [0,1]$ defines a cosine similarity threshold that induces a spherical cap around μ^+.

Only unlabeled instances within this cap—i.e., overly aligned with the prototype—are softly repelled. This discourages ambiguous samples from crowding the positive region, while leaving the rest untouched.

Intuitively, the repulsion margin creates a *confidence-aware exclusion zone* around the prototype. It acts as a curriculum: alignment precedes repulsion, ensuring that latent positives are not prematurely pushed away. Over time, this mechanism refines the boundary between positives and the rest of the space.

The full training loss is defined as:

$$\mathcal{L}^{(t)} = \begin{cases} \mathcal{L}_{\text{stage1}}, & \text{if } t < s \\ \mathcal{L}_{\text{stage2}}, & \text{otherwise} \end{cases} \tag{3}$$

This stagewise schedule allows the model to first consolidate the geometry of the positive class, then apply discriminative pressure to the unlabeled set.

The prototype vector is updated after each batch using an exponential moving average (EMA):

$$\mu^+ \leftarrow \text{normalize}\left((1-\alpha)\cdot \mu^+ + \alpha \cdot \frac{1}{|\mathcal{P}|} \sum_{i \in \mathcal{P}} \mathbf{z}_i \right), \tag{4}$$

where $\alpha \in (0,1)$ is a smoothing factor. Normalization ensures that μ^+ remains on the hypersphere $\mathbb{S}^{k-1}$. The EMA mechanism provides stable updates and tracks the moving center of the positive cluster without requiring backpropagation through the update.

The overall training process alternates between computing minibatch losses and updating model parameters and the prototype. The parameter κ is a learnable scalar parameter updated jointly with the encoder. In practice, we apply normalization after each epoch, not after every batch-level update. This allows us to track the evolution of the prototype's norm during training, while still ensuring unit-norm consistency across stages. The full procedure is summarized in Algorithm 1.

Algorithm 1: Two-Stage Angular Alignment for PU Learning.

Data: Positive set $\mathcal{P}$, unlabeled set $\mathcal{U}$, encoder f_θ, margin m, learning rate η, EMA update rate α
Result: Trained encoder f_θ and positive prototype μ^+
Initialize prototype $\mu^+ \sim \mathbb{S}^{k-1}$;
Set epoch counter $t \leftarrow 0$;
while *not converged* **do**
  foreach *minibatch* $\mathcal{B} \subset \mathcal{P} \cup \mathcal{U}$ **do**
    foreach $\mathbf{x}_i \in \mathcal{B}$ **do**
      $\mathbf{z}_i \leftarrow \text{normalize}(f_\theta(\mathbf{x}_i))$;
    end
    Compute stage-specific loss $\mathcal{L}^{(t)}$ (Eq. (3));
    $\theta \leftarrow \theta - \eta \nabla_\theta \mathcal{L}^{(t)}$;
    $\mathcal{B}^+ \leftarrow \mathcal{B} \cap \mathcal{P}$;
    if $|\mathcal{B}^+| > 0$ **then**
      $\bar{\mathbf{z}}^+ \leftarrow \frac{1}{|\mathcal{B}^+|} \sum_{\mathbf{x}_i \in \mathcal{B}^+} \mathbf{z}_i$;
      $\mu^+ \leftarrow$ normalize$((1-\alpha)\mu^+ + \alpha \bar{\mathbf{z}}^+)$;
    end
  end
  $t \leftarrow t+1$;
end

4 THEORETICAL INTUITION

We provide a geometric perspective on the proposed two-stage angular alignment framework. Modeling normalized embeddings as points on the unit hypersphere $\mathbb{S}^{k-1} \subset \mathbb{R}^k$, we demonstrate that our loss function inherently encourages class-wise concentration, inter-class separation, and prototype stability.

We assume that positive embeddings follow a von Mises–Fisher (vMF) distribution centered at $\mu^* \in \mathbb{S}^{k-1}$ with concentration $\kappa > 0$. This reflects the intuitive notion of compact angular clusters for positives.

Proposition 1 (vMF Equivalence)**.** *Minimizing the Stage 1 loss*

$$\mathcal{L}_{stage1} = -\kappa \cdot \frac{1}{n_P} \sum_{i=1}^{n_P} \mathbf{z}_i^\top \mu^+$$

is equivalent to computing the maximum likelihood estimator of the mean direction μ^* under a vMF distribution.

Sketch. Maximizing the log-likelihood of a vMF-distributed set of unit vectors $\{\mathbf{z}_i\}$ yields the sample mean direction:

$$\mu^* = \frac{\sum_i \mathbf{z}_i}{\|\sum_i \mathbf{z}_i\|}.$$

The Stage 1 loss encourages μ^+ to align with this mean, recovering the MLE of μ^* up to scaling. Full proof omitted for brevity. □

This shows that the prototype μ^+ converges toward the average positive direction, promoting intra-class compactness (Fig. 1, left).

In Stage 2, we introduce a cosine margin $m \in [0,1]$ to softly repel unlabeled embeddings that are too similar to the prototype. Geometrically, this defines a spherical cap:

$$\mathcal{R}_m = \{\mathbf{z} \in \mathbb{S}^{k-1} \mid \mathbf{z}^\top \mu^+ \geq m\}.$$

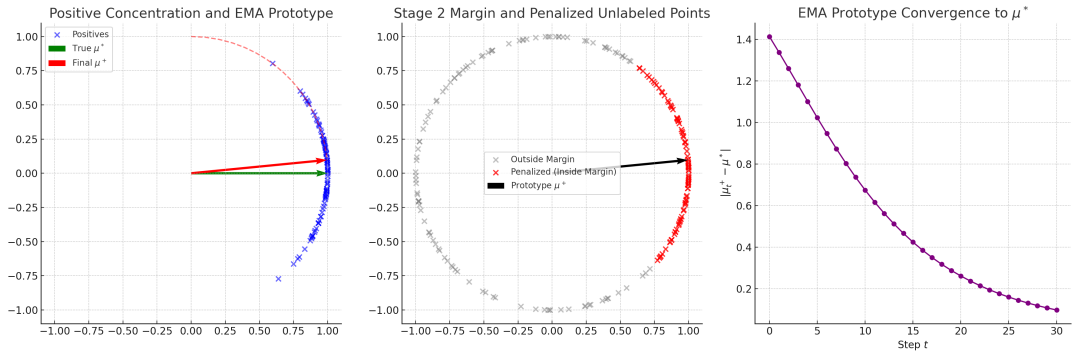

Figure 1: **(Left)** Labeled positives (blue) drawn from a von Mises–Fisher distribution cluster around the true mean μ^* (green). The learnable prototype μ^+ (red) converges toward μ^* via EMA updates. **(Middle)** In Stage 2, a cosine margin m defines a spherical cap where overly similar unlabeled embeddings (red) are softly repelled. **(Right)** Prototype convergence over time, measured as $\|\mu_t^+ - \mu^*\|$.

Lemma 2 (Sketch)**.** *Let $p_U(\mathbf{z})$ be the density of unlabeled embeddings. The repulsion loss:*

$$\mathcal{L}_{repel} = \kappa \cdot \mathbb{E}_{\mathbf{z}\sim p_U}\left[\max(0, \mathbf{z}^\top \mu^+ - m)\right]$$

is minimized when the mass of p_U inside $\mathcal{R}_m$ is reduced.

Sketch. The loss integrates over $\mathcal{R}_m$ with weight proportional to angular proximity. If p_U places less mass inside this margin, the loss is lower. See Fig. 1 (middle). □

This term discourages ambiguous unlabeled samples from intruding into the confident positive region, while avoiding hard negative assignments.

5 EXPERIMENTAL EVALUATION

We evaluate the proposed two-stage angular alignment method on two benchmark datasets commonly used in PU learning: CIFAR-10 and SVHN. Our evaluation is designed to measure effectiveness on the visual domain, compare against competitive baselines, and analyze the individual contributions of each design component via ablation studies and sensitivity analysis.

Following standard PU protocols (Zhao et al., 2022; Sevetlidis et al., 2024b; Yuan et al., 2025), we form a binary task on the **CIFAR-10** with non-animal classes (airplane, automobile, ship, truck) as positives and animal classes (cat, deer, dog, frog, horse, bird) as negatives. We label 1,000 positives; all remaining data (unlabeled positives plus all negatives) is treated as unlabeled. SVHN is split into even digits (0,2,4,6,8) which are regarded as positives and odd digits (1,3,5,7,9) as negatives respectively. We label 1,000 positives and treat the rest as unlabeled. SVHN is challenging due to cluttered backgrounds and large intra-class variation.

We evaluate on a held-out, fully labeled test set using **F1** (at the F1-maximizing threshold), with **precision** and **recall** reported at the same threshold, and **AUROC**. All results are averaged over ten random seeds, and thresholds are chosen on a disjoint validation split.

We employ a VGG11 encoder pretrained on ImageNet. The convolutional backbone is frozen, and a lightweight head is trained to produce k-dimensional embeddings. Unless otherwise specified, $k = 16$. Our method uses a cosine similarity classifier with a learnable prototype μ^+, updated via exponential moving average in Stage 2. In Stage 1, the prototype is initialized from labeled positives and used to estimate alignment; Stage 2 introduces a margin-based repulsion mechanism from the prototype to separate unlabeled instances. The alignment weight κ was selected via grid search on the validation set within reasonable ranges, and we observed stable performance across a broad range of values, indicating that the method is not overly sensitive to precise tuning. Training is conducted using the Adam optimizer with learning rate 10^{-4} for 20 epochs, batch size 128. The decision threshold is selected by maximizing F1 on validation data. Evaluation is based on similarity between test embeddings and the learned prototype.

We compare our angular alignment method against representative PU baselines spanning risk-estimation, bias correction, self-training, and contrastive/prototype learning: **uPU**(Du Plessis et al., 2014) uses an unbiased PU risk estimator and argues non-convex losses (e.g., ramp) reduce biased boundaries, with generalization and prior-robustness analysis; **nnPU**(Kiryo et al., 2017) enforces a non-negative correction to avoid negative empirical risks and overfitting in uPU, enabling stable deep training;

RP(Northcutt et al., 2017) applies confident learning to estimate label/noise structure and filter errors, improving positive selection reliability; **PUSB**(Kato et al., 2019) handles selection bias by learning a score that preserves class-posterior ordering, relaxing SCAR; **PUbN**(Hsieh et al., 2019) incorporates biased negatives via a risk estimator combining importance weighting with a PU-based preliminary estimate; **Self-PU**(Chen et al., 2020b) iteratively refines pseudo-labels via self-paced selection, loss calibration, and self-distillation; **aPU**(Hammoudeh and Lowd, 2020) targets arbitrarily shifted positives (abs-PU, PURR) while assuming negatives are fixed; **VPU**(Chen et al., 2020a) avoids explicit class-prior estimation via a variational objective that directly approaches Bayes-optimal decision rules; **ImbPU**(Su et al., 2021) reweights positive/unlabeled risks to align PU training with balanced PN risk and improve F1 under imbalance; **Dist-PU**(Zhao et al., 2022) matches the expected predicted label distribution to the ground-truth distribution to mitigate negative-prediction bias; **Dense-PU**(Sevetlidis et al., 2024b) uses dense (spatial) loss weighting for PU semantic segmentation with priors/consistency; **puNCE**(Acharya et al., 2022) introduces a class-prior reweighted contrastive objective for PU representation learning; **PiCO**(Wang et al., 2022) couples contrastive learning with prototype-based label disambiguation in an EM-like loop; **WConPU**(Yuan et al., 2025) extends weighted contrastive PU learning with prototype-based hard-negative mining and an EM-style interpretation; and **AngularPU**(Sevetlidis et al., 2025) performs prototype-based PU classification on the unit hypersphere by thresholding cosine similarity to a learnable positive prototype, and adds an unlabeled-only cosine-uniformity regularizer to prevent unlabeled embeddings from collapsing near the prototype (no class prior or pseudo-labeling).

5.1 Main Results

Table 1 presents the performance of our method across CIFAR-10 and SVHN, compared to recent PU learning techniques. Across all datasets, our method demonstrates strong performance, particularly in F1 score and AUROC, indicating both discriminative power and robustness in handling class imbalance. On **CIFAR-10**, our approach achieves competitive performance (F1 = 95.6). On **SVHN**, our method yields F1 = 89.70, AUROC = 94.70, and competitive precision-recall tradeoffs. This highlights its capacity to cope with noisy and visually diverse data, benefiting from the stage-wise alignment.

While our method achieves strong overall results across CIFAR-10 and SVHN, it trails slightly behind the latest contrastive and cost-sensitive PU learning methods such as PiCO and WConPU in certain metrics. Specifically, on CIFAR-10, our F1 score is slightly lower than that of PiCO and WConPU. These baselines often rely on iterative pseudo-label refinement, contrastive training pipelines, or cost-sensitive risk estimators that leverage fine-grained supervisory signals.

In contrast, our method trades peak performance for architectural simplicity, interpretability, and efficiency. It avoids pseudo-labeling, class-prior estimation, and extensive augmentation regimes, relying instead on a geometric alignment mechanism governed by angular similarity and a lightweight prototype. This makes it more suitable for deployment in scenarios with limited computational resources or where model transparency is critical.

The two-stage loss schedule further contributes to interpretability by aligning with a human-like learn-

Table 1: Performance comparison on CIFAR-10 and SVHN.

Method	**F1**	**Prec**	**Rec**	**AUC**
		CIFAR-10		
uPU	85.12 ± 0.43	87.21 ± 2.09	83.02 ± 1.98	94.98 ± 0.62
nnPU	86.11 ± 0.49	86.21 ± 1.02	86.03 ± 1.22	95.13 ± 0.55
RP	85.93 ± 0.30	86.02 ± 1.05	85.72 ± 1.61	95.21 ± 0.23
PUSB	86.18 ± 0.51	86.15 ± 0.56	86.22 ± 0.45	95.15 ± 0.50
PUbN	87.18 ± 0.54	87.85 ± 0.98	86.56 ± 1.87	94.44 ± 0.35
Self-PU	86.77 ± 1.12	86.26 ± 0.76	87.22 ± 2.16	95.52 ± 0.46
aPU	86.41 ± 0.41	86.31 ± 1.33	86.33 ± 0.71	95.11 ± 0.39
VPU	84.42 ± 1.12	86.71 ± 1.46	82.88 ± 2.93	94.55 ± 0.46
ImbPU	86.77 ± 0.56	86.72 ± 0.89	86.91 ± 0.78	95.53 ± 0.26
Dist-PU	89.85 ± 0.62	89.87 ± 1.09	89.84 ± 0.81	96.92 ± 0.45
Dense-PU	91.96 ± 0.80	92.68 ± 1.31	91.25 ± 1.12	93.22 ± 0.99
puNCE	94.21 ± 0.44	95.11 ± 0.88	93.43 ± 0.45	98.59 ± 0.53
PiCO	94.75 ± 0.49	94.89 ± 0.76	93.97 ± 0.47	98.67 ± 0.44
WConPU	96.43 ± 0.29	96.87 ± 0.54	96.02 ± 0.32	99.49 ± 0.22
AngularPU	92.87 ± 0.71	92.26 ± 0.64	93.49 ± 0.13	96.61 ± 0.50
Ours	95.59 ± 0.41	91.75 ± 0.23	94.28 ± 0.51	96.10 ± 0.37
		SVHN		
uPU	81.12 ± 0.56	87.11 ± 2.39	75.93 ± 2.68	91.93 ± 0.62
nnPU	82.01 ± 0.58	86.78 ± 1.15	77.25 ± 1.42	92.02 ± 0.52
RP	80.10 ± 0.32	84.01 ± 1.01	76.12 ± 1.51	89.75 ± 0.23
PUSB	82.11 ± 0.51	86.81 ± 0.51	78.01 ± 0.51	91.89 ± 0.52
PUbN	83.16 ± 0.54	88.26 ± 0.98	83.57 ± 1.87	92.03 ± 0.35
Self-PU	82.55 ± 1.06	86.16 ± 0.78	79.22 ± 2.35	91.73 ± 0.58
aPU	82.33 ± 0.56	86.29 ± 1.30	81.21 ± 0.79	91.56 ± 0.42
VPU	75.36 ± 2.84	79.56 ± 1.41	79.56 ± 1.41	73.31 ± 0.91
ImbPU	82.99 ± 0.56	86.69 ± 0.87	81.18 ± 0.82	91.79 ± 0.27
Dist-PU	83.66 ± 0.56	89.06 ± 0.89	84.36 ± 0.76	92.92 ± 0.49
Dense-PU	85.37 ± 0.91	89.32 ± 0.78	82.72 ± 0.99	93.25 ± 0.64
puNCE	87.01 ± 0.55	90.35 ± 0.92	83.81 ± 1.99	94.87 ± 0.35
PiCO	87.51 ± 0.44	90.47 ± 0.79	85.74 ± 0.64	95.58 ± 0.54
WConPU	90.45 ± 0.35	93.77 ± 0.67	87.54 ± 0.67	96.97 ± 0.59
AngularPU	89.27 ± 0.13	88.33 ± 0.22	90.27 ± 0.12	95.85 ± 0.96
Ours	89.70 ± 0.37	87.54 ± 0.30	89.37 ± 0.23	94.70 ± 0.76

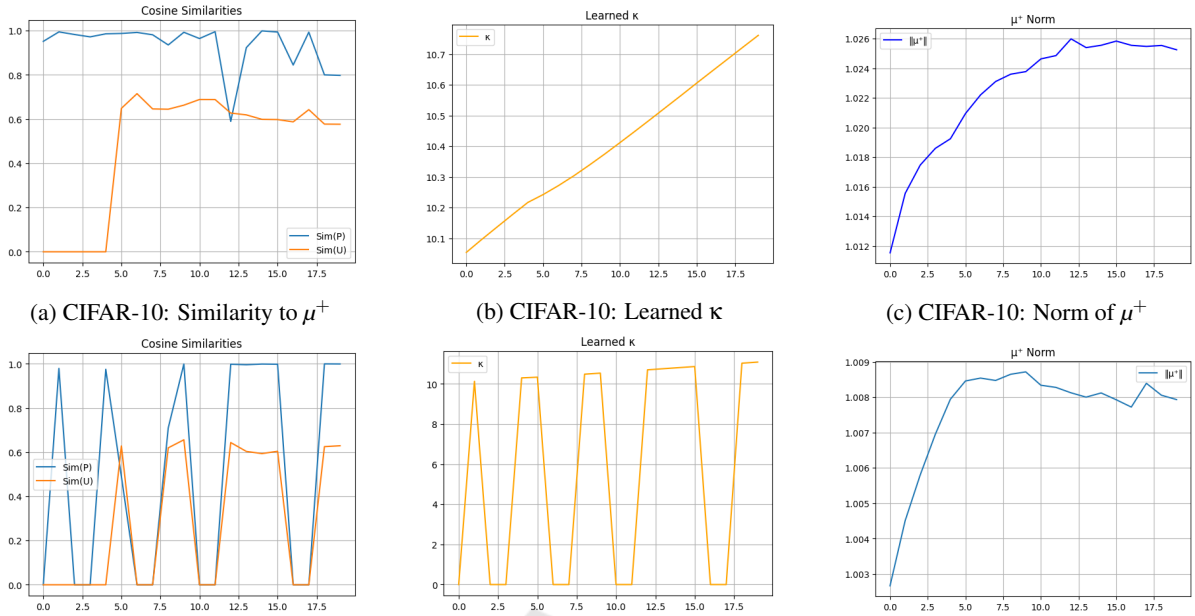

(a) CIFAR-10: Similarity to μ^+ (b) CIFAR-10: Learned κ (c) CIFAR-10: Norm of μ^+

(d) SVHN: Similarity to μ^+ (e) SVHN: Learned κ (f) SVHN: Norm of μ^+

Figure 2: Training dynamics across datasets. Each row corresponds to a dataset (CIFAR-10, SVHN), showing (left) cosine similarities of positives and unlabeled samples to μ^+, (center) the learned concentration parameter κ, and (right) the norm of the prototype vector μ^+.

ing curriculum. In the first stage, the model consolidates knowledge by pulling labeled positives toward μ^+. In the second stage, it begins to identify non-positive instances by repelling unlabeled points that are overly similar to the prototype. This progressive strategy ensures that the model's decision boundary evolves in a smooth and semantically meaningful manner, avoiding abrupt shifts caused by premature repulsion or noisy pseudo-labels. This behavior is reflected in the training dynamics of μ^+, κ, and cosine similarities, as shown in Figure 2. Figure 3 shows how the directional prototype μ^+ sits near the cluster of positive samples in the embedding space, even under varying visual domains.

The prototype vector μ^+ is updated via an exponential moving average (EMA), which ensures smooth adaptation and interpretability over time. Unlike fully learned prototypes, which may fluctuate in non-transparent ways, our update mechanism maintains a running summary of the positive class distribution. This enables monitoring and visualization of μ^+ convergence, and provides a stable reference point for understanding classification decisions. Additionally, the fixed cosine margin m used during Stage 2 offers a tunable and interpretable control over model sensitivity to ambiguous samples.

To better understand the semantic structure induced by our angular alignment loss, we visualize labeled positive samples with the lowest alignment to the learned prototype μ^+. As shown in Figure 4, these samples tend to be visually ambiguous, atypical in pose or lighting, or otherwise distinct from the majority of positives. This observation suggests that our prototype captures a robust central concept of the positive class, and that cosine similarity to μ^+ provides a meaningful measure of confidence.

6 ABLATION STUDY

We conduct a focused ablation study to evaluate the contribution of key design components of our framework. All experiments are reported as average F1 scores over ten seeds on CIFAR-10 and SVHN.

We observed that embedding dimensionality has a mild effect: moderate dimensions (e.g., $k = 16$) provide stable performance, while too small or too large embeddings slightly reduce accuracy. Detailed results are omitted for brevity. Our method relies on a two-stage training schedule: early alignment followed by margin-based repulsion. Table 2 shows that performance peaks when Stage 1 spans roughly half to three-quarters of training. Transitions that are too early (e.g., 25%) degrade performance by destabilizing the prototype, while excessively late transitions underutilize the repulsion term.

We also study the impact of learning the prototype μ^+ and the concentration parameter κ. Table 3 shows

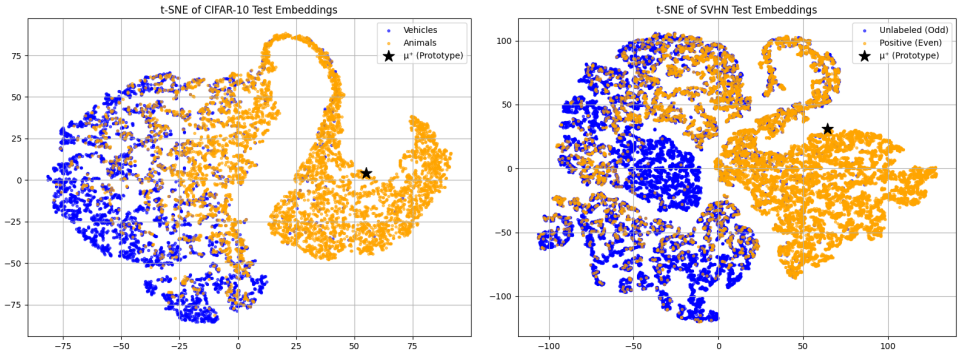

Figure 3: t-SNE projections of the learned feature space for CIFAR-10, SVHN. We visualize test-time embeddings colored by true label, along with the prototype μ^+ as a black star.

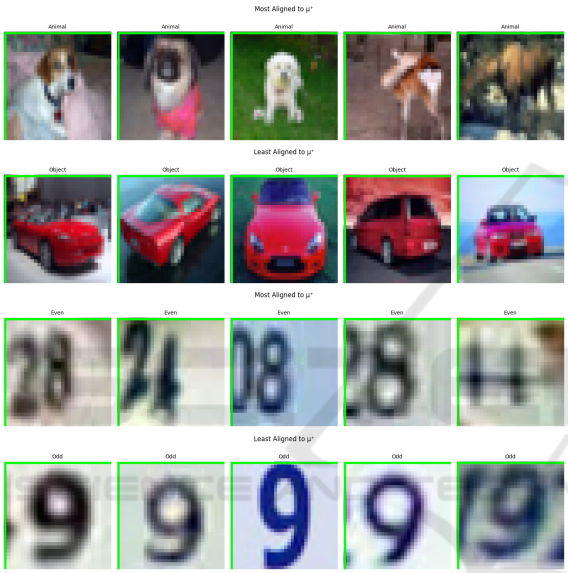

Figure 4: **Samples with the most and lowest alignment to the positive prototype μ^+ i.e., the positive samples that achieve the highest and lowest cosine similarity to μ^+ across two datasets. (Top: CIFAR-10 animals, Bottom: SVHN even digits.**

Table 2: Effect of Stage 1 duration measured by **F1 score**.

Epochs	Stage 1 Ratio	CIFAR-10	SVHN
10	0.25	0.923	0.875
10	0.50	0.933	0.869
10	0.75	0.921	0.879
20	0.50	0.925	**0.897**
20	0.75	**0.937**	0.856

that jointly learning both yields the highest performance. Fixing μ^+ or κ reduces the model's flexibility and accuracy.

Table 3: Impact of prototype and κ learning on F1 score.

Learn μ^+	Learn κ	CIFAR-10	SVHN
✗	✗	0.927	0.869
✗	✓	0.898	0.841
✓	✗	0.904	0.859
✓	✓	**0.929**	**0.881**

7 CONCLUSION

We proposed a geometric framework for PU learning based on two-stage angular alignment in normalized embedding space. Our method combines a learnable prototype vector with a curriculum-style loss, transitioning from positive attraction to margin-based repulsion, enabling effective learning without class priors, pseudo-labels, or negative samples. Operating in the angular domain with EMA-updated prototypes, the model defines interpretable decision boundaries via cosine similarity. Compared to prior PU approaches, it is simpler, lightweight, and analytically tractable. To our knowledge, this is the first PU method to jointly exploit hyperspherical geometry, a learnable prototype, and a curriculum-based loss schedule. This combination enhances transparency and robustness, particularly in settings where model interpretability is essential.

ACKNOWLEDGEMENTS

This work has been partially supported by the ARGUS EU project (Grant Agreement No. 101132308), funded by the European Union. To improve the readability of the English language in narrative sections throughout the paper we have used chatGPT (https://chat.openai.com) and Paperpal Preflight (https://preflight.paperpal.com/partner/ieee/access).

REFERENCES

Acharya, A., Sanghavi, S., Jing, L., Bhushanam, B., Choudhary, D., Rabbat, M., and Dhillon, I. (2022). Positive unlabeled contrastive learning. *arXiv preprint arXiv:2206.01206*.

Bekker, J. and Davis, J. (2020). Learning from positive and unlabeled data: A survey. *Machine learning*, 109(4):719–760.

Bekker, J., Robberechts, P., and Davis, J. (2019). Beyond the selected completely at random assumption for learning from positive and unlabeled data. In *Joint European conference on machine learning and knowledge discovery in databases*, pages 71–85. Springer.

Chen, H., Liu, F., Wang, Y., Zhao, L., and Wu, H. (2020a). A variational approach for learning from positive and unlabeled data. *Advances in Neural Information Processing Systems*, 33:14844–14854.

Chen, X., Chen, W., Chen, T., Yuan, Y., Gong, C., Chen, K., and Wang, Z. (2020b). Self-pu: Self boosted and calibrated positive-unlabeled training. In *International Conference on Machine Learning*, pages 1510–1519. PMLR.

Dorigatti, E., Goschenhofer, J., Schubert, B., Rezaei, M., and Bischl, B. (2024). Uncertainty-aware pseudo-label selection for positive-unlabeled learning.

Du Plessis, M. C., Niu, G., and Sugiyama, M. (2014). Analysis of learning from positive and unlabeled data. *Advances in neural information processing systems*, 27.

Elkan, C. and Noto, K. (2008). Learning classifiers from only positive and unlabeled data. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 213–220.

Hammoudeh, Z. and Lowd, D. (2020). Learning from positive and unlabeled data with arbitrary positive shift. *Advances in Neural Information Processing Systems*, 33:13088–13099.

He, F., Liu, T., Webb, G. I., and Tao, D. (2020). Instance-dependent pu learning by bayesian optimal relabeling.

Hsieh, Y.-G., Niu, G., and Sugiyama, M. (2019). Classification from positive, unlabeled and biased negative data. In *International conference on machine learning*, pages 2820–2829. PMLR.

Karamanlis, I., Kokkalis, A., Profillidis, V., Botzoris, G., Kiourt, C., Sevetlidis, V., and Pavlidis, G. (2023). Deep learning-based black spot identification on greek road networks. *Data*, 8(6):110.

Kato, M., Teshima, T., and Honda, J. (2019). Learning from positive and unlabeled data with a selection bias. In *International conference on learning representations*.

Kiryo, R., Niu, G., Du Plessis, M. C., and Sugiyama, M. (2017). Positive-unlabeled learning with non-negative risk estimator. *Advances in neural information processing systems*, 30.

Northcutt, C. G., Wu, T., and Chuang, I. L. (2017). Learning with confident examples: Rank pruning for robust classification with noisy labels. *arXiv preprint arXiv:1705.01936*.

Sevetlidis, V., Pavlidis, G., Balaska, V., Psomoulis, A., Mouroutsos, S., and Gasteratos, A. (2023). Defect detection using weakly supervised learning. In *2023 IEEE International Conference on Imaging Systems and Techniques (IST)*, pages 1–5. IEEE.

Sevetlidis, V., Pavlidis, G., Balaska, V., Psomoulis, A., Mouroutsos, S. G., and Gasteratos, A. (2024a). Enhancing weakly supervised defect detection through anomaly-informed weighted training. *IEEE Transactions on Instrumentation and Measurement*.

Sevetlidis, V., Pavlidis, G., and Gasteratos, A. (2025). Angular regularization for positive-unlabeled learning on the hypersphere. *Transactions on Machine Learning Research*. Featured Certification, J2C Certification.

Sevetlidis, V., Pavlidis, G., Mouroutsos, S. G., and Gasteratos, A. (2024b). Dense-pu: Learning a density-based boundary for positive and unlabeled learning. *IEEE Access*, 12:90287–90298.

Sevetlidis, V., Pavlidis, G., Mouroutsos, S. G., Karakatsanis, T., and Gasteratos, A. (2024c). Positive and unlabelled learning with data approximation for photovoltaic defect detection. In *2024 IEEE International Conference on Imaging Systems and Techniques (IST)*, pages 1–6.

Su, G., Chen, W., and Xu, M. (2021). Positive-unlabeled learning from imbalanced data. In *IJCAI*, pages 2995–3001.

Tang, Z., Pei, S., Zhang, Z., Zhu, Y., Zhuang, F., Hoehndorf, R., and Zhang, X. (2022). Positive-unlabeled learning with adversarial data augmentation for knowledge graph completion. *arXiv preprint arXiv:2205.00904*.

Wang, H., Xiao, R., Li, Y., Feng, L., Niu, G., Chen, G., and Zhao, J. (2022). Pico: Contrastive label disambiguation for partial label learning. *ICLR*, 1(2):5.

Yuan, B., Gong, C., Tao, D., and Yang, J. (2025). Weighted contrastive learning with hard negative mining for positive and unlabeled learning. *IEEE Transactions on Neural Networks and Learning Systems*.

Zhao, Y., Xu, Q., Jiang, Y., Wen, P., and Huang, Q. (2022). Dist-pu: Positive-unlabeled learning from a label distribution perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14461–14470.

Zheng, Y., Peng, H., Zhang, X., Zhao, Z., Gao, X., and Li, J. (2019). Ddi-pulearn: a positive-unlabeled learning method for large-scale prediction of drug-drug interactions. *BMC bioinformatics*, 20(Suppl 19):661.

Žunkovič, B. (2023). Positive unlabeled learning with tensor networks. *Neurocomputing*, 552:126556.