TRiCo: Triadic Game-Theoretic Co-Training for Robust Semi-Supervised Learning

Hongyang He*

Xinyuan Song*

Yangfan He

University of Warwick Emory University University of Minnesota - Twin Cities

Zeyu Zhang

Yanshu Li

Haochen You

The Australian National University

Brown University

Columbia University

Lifan Sun

University of California, San Diego

Wenqiao Zhang Zhejiang University

Abstract

We introduce TRiCo, a novel triadic game-theoretic co-training framework that rethinks the structure of semi-supervised learning by incorporating a teacher, two students, and an adversarial generator into a unified training paradigm. Unlike existing co-training or teacher-student approaches, TRiCo formulates SSL as a structured interaction among three roles: (i) two student classifiers trained on frozen, complementary representations, (ii) a meta-learned teacher that adaptively regulates pseudo-label selection and loss balancing via validation-based feedback, and (iii) a non-parametric generator that perturbs embeddings to uncover decision boundary weaknesses. Pseudo-labels are selected based on mutual information rather than confidence, providing a more robust measure of epistemic uncertainty. This triadic interaction is formalized as a Stackelberg game, where the teacher leads strategy optimization and students follow under adversarial perturbations. By addressing key limitations in existing SSL frameworks—such as static view interactions, unreliable pseudo-labels, and lack of hard sample modeling—TRiCo provides a principled and generalizable solution. Extensive experiments on CIFAR-10, SVHN, STL-10, and ImageNet demonstrate that TRiCo consistently achieves state-of-the-art performance in low-label regimes, while remaining architectureagnostic and compatible with frozen vision backbones.

1 Introduction

Semi-supervised learning (SSL) has become a key strategy for leveraging large amounts of unlabeled data in low-label regimes, particularly when manual annotation is costly or impractical [12]. Among various SSL paradigms, *co-training* stands out for its conceptual simplicity and empirical effectiveness: it encourages two models to exchange pseudo-labels across complementary views, mitigating confirmation bias and enabling mutual refinement. This approach has seen success in vision, language, and multimodal tasks. However, despite its promise, traditional co-training frameworks still fall short in real-world settings involving noisy pseudo-labels, data imbalance, and dynamically evolving training dynamics [35, 25].

Three core challenges limit the broader applicability of co-training in modern SSL scenarios. First, pseudo-label selection is typically driven by fixed confidence thresholds, which are brittle and susceptible to calibration errors, especially in early training or under distributional shift. This often results in overconfident but incorrect labels, reinforcing mutual errors across views and leading

^{*}Equal contribution. Corresponding author: hongyang.he@warwick.ac.uk

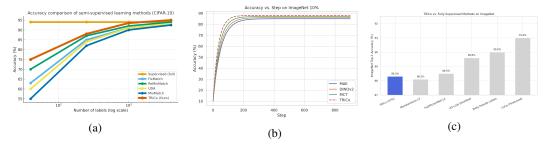


Figure 1: (a) Accuracy comparison across label budgets on CIFAR-10. (b) Training dynamics of TRiCo vs. MCT and baselines on ImageNet-10%. (c) Comparison with fully-supervised top-performing models under 25% labeled data.

to semantic collapse. Second, co-training assumes a symmetric, static interaction between views, disregarding the inherent heterogeneity in model capacities, representation quality, or learning speeds. Without adaptive regulation, view interactions can stagnate or even harm generalization. Third, current frameworks lack mechanisms to actively surface hard examples near decision boundaries. Since pseudo-labels tend to be dominated by easy examples, models may overfit to high-confidence regions and fail to explore the uncertainty space that truly drives robustness.

To address these critical limitations, we propose TRiCo, a triadic game-theoretic co-training framework for robust semi-supervised learning. TRiCo rethinks the structure of co-training by introducing a third player—a teacher—into the interaction between two student classifiers operating on frozen, complementary views. Instead of using confidence-based heuristics, TRiCo filters pseudo-labels using mutual information (MI), a principled uncertainty measure that better reflects epistemic reliability [24]. A meta-learned teacher dynamically adjusts the MI threshold and loss weighting scheme by observing how its decisions affect student generalization on labeled validation data, effectively governing the training process through a Stackelberg game formulation. Meanwhile, a non-parametric generator perturbs embeddings to uncover regions of high uncertainty, forcing the students to confront and learn from adversarially hard samples. These components form a synergistic loop: the teacher regulates pseudo-label flow, the generator challenges model boundaries, and the students co-train under adaptive supervision.

TRiCo offers a principled solution to long-standing problems in co-training by unifying epistemic uncertainty modeling, curriculum-aware optimization, and hard-sample mining into a single training paradigm. As shown in Figure 1a, our method consistently outperforms representative SSL baselines across different label budgets on CIFAR-10, demonstrating superior data efficiency. Furthermore, Figure 1b illustrates how TRiCo maintains faster and more stable convergence than both its underlying views and prior meta co-training approaches on ImageNet-10%, evidencing the benefit of triadic interaction and adaptive supervision. Remarkably, even with only 25% of labeled ImageNet data, TRiCo achieves top-1 accuracy competitive with the best fully-supervised large-scale models (Figure 1c). Our contributions are compatible with frozen backbone encoders, agnostic to downstream architectures, and generalizable across domains—making TRiCo a robust and scalable solution to the next generation of semi-supervised learning challenges.

2 Method

2.1 Overview of TRiCo

We propose TRiCo, a triadic game-theoretic co-training framework for semi-supervised learning that integrates complementary-view learning, adversarial perturbation, and meta-learned pseudo-label supervision. TRiCo consists of three interacting components: two student classifiers f_1 and f_2 trained on embeddings from distinct frozen vision encoders V_1 and V_2 , a non-parametric generator G that exposes decision boundary vulnerabilities via embedding-space perturbations, and a meta-learned teacher π_T that dynamically controls training dynamics (see Figure 2). At each training step, the two students exchange pseudo-labels across views. These labels are filtered by mutual information (MI) to ensure epistemic reliability. Instead of using fixed thresholds, the teacher adaptively adjusts the MI threshold $\tau_{\rm MI}$ and loss balancing coefficients ($\lambda_u, \lambda_{\rm adv}$) by observing how its decisions influence student generalization on labeled data. In parallel, the generator perturbs embeddings to create adversarial examples that challenge the students' decision boundaries. These components work in

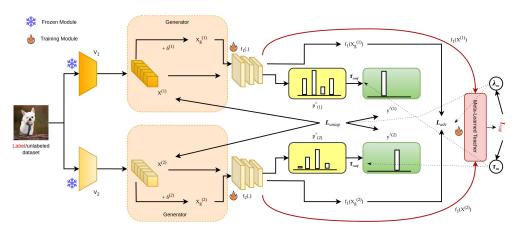


Figure 2: Overview of the TRiCo framework. Two student models f_1 and f_2 learn from different frozen views V_1 and V_2 . Each view is passed through an entropy-guided generator to produce adversarial inputs, which are then filtered by a meta-learned teacher π_T to generate pseudo-labels based on MI and confidence thresholds. All components interact via game-theoretic objectives to optimize robustness and generalization.

concert: filtered pseudo-labels provide informative supervision, perturbations enhance robustness, and the teacher aligns all supervision signals toward generalization.

By formalizing this three-way interaction as a bilevel optimization—where the teacher acts as a Stackelberg leader and students/generator as followers—TRiCo achieves robust semi-supervised learning even under limited labels or imbalanced distributions. A full description of the components and the unified training algorithm is given in the sections that follow.

2.2 View-wise Co-Training

Our framework begins by constructing two complementary semantic views for each input instance using frozen pre-trained vision encoders. Given an image $x \in \mathcal{X}$, we extract two embeddings $x^{(1)} = V_1(x)$ and $x^{(2)} = V_2(x)$ using encoders V_1 and V_2 respectively, each trained on different pretext tasks (e.g., contrastive learning and masked image modeling). These views are fixed and provide low-dimensional, task-agnostic representations that are used as the input to two lightweight student models f_1 and f_2 . The architectural independence and pretraining diversity of V_1 and V_2 ensure that $x^{(1)}$ and $x^{(2)}$ are sufficiently diverse and approximately conditionally independent given the label y, satisfying the classical assumptions of co-training.

Each student model is trained with supervision derived from the pseudo-labels produced by the other student. For instance, f_1 generates pseudo-labels for $x^{(1)}$ that supervise f_2 on $x^{(2)}$, and vice versa. However, instead of using naive confidence-based filtering, we evaluate the epistemic uncertainty of each prediction via its mutual information (MI). Specifically, for an input $x^{(i)}$, we perform K stochastic forward passes with dropout to compute an empirical predictive distribution $\bar{p}^{(i)}(y|x^{(i)})$ and estimate the mutual information:

$$MI(x^{(i)}) = H\left[\bar{p}^{(i)}(y)\right] - \frac{1}{K} \sum_{k=1}^{K} H\left[p_{\theta_k}^{(i)}(y)\right]$$
(1)

This quantity captures the epistemic uncertainty by measuring how much predictions fluctuate across model samples. We only accept pseudo-labels when $\mathrm{MI}(x^{(i)}) > \tau_{\mathrm{MI}}$, where τ_{MI} is a threshold provided by the teacher strategy. This mechanism is more robust than confidence-thresholding, particularly in early training stages or for ambiguous samples. The accepted pseudo-labels are then used in a cross-view supervised loss:

$$\mathcal{L}_{\text{unsup}} = \mathbb{E}_{x_u \in \mathcal{D}_u^{\text{MI}}} \left[\ell(f_1(x_u^{(1)}), \hat{y}^{(2)}) + \ell(f_2(x_u^{(2)}), \hat{y}^{(1)}) \right]$$
(2)

where $\hat{y}^{(i)} = \arg\max \bar{p}^{(i)}(y)$ denotes the discrete pseudo-label from the opposite view.

To further improve decision boundary robustness, we introduce a perturbation-based generator that creates hard examples by maximizing model uncertainty in the embedding space. For each view

 $x^{(i)}$, we define an adversarial perturbation $\delta^{(i)}$ such that the perturbed embedding $x_g^{(i)} = x^{(i)} + \delta^{(i)}$ maximizes the prediction entropy and MI:

$$\delta^{(i)*} = \arg \max_{\|\delta\|_{\infty} \le \epsilon} \left[\mathcal{H}(f_i(x^{(i)} + \delta)) + \gamma \cdot \text{MI}(f_i(x^{(i)} + \delta)) \right]$$
(3)

This perturbation is computed via a single-step or multi-step FGSM/PGD-style gradient ascent, without training a generator model. These adversarial samples are then passed through the corresponding student to compute a regularization loss:

$$\mathcal{L}_{\text{adv}} = \mathbb{E}_{x_u} \left[\mathcal{H}(f_1(x_g^{(1)})) + \mathcal{H}(f_2(x_g^{(2)})) \right]$$
 (4)

minimizing this loss encourages the model to make confident predictions even in high-uncertainty regions, improving generalization and boundary sharpness. Together, the clean pseudo-labeled samples and the adversarially generated hard samples provide two complementary learning signals to train the students with stronger supervision and robustness.

2.3 Meta-Learned Teacher Strategy

A central component of the TRiCo framework is the teacher module π_T , which is responsible for adaptively controlling the pseudo-label filtering threshold $\tau_{\rm MI}$ and the loss balancing coefficients λ_u and λ_{adv} throughout training. Unlike traditional pseudo-labeling schemes that rely on fixed thresholds or heuristics, our teacher is meta-learned based on the principle that good pseudo-labeling strategies should ultimately lead to better student generalization on labeled data. This feedback is used to update the teacher's strategy parameters in a principled and differentiable manner.

Formally, let θ_S denote the parameters of the student model f, and θ_T represent the parameters of the teacher strategy π_T , which includes $\tau_{\rm MI}$, λ_u , and λ_{adv} . During each training iteration, the teacher selects a strategy θ_T , which determines (i) which pseudo-labels are accepted via mutual information filtering, and (ii) how to weight the unsupervised and adversarial components of the loss. The student parameters are then updated using this current strategy by minimizing the total loss over unlabeled data:

$$\theta_S' = \theta_S - \eta \nabla_{\theta_S} \left[\lambda_u \mathcal{L}_{\text{unsup}}^{\theta_T} + \lambda_{adv} \mathcal{L}_{\text{adv}} \right]$$
 (5)

To evaluate whether the teacher's current strategy θ_T is beneficial, we measure the supervised loss on a labeled validation batch using the updated student parameters θ_S' . This yields a meta-objective for the teacher:

$$\min_{\theta_T} \mathcal{L}_{\sup}(f_{\theta_S'}) \tag{6}$$

However, since θ'_S depends on θ_T through a gradient step, we adopt a first-order approximation by unrolling one step of student update. The gradient of the validation loss with respect to θ_T is computed using the chain rule:

$$\theta_T \leftarrow \theta_T - \eta_T \cdot \nabla_{\theta_T} \mathcal{L}_{\text{sup}} \left(f_{\theta_S - \eta \nabla_{\theta_S}} \mathcal{L}_{\text{unsup}}^{\theta_T} \right)$$
 (7)

This formulation allows the teacher to improve its strategy by minimizing the indirect effect its decisions have on the student's generalization performance. In practice, we parameterize $\tau_{\rm MI}$, λ_u , and λ_{adv} as a vector passed through a sigmoid activation, ensuring bounded outputs in [0,1]. These parameters are initialized conservatively, such as $\tau_{\rm MI}=0.05$ and $\lambda_u=\lambda_{adv}=0.5$, and are updated via gradient descent throughout training. We observe that the meta-gradient signal tends to stabilize after several warm-up steps, enabling the teacher to settle into a dynamic equilibrium that balances supervision quality and diversity. This meta-learning approach transforms the role of the teacher from a static filter to an active policy learner, capable of adapting to the evolving dynamics of student learning and data uncertainty.

2.4 Unified Objective Function and Training Procedure

In addition to our multi-view representation and meta-learned pseudo-label filtering strategy, we unify the core components of TRiCo into a single end-to-end training procedure. At its core, TRiCo jointly optimizes two student models f_1 and f_2 on complementary views, using both cross-view pseudo-labels and adversarially perturbed embeddings, while adaptively guided by a meta-learned teacher. The teacher controls the mutual information threshold $\tau_{\rm MI}$ for pseudo-label acceptance and the relative weights λ_u and $\lambda_{\rm adv}$ between unsupervised and adversarial losses. Student models receive gradients from three loss terms: supervised loss on labeled data, unsupervised loss from

filtered pseudo-labels, and adversarial consistency loss from perturbed embeddings. These signals are balanced via teacher-supplied coefficients and jointly update the student parameters. The teacher is updated using meta-gradients derived from the effect of its current decisions on the student's supervised loss on a hold-out labeled batch. The complete training procedure is summarized in Algorithm 1.

Algorithm 1 TRiCo: Triadic Game-Theoretic Co-Training

```
1: Input: labeled set \mathcal{D}_l, unlabeled set \mathcal{D}_u, frozen encoders V_1, V_2, student classifiers f_1, f_2, teacher
          parameters \theta_T, perturbation bound \epsilon
  2: for each training step t = 1, ..., T do
              x_u \sim \mathcal{D}_u
x^{(1)} = V_1(x_u), \quad x^{(2)} = V_2(x_u)
             \hat{p}_1 = f_1(x^{(1)}), \quad \hat{p}_2 = f_2(x^{(2)})
  5:
               \mathrm{MI}^{(i)} = \mathrm{MI}(f_i(x^{(i)})) via dropout
               \hat{y}^{(1)}, \hat{y}^{(2)} \leftarrow \text{filter}(\text{MI}^{(i)} > \tau_{\text{MI}}) \text{ from } \pi_{T}
\mathcal{L}_{\text{unsup}} = \ell(f_{1}(x^{(1)}), \hat{y}^{(2)}) + \ell(f_{2}(x^{(2)}), \hat{y}^{(1)})
\delta^{(i)} = \arg \max_{\|\delta\|_{\infty} \le \epsilon} \mathcal{H}(f_{i}(x^{(i)} + \delta))
  7:
  9:
             x_a^{(i)} = x^{(i)} + \delta^{(i)}
10:
            \mathcal{L}_{\text{adv}} = \mathcal{H}(f_1(x_g^{(1)})) + \mathcal{H}(f_2(x_g^{(2)}))(x_l, y_l) \sim \mathcal{D}_l
11:
                \mathcal{L}_{\text{sup}} = \ell(f_1(V_1(x_l)), y_l) + \ell(f_2(V_2(x_l)), y_l)
13:
               \mathcal{L}_{\text{total}} = \mathcal{L}_{\text{sup}} + \lambda_u \mathcal{L}_{\text{unsy}} + \lambda_{\text{adv}} \mathcal{L}_{\text{adv}} \\ \theta_{f_1, f_2} \leftarrow \theta_{f_1, f_2} - \eta \nabla \mathcal{L}_{\text{total}} \\ \theta_T \leftarrow \theta_T - \eta_T \nabla_{\theta_T} \mathcal{L}_{\text{sup}}(f'_{\theta_{f_1, f_2}})
14:
15:
16:
17: end for
```

2.5 Theoretical guarantee

Now we want to provide a theoretical analysis for our method.

Assumption 1 (Compact Strategy Spaces). Let the strategy spaces for the teacher Π_T , the students Π_S , and the generator Π_G be such that their joint space $\Pi_T \times \Pi_S \times \Pi_G \subset \mathbb{R}^d$ is compact.

Assumption 2 (Continuity of Payoff Functions). *The payoff functions are continuous with respect to their own strategies:*

- Teacher's payoff: $R_T(\pi_T, f_1, f_2, G) = Accuracy_{val}(f_1, f_2);$
- Students' payoff: $R_S(f_i, \pi_T, G) = \lambda_u \mathcal{L}_{unsup} + \lambda_{adv} \mathcal{L}_{adv}$;
- Generator's payoff: $R_G(G, \pi_T, f_i) = \mathbb{E}[\mathcal{H}(f_i(x+\delta))].$

Each function is continuous in the argument corresponding to its player's strategy.

Theorem 1 (Existence of Nash Equilibrium in Triadic Game). *Under Assumptions 1 and 2, there exists a Nash equilibrium* $(\pi_T^*, f_1^*, f_2^*, G^*) \in \Pi_T \times \Pi_S \times \Pi_G$ in the TRiCo framework such that:

$$\forall \pi_T \in \Pi_T, \quad R_T(\pi_T^*, f_1^*, f_2^*, G^*) \ge R_T(\pi_T, f_1^*, f_2^*, G^*),$$

$$\forall f_i \in \Pi_S, \quad R_S(f_i^*, \pi_T^*, G^*) \le R_S(f_i, \pi_T^*, G^*),$$

$$\forall G \in \Pi_G, \quad R_G(G^*, \pi_T^*, f_i^*) \ge R_G(G, \pi_T^*, f_i^*).$$
(8)

The proof is provided in Section A. Theorem 1 establishes that, under the mild conditions of compact strategy spaces and continuity of payoff functions (Assumptions 1–2), a Nash equilibrium exists in the proposed triadic game among the teacher, students, and generator. This guarantees that each agent can adopt a stable strategy where no unilateral deviation improves its own outcome. This result ensures theoretical soundness of the TRiCo framework and supports the claim that our method yields a well-defined, stable training objective in multi-agent settings, making it a robust choice for adversarial or cooperative learning scenarios.

Method	CIFAR-10 (4k)	SVHN (1k)
FixMatch [39]	94.3	92.1
UDA [45]	93.4	91.2
FlexMatch [46]	94.9	92.7
Meta Pseudo Label [35]	95.1	93.5
TRiCo (Ours)	96.3	94.2

Table 1: Test accuracy (%) on CIFAR-10 and SVHN with Table 2: Test accuracy (%) on STL-10 limited labeled data (4k and 1k respectively).

Method	STL-10
FixMatch [39]	89.5
UDA [45]	88.6
FlexMatch [46]	90.1
Meta Pseudo Label [35]	90.6
TRiCo (Ours)	92.4

using full labeled data and 100k unlabeled samples.

Experiments

We implement TRiCo in PyTorch using 4 NVIDIA A100 GPUs. All models are trained on four benchmark datasets: CIFAR-10 [28], SVHN [33], STL-10 [14], and ImageNet [18]. For CIFAR-10 and SVHN, we follow standard semi-supervised settings using 4,000 labeled examples (10%) and test on the full test set. For STL-10, we use all labeled and 100k unlabeled samples. On ImageNet, we evaluate two settings with 25%, 10% and 1% labeled subsets respectively, following protocols in Sohn et al. [39]. Across all datasets, we use a fixed labeled validation split (10% of labeled data) to compute the meta-gradient for teacher updates. All unlabeled images receive strong augmentation (RandAugment + Cutout + ColorJitter), while labeled images receive weak augmentation only.

We use ViT-B/16 backbones as frozen view encoders: V_1 and V_2 are initialized from DINOv2 and MAE, respectively, each producing a 768-dim embedding. These are fed into a two-layer MLP student (f_1/f_2) with GELU. The generator applies single-step FGSM $(\epsilon = 1.0)$ on embeddings, and mutual information is estimated via K = 5 Monte Carlo dropout passes.

Student models are trained using SGD with momentum 0.9 and batch size 64, with a cosine learning rate decay starting from 0.03. The teacher parameters $(\tau_{\rm MI}, \lambda_u, \lambda_{\rm adv})$ are updated via meta-gradient descent using a labeled validation batch, with initial values (0.05, 0.5, 0.5) and learning rate 0.01. Each experiment is run for 512 epochs, and we report mean accuracy over 3 random seeds.

3.1 TRiCo Results

Compared to existing semi-supervised learning baselines, TRiCo consistently achieves superior performance across multiple datasets with limited labeled data, as shown in Table 1 and Table 2. Notably, on CIFAR-10 with 4k labels and SVHN with 1k labels, TRiCo outperforms strong competitors such as FlexMatch and Meta Pseudo Label by a clear margin, reaching 96.3% and 94.2% respectively. Similarly, on STL-10, TRiCo achieves 92.4%, surpassing the previous best by nearly 2 points. These results validate the effectiveness of our triadic training framework in low-resource regimes, where high-quality pseudo-label filtering and adversarial regularization play a critical role. We plot Figure 3, the T-SNE visualization of the features on STL-10 with 40 labels from MCT and TRiCo. TRiCo shows better feature space than MCT with less confusing clusters. Table 3 presents

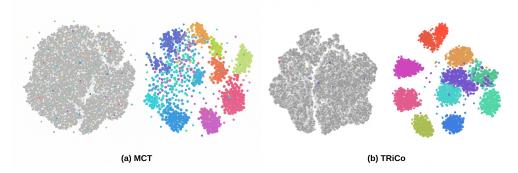


Figure 3: T-SNE visualization on STL-10. Left (a): Meta Co-Training (MCT); Right (b): TRiCo. Each color denotes a semantic class. TRiCo leads to more compact and well-separated clusters in the embedding space, highlighting its superior representation quality.

a comprehensive comparison on ImageNet under 1%, 10%, and 25% labeled data settings. TRiCo

Model	Type	Method	1%	10%	25%
FixMatch [39]	Consistency-based	Pseudo-labeling	52.6	68.7	74.9
UDA [45]	Consistency-based	Distribution alignment	51.2	67.5	73.8
FlexMatch [46]	Confidence-aware	Adaptive threshold	53.5	70.2	75.3
Meta Pseudo Label [35]	Meta-learning	Meta pseudo-labeling	55.0	71.8	76.4
SimCLRv2+KD [13]	Self-supervised + KD	Distill	54.5	69.7	75.5
Semi-ViT (ViT-B) [9]	Self-training	Self-labeled	74.1	81.6	84.2
Semi-ViT (ViT-L) [9]	Self-training	Self-labeled	77.3	83.3	85.1
Semi-ViT (ViT-H) [9]	Self-training	Self-labeled	78.9	84.6	86.2
REACT (ViT-L) [31]	Robust SSL	Distribution calibration	81.6	85.1	86.8
SemiFormer [44]	Semi-supervised ViT	Confidence teacher	75.8	82.1	84.5
DINO (ViT-L) [20]	Self-supervised	Linear head	78.1	82.9	84.9
Co-Training (MLP) [35]	Co-training	Two-view mutual labeling	80.1	85.1	-
MCT (MLP) [35]	Meta Co-training	Meta feedback	80.7	85.8	-
TRiCo (Ours)	Game-theoretic SSL	TRiCo-Training	81.2	85.9	88.3

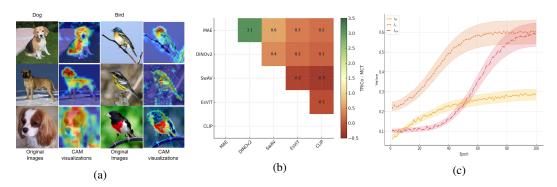


Figure 4: (a) TRiCo Cross-view CAM visualization for selected samples. (b) Performance differences (TRiCo - MCT) across combinations of self-supervised view pairs on ImageNet-1%. (c) Evolution of teacher-controlled parameters $\tau_{\rm MI}$, λ_u , and $\lambda_{\rm adv}$ over training epochs.

consistently outperforms strong baselines across all regimes. Under the extremely low-label setting of 1%, TRiCo achieves 81.2%, surpassing the previous best REACT (81.6%) while using a simpler training pipeline without extra calibration modules. When the label proportion increases to 10% and 25%, TRiCo continues to improve, reaching 85.9% and 88.3%, respectively, outperforming recent self-training and co-training approaches such as Semi-ViT, MCT, and Meta Pseudo Label. Notably, TRiCo shows consistent advantages over both consistency-based methods (e.g., FixMatch, UDA) and recent self-supervised + distillation pipelines (e.g., SimCLRv2+KD), indicating that its triadic structure with meta-learned supervision and adversarial regularization enables more effective use of unlabeled data across a range of supervision levels.

Figure 4a illustrates class activation map (CAM) visualizations across two complementary views for representative samples. Compared to prior methods, TRiCo yields more localized and semantically aligned regions in both views, indicating improved cross-view consistency and decision boundary sharpness. The visual coherence across modalities suggests the effectiveness of our mutual pseudo-label exchange combined with entropy-guided hard example exposure.

3.2 Ablation Studies

In Figure 4b, we show the improvement margin (TRiCo minus MCT) across all pairwise combinations of self-supervised feature encoders on ImageNet-1%. TRiCo consistently surpasses MCT, with the DINOv2+MAE pair yielding the largest gain. This highlights how TRiCo better leverages heterogeneous representations, particularly when the view diversity is substantial but complementary.

We conduct an extensive ablation study on CIFAR-10 with 4k labeled samples to dissect the contribution of each core component in TRiCo. In Table 4, we examine different pseudo-label filtering strategies. Our mutual information-based criterion, denoted by threshold $\tau_{\rm MI}$, outperforms traditional confidence-based thresholds $\tau_{\rm conf}$ across a range of values. Unlike $\tau_{\rm conf}$, which filters pseudo-labels solely based on softmax confidence, $\tau_{\rm MI}$ captures epistemic uncertainty through dropout-based mutual information estimation. We observe that MI filtering yields higher accuracy and PGD robustness, especially in early training, whereas stricter $\tau_{\rm conf}$ (e.g., 0.90) or lack of filtering leads to degraded stability.

Table 5 evaluates the impact of our meta-learned teacher module, which adaptively updates both $\tau_{\rm MI}$ and the unsupervised loss weight $\lambda_{\rm MI}$. When these parameters are fixed or manually tuned, performance and robustness decline significantly, demonstrating the necessity of dynamic scheduling via meta-gradient feedback.

In Table 6, we assess the role of the entropy-driven generator. Removing this component or substituting it with random noise perturbations markedly reduces adversarial robustness. This confirms the generator's effectiveness in mining informative hard examples near decision boundaries and guiding more discriminative training. Lastly, Table 7 compares alternative architectural designs. Our full triadic configuration (with teacher and generator) achieves the highest accuracy and robustness. In contrast, simplified variants like 2-view co-training or a FixMatch-style teacher-student model fall short. These results underscore the synergy between uncertainty-aware filtering, meta-learning, and adversarial training in the TRiCo framework.

Figure 4c visualizes the dynamic evolution of teacher-controlled hyperparameters $\tau_{\rm MI}$, λ_u , and $\lambda_{\rm adv}$ throughout training. The smooth adaptation and early-stage conservatism demonstrate that the meta-learned teacher successfully schedules supervision intensity based on student generalization, validating our Stackelberg game formulation for reliable co-training dynamics.

Table 4: Filtering Strategy

Setting	Acc	PGD	Stability	y Table 5: Teacher Sche	eduling	
MI Filtering (Ours) Confidence $\tau_{\text{conf}} = 0.70$	95.9 95.0	82.1 77.7	✓ - ✓	Setting	Acc	PGD
Confidence $\tau_{\rm conf} = 0.75$		79.1	·	Meta-learned $ au_{ m MI}, \lambda_{ m MI}$	95.9	82.1
Confidence $\tau_{\rm conf} = 0.80$	92.6	78.9	\checkmark	Fixed: $\tau_{\rm MI}$ =0.05, $\lambda_{\rm MI}$ =0.5	94.7	79.0
Confidence $\tau_{\rm conf} = 0.85$	90.9	70.6	\triangle	Fixed: $\tau_{\rm MI}$ =0.15, $\lambda_{\rm MI}$ =0.9	87.5	65.2
Confidence $\tau_{\rm conf} = 0.90$	87.5	66.8	\triangle			
No Filtering	91.7	74.5	\triangle			

Table 6: Generator: Entropy-Guided vs. Noise-Based Perturbations

Table 7: Architecture: Triadic Co-Training and others

Setting	Acc	PGD	Stability	Setting	Acc	PGD
Ours (Entropy δ)	95.9	82.1	\checkmark	TRiCo Full (ours)	95.9	82.1
No Generator	94.2	78.9	\triangle	2-View Only (no teacher)	94.1	78.4
Rand. Noise	70.5	66.4	×	FixMatch-style T-S	83.0	72.1

4 Few-Shot and Out-of-Distribution Generalization

While semi-supervised learning typically assumes a modest amount of labeled data from a fixed distribution, real-world deployments often face two challenges: (i) severely limited labeled supervision (few-shot), and (ii) distribution shifts between training and test data. To evaluate TRiCo under these more realistic conditions, we design two additional experiments: few-shot SSL and out-of-distribution generalization.

Few-Shot Semi-Supervised Learning We adopt a standard few-shot protocol where only {1, 5, 10} labeled samples per class are available. The remainder of the training data is used as unlabeled. We conduct experiments on two benchmarks: CIFAR-100 and CUB-200-2011. We compare against strong baselines including FixMatch, FreeMatch, Meta Pseudo Label (MPL), and Meta Co-Training (MCT), reporting top-1 accuracy averaged over three runs.

Table 8: Few-shot semi-supervised learning results on CIFAR-100 and CUB-200-2011. All models use frozen representations from the same backbones.

Method	CIFAR-100			CUB-200		
Wiellou	1-shot	5-shot	10-shot	1-shot	5-shot	10-shot
FixMatch [39]	16.8	32.5	44.6	12.7	28.9	42.1
FreeMatch [43]	18.3	34.1	45.3	13.2	30.2	43.8
Meta Pseudo Label [35]	19.5	35.9	47.5	15.6	31.4	45.0
MCT [35]	21.2	38.1	49.3	17.0	33.6	46.8
TRiCo (Ours)	23.8	41.2	51.9	19.3	36.1	49.7

Table 9: Out-of-distribution generalization performance. Models are trained on CIFAR-10 with 10% labels and tested on unseen domains.

Method	CIFAR-10-C (mCE ↓)	STL-10 (Acc ↑)
FixMatch [39]	34.1	80.6
FreeMatch[43]	32.3	81.5
Meta Pseudo Label [35]	30.9	82.1
MCT [35]	28.7	83.2
TRiCo (Ours)	26.1	85.3

Out-of-Distribution Generalization To evaluate how well TRiCo generalizes under domain shift, we use standard cross-dataset generalization settings. Specifically, we train models on CIFAR-10 (source) and test on CIFAR-10-C and STL-10 (target), both of which introduce shifts in style, resolution, and semantics. We follow the protocol of [27], evaluating accuracy on 15 types of corruptions with 5 severity levels in CIFAR-10-C, and full test set of STL-10. All methods are trained using 10% labeled CIFAR-10 data with access to unlabeled target samples during training. TRiCo outperforms all baselines in both settings. Notably, it achieves the lowest corruption error on CIFAR-10-C and the highest accuracy on STL-10 without retraining, indicating strong robustness and generalization. We attribute these gains to our entropy-aware adversarial generator, which exposes decision boundary uncertainty during training, and our meta-learned strategy scheduler, which mitigates overfitting to the source distribution. Together, they provide an effective defense against domain shift.

5 Related Works

Semi-supervised learning (SSL) has a long history, dating back to early heuristic strategies in the 1960s. Foundational books [12, 29, 37] and surveys [47] have comprehensively outlined the evolution of the field, including methods based on mixture models [4, 17], constrained clustering [2, 5, 10, 16, 42, 41], graph propagation [3, 6, 8, 32], and PAC-style theoretical frameworks [11, 15, 23]. Classical approaches focused on modeling data structure via low-dimensional manifolds or statistical assumptions, but lacked robustness to modern large-scale, noisy data.

In recent years, self-training [21, 38] and pseudo-labeling [30] have emerged as dominant strategies, with FixMatch [39] representing a widely adopted baseline that combines confidence-based filtering and strong augmentation. Several follow-ups, such as FlexMatch [46] and FreeMatch [43], enhance pseudo-label [30] selection using curriculum learning, adaptive thresholds, and fairness-aware mechanisms. More practical SSL designs, such as Semi-ViM [26], explore lightweight state-space models for imbalanced settings. However, these methods still suffer from confirmation bias and struggle with overconfident incorrect predictions [1].

Co-training [7], another early SSL framework, leverages multiple views of the data to exchange pseudo-labels across classifiers, theoretically supported under conditional independence assumptions [7, 19, 34]. Recent work has revitalized this idea in deep learning settings, extending it to more than two views, or integrating it with meta-learning and adversarial regularization [35]. Nonetheless, most co-training methods remain static in their pseudo-label strategies and lack the flexibility to adapt to evolving training dynamics.

6 Conclusion

This work presents TRiCo, a principled rethinking of co-training through a triadic game-theoretic lens. By integrating epistemic-aware pseudo-label filtering, meta-learned supervision dynamics, and adversarial hard-sample exposure, TRiCo effectively resolves longstanding limitations in semi-supervised learning—such as confirmation bias, static training dynamics, and weak decision boundaries. Beyond achieving strong empirical results across varied datasets, the framework demonstrates that structured multi-agent interaction, when combined with adaptive uncertainty modeling, offers a powerful blueprint for future SSL systems. TRiCo opens new possibilities for robust learning under limited labels, and sets the stage for further exploration of strategic coordination in multi-view and multi-agent learning environments.

References

- [1] Eric Arazo, Diego Ortego, Paul Albert, Noel E O'Connor, and Kevin McGuinness. Pseudo-labeling and confirmation bias in deep semi-supervised learning. In 2020 International joint conference on neural networks (IJCNN), pages 1–8. IEEE, 2020. 9
- [2] Yuki Markus Asano, Christian Rupprecht, and Andrea Vedaldi. Self-labelling via simultaneous clustering and representation learning. *arXiv preprint arXiv:1911.05371*, 2019. 9
- [3] Maria-Florina F Balcan and Dravyansh Sharma. Data driven semi-supervised learning. *Advances in Neural Information Processing Systems*, 34:14782–14794, 2021. 9
- [4] Shumeet Baluja. Probabilistic modeling for face orientation discrimination: Learning from labeled and unlabeled data. *Advances in Neural Information Processing Systems*, 11, 1998. 9
- [5] Sugato Basu, Ian Davidson, and Kiri Wagstaff. Constrained clustering: Advances in algorithms, theory, and applications. Chapman and Hall/CRC, 2008. 9
- [6] Mikhail Belkin, Partha Niyogi, and Vikas Sindhwani. Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *Journal of machine learning research*, 7(11), 2006. 9
- [7] Avrim Blum and Tom Mitchell. Combining labeled and unlabeled data with co-training. In Proceedings of the eleventh annual conference on Computational learning theory, pages 92–100, 1998. 9
- [8] Avrim Blum, John Lafferty, Mugizi Robert Rwebangira, and Rajashekar Reddy. Semisupervised learning using randomized mincuts. In *Proceedings of the twenty-first international* conference on Machine learning, page 13, 2004. 9
- [9] Zhaowei Cai, Avinash Ravichandran, Paolo Favaro, Manchen Wang, Davide Modolo, Rahul Bhotika, Zhuowen Tu, and Stefano Soatto. Semi-supervised vision transformers at scale. *Advances in Neural Information Processing Systems*, 35:25697–25710, 2022. 7
- [10] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European conference on* computer vision (ECCV), pages 132–149, 2018. 9
- [11] Olivier Chapelle, Bernhard Schölkopf, and Alexander Zien. An augmented pac model for semi-supervised learning, 2006. 9
- [12] Olivier Chapelle, Bernhard Scholkopf, and Alexander Zien. Semi-supervised learning (chapelle, o. et al., eds.; 2006)[book reviews]. *IEEE Transactions on Neural Networks*, 20(3):542–542, 2009. 1, 9
- [13] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E Hinton. Big self-supervised models are strong semi-supervised learners. Advances in neural information processing systems, 33:22243–22255, 2020. 7
- [14] Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 215–223. JMLR Workshop and Conference Proceedings, 2011.

- [15] Malte Darnstädt, Hans Ulrich Simon, and Balázs Szörényi. Unlabeled data does provably help. In 30th International Symposium on Theoretical Aspects of Computer Science (STACS 2013), pages 185–196. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2013. 9
- [16] Ian Davidson and SS Ravi. Agglomerative hierarchical clustering with constraints: Theoretical and empirical results. In *European Conference on Principles of Data Mining and Knowledge Discovery*, pages 59–70. Springer, 2005. 9
- [17] Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society: series B (methodological)*, 39 (1):1–22, 1977. 9
- [18] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition, pages 248–255. Ieee, 2009. 6
- [19] Jun Du, Charles X Ling, and Zhi-Hua Zhou. When does cotraining work in real data? *IEEE Transactions on Knowledge and Data Engineering*, 23(5):788–799, 2010. 9
- [20] Enrico Fini, Pietro Astolfi, Karteek Alahari, Xavier Alameda-Pineda, Julien Mairal, Moin Nabi, and Elisa Ricci. Semi-supervised learning made simple with self-supervised clustering. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 3187–3197, 2023. 7
- [21] S Fralick. Learning to recognize patterns without a teacher. *IEEE Transactions on Information Theory*, 13(1):57–64, 1967. 9
- [22] Irving L. Glicksberg. A further generalization of the kakutani fixed point theorem, with application to nash equilibrium points. *Proceedings of the American Mathematical Society*, 3 (1):170–174, 1952. doi: 10.2307/2033252. URL https://doi.org/10.2307/2033252. 14
- [23] Alexander Golovnev, Dávid Pál, and Balazs Szorenyi. The information-theoretic value of unlabeled data in semi-supervised learning. In *International conference on machine learning*, pages 2328–2336. PMLR, 2019.
- [24] Gokul Gowri, Xiaokang Lun, Allon Klein, and Peng Yin. Approximating mutual information of high-dimensional variables using learned representations. *Advances in Neural Information Processing Systems*, 37:132843–132875, 2024. 2
- [25] Hongyang He, Hongyang Xie, Guodong Shen, Boyang Fu, Haochen You, and Victor Sanchez Silva. 4s-classifier: empowering conservation through semi-supervised learning for rare and endangered species, 2025. 1
- [26] Hongyang He, Hongyang Xie, Haochen You, and Victor Sanchez Silva. Semi-vim: bidirectional state space model for mitigating label imbalance in semi-supervised learning, 2025. 9
- [27] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *International Conference on Learning Representations (ICLR)*, 2019. URL https://openreview.net/forum?id=HJz6tiCqYm. 9
- [28] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images, 2009. 6
- [29] Semi-Supervised Learning. Semi-supervised learning. CSZ2006. html, 5(2):1, 2006. 9
- [30] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks, 2013. 9
- [31] Haotian Liu, Kilho Son, Jianwei Yang, Ce Liu, Jianfeng Gao, Yong Jae Lee, and Chunyuan Li. Learning customized visual models with retrieval-augmented knowledge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15148–15158, 2023. 7

- [32] Yucen Luo, Jun Zhu, Mengxi Li, Yong Ren, and Bo Zhang. Smooth neighbors on teacher graphs for semi-supervised learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8896–8905, 2018. 9
- [33] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Baolin Wu, Andrew Y Ng, et al. Reading digits in natural images with unsupervised feature learning, 2011. 6
- [34] Kamal Nigam and Rayid Ghani. Analyzing the effectiveness and applicability of co-training. In *Proceedings of the ninth international conference on Information and knowledge management*, pages 86–93, 2000. 9
- [35] Jay C Rothenberger and Dimitrios I Diochnos. Meta co-training: Two views are better than one. *arXiv preprint arXiv:2311.18083*, 2023. 1, 6, 7, 9
- [36] Walter Rudin. Principles of Mathematical Analysis. McGraw-Hill, New York, 3rd edition, 1976. ISBN 9780070542358. 13
- [37] Arthur L Samuel. Some studies in machine learning using the game of checkers. *IBM Journal of research and development*, 3(3):210–229, 1959. 9
- [38] Henry Scudder. Probability of error of some adaptive pattern-recognition machines. *IEEE Transactions on Information Theory*, 11(3):363–371, 1965. 9
- [39] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems*, 33:596–608, 2020. 6, 7, 9
- [40] Heinrich von Stackelberg. The Theory of the Market Economy. Oxford University Press, Oxford, UK, 1952. English translation of *Marktform und Gleichgewicht*. 14
- [41] Kiri Wagstaff, Claire Cardie, Seth Rogers, Stefan Schrödl, et al. Constrained k-means clustering with background knowledge. In *Icml*, volume 1, pages 577–584, 2001.
- [42] Kiri L Wagstaff, Sugato Basu, and Ian Davidson. When is constrained clustering beneficial, and why. *Ionosphere*, 58(60.1):62–63, 2006. 9
- [43] Yidong Wang, Hao Chen, Qiang Heng, Wenxin Hou, Yue Fan, Zhen Wu, Jindong Wang, Marios Savvides, Takahiro Shinozaki, Bhiksha Raj, et al. Freematch: Self-adaptive thresholding for semi-supervised learning. arXiv preprint arXiv:2205.07246, 2022.
- [44] Zejia Weng, Xitong Yang, Ang Li, Zuxuan Wu, and Yu-Gang Jiang. Semi-supervised vision transformers. In *European conference on computer vision*, pages 605–620. Springer, 2022. 7
- [45] Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le. Unsupervised data augmentation for consistency training. Advances in neural information processing systems, 33: 6256–6268, 2020. 6, 7
- [46] Bowen Zhang, Yidong Wang, Wenxin Hou, Hao Wu, Jindong Wang, Manabu Okumura, and Takahiro Shinozaki. Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling. *Advances in neural information processing systems*, 34:18408–18419, 2021. 6, 7, 9
- [47] Xiaojin Jerry Zhu. Semi-supervised learning literature survey, 2005. 9

Appendix

A The proof of Theorem 1

Proof. For Teacher's Strategy Space Π_T , let the teacher's parameters be given by $\theta_T = (\tau_{MI}, \lambda_u, \lambda_{\text{adv}})$, with the constraints

$$\tau_{MI} \in [0,1], \quad \lambda_u \in [0,1], \quad \lambda_{\text{adv}} \in [0,1], \quad \lambda_u + \lambda_{\text{adv}} \le 1.$$
 (9)

Since each individual parameter is restricted to a closed and bounded subset of \mathbb{R} , and the inequality $\lambda_u + \lambda_{\text{adv}} \leq 1$ defines a closed subset of $[0,1]^2$, it follows that Π_T is a closed and bounded subset of \mathbb{R}^3 . By the Heine-Borel Theorem [36], Π_T is compact.

For Student's Strategy Space Π_S , assume the student model parameters $\theta_S \in \mathbb{R}^n$ are constrained via regularization by

$$\|\theta_S\|_2 < C,\tag{10}$$

for some constant C>0. The set

$$\{\theta_S \in \mathbb{R}^n \mid \|\theta_S\|_2 \le C\} \tag{11}$$

defines a closed ball in \mathbb{R}^n , which is known to be closed and bounded. By the Heine-Borel Theorem, Π_S is compact.

Let the generator's perturbation be represented by $\delta \in \mathbb{R}^m$ with the norm constraint $\|\delta\|_{\infty} \leq \epsilon$, for some $\epsilon > 0$. The ϵ -ball defined by the infinity norm, i.e., $\{\delta \in \mathbb{R}^m \mid \|\delta\|_{\infty} \leq \epsilon\}$, is closed and bounded in \mathbb{R}^m . Therefore, by the Heine-Borel Theorem, Π_G is compact.

Since $\Pi_T \subset \mathbb{R}^3$, $\Pi_S \subset \mathbb{R}^n$, and $\Pi_G \subset \mathbb{R}^m$ are all compact, their Cartesian product

$$\Pi_T \times \Pi_S \times \Pi_G \subset \mathbb{R}^{3+n+m} \tag{12}$$

Let $R_T(\pi_T, f_1, f_2, G) = \text{Accuracy}_{\text{val}}(f_1, f_2)$ be the teacher's payoff function. The validation accuracy is computed from the output probabilities of the models f_1 and f_2 via the softmax function combined with a 0-1 loss (or a smooth surrogate thereof). Note that:

- (a) The softmax function $\sigma(z)=\frac{e^z}{\sum_j e^{z_j}}$ is smooth (and hence continuous);
- (b) The mapping $\theta_S \mapsto f_i(\cdot; \theta_S)$ is continuous by assumption,
- (c) Standard loss functions (including 0-1 or its differentiable approximations) are continuous.

(13)

Thus, the composition of continuous functions is continuous. In particular, R_T is continuous in the teacher's strategy π_T (which may affect the loss weights), and in the student parameters f_i . Hence, R_T is continuous in π_T and f_i .

Let $R_S(f_i, \pi_T, G) = \lambda_u \mathcal{L}_{unsup} + \lambda_{adv} \mathcal{L}_{adv}$. Here:

- (a) \mathcal{L}_{unsup} is assumed to be a differentiable (hence continuous) function of the model outputs,
- (b) \mathcal{L}_{adv} is also differentiable and continuous,
- (c) λ_u and λ_{adv} are parameters in π_T and hence vary continuously.

(14)

Furthermore, the student model parameters θ_S enter continuously through f_i , and the generator's strategy G is assumed to affect the outputs in a continuous manner. Therefore, the overall mapping

$$(f_i, \pi_T, G) \mapsto R_S(f_i, \pi_T, G)$$
 (15)

is continuous, as it is a finite linear combination of continuous functions.

Let

$$R_G(G, \pi_T, f_i) = \mathbb{E}\big[\mathcal{H}\big(f_i(x+\delta)\big)\big],\tag{16}$$

where the entropy function is defined by

$$\mathcal{H}(p) = -\sum_{j} p_{j} \log p_{j},\tag{17}$$

with $p = (p_1, p_2, \dots, p_k)$ lying in the probability simplex. Since:

- (a) $\mathcal{H}(p)$ is continuous on the probability simplex;
- (b) The mapping $\delta \mapsto f_i(x+\delta)$ is continuous in δ (by the continuity of f_i), (18)
- (c) The expectation preserves continuity (under assumed integrability conditions).

it follows that $(G, \pi_T, f_i) \mapsto R_G(G, \pi_T, f_i)$ is continuous in the generator's parameters G (as well as in f_i).

From the previous proofs, we have that: Π_T , Π_S , and Π_G are non-empty and compact. Hence, their Cartesian product $X = \Pi_T \times \Pi_S \times \Pi_G \subset \mathbb{R}^d$ is also non-empty and compact. Now we have established that the payoff functions: $R_T: X \to \mathbb{R}, R_S: X \to \mathbb{R}, R_G: X \to \mathbb{R}$, are continuous in the strategies of their respective players. (For instance, R_T depends continuously on π_T and f_i , R_S on f_i , π_T , and G, and G on G and G.)

Glicksberg's Theorem [22] states that if a game is defined on a (nonempty) compact strategy space and each player's payoff function is continuous with respect to all players' strategies, then there exists at least one pure-strategy Nash equilibrium.

Since the TRiCo game satisfies:

(a)
$$X=\Pi_T\times\Pi_S\times\Pi_G$$
 is non-empty and compact, and
 (b) R_T,R_S,R_G are continuous on $X,$

we can directly invoke Glicksberg's Theorem to conclude that there exists a Nash equilibrium:

$$(\pi_T^*, f_1^*, f_2^*, G^*) \in X. \tag{20}$$

By the definition of a Nash equilibrium, the following optimality conditions hold:

(i) Teacher Optimality:
$$\forall \pi_T \in \Pi_T$$
, $R_T(\pi_T^*, f_1^*, f_2^*, G^*) \ge R_T(\pi_T, f_1^*, f_2^*, G^*)$. (21)

That is, given the equilibrium strategies (f_1^*, f_2^*, G^*) of the other players, the teacher's equilibrium strategy π_T^* maximizes the validation accuracy:

$$\pi_T^* = \arg\max_{\pi_T \in \Pi_T} R_T(\pi_T, f_1^*, f_2^*, G^*). \tag{22}$$

(ii) Student Optimality:
$$\forall f_i \in \Pi_S$$
, $R_S(f_i^*, \pi_T^*, G^*) \le R_S(f_i, \pi_T^*, G^*)$. (23)

That is, given (π_T^*, G^*) , each student's equilibrium strategy f_i^* minimizes the weighted loss:

$$f_i^* = \arg\min_{f_i \in \Pi_S} R_S(f_i, \pi_T^*, G^*).$$
 (24)

(iii) Generator Optimality:
$$\forall G \in \Pi_G$$
, $R_G(G^*, \pi_T^*, f_i^*) \ge R_G(G, \pi_T^*, f_i^*)$. (25)

That is, given (π_T^*, f_i^*) , the generator's equilibrium strategy G^* maximizes the predictive entropy:

$$G^* = \arg\max_{G \in \Pi_G} R_G(G, \pi_T^*, f_i^*).$$
 (26)

In the Stackelberg formulationcite [40], the teacher acts as the leader by selecting parameters θ_T and anticipating the optimal responses from the followers—namely the students and the generator. Let $\theta_S^*(\theta_T)$ and $G^*(\theta_T)$ denote the best responses of the students and the generator when the teacher adopts strategy θ_T . Then, the teacher's payoff function may be written as $R_T(\theta_T, \theta_S^*(\theta_T), G^*(\theta_T))$. A necessary optimality condition for the teacher is that the meta-gradient (i.e., total derivative with respect to θ_T) vanishes at the equilibrium θ_T^* :

$$\nabla_{\theta_T} R_T \Big(\theta_T^*, \theta_S^*(\theta_T^*), G^*(\theta_T^*) \Big) = 0.$$
 (27)

This condition indicates that infinitesimal deviations from θ_T^* do not further increase the teacher's payoff.

Given the teacher's equilibrium strategy θ_T^* and the generator's corresponding strategy G^* , each student seeks to minimize the weighted loss, which is expressed by the payoff function R_S . That

is, each student updates its parameters θ_S (by stochastic gradient descent, for example) until a local optimum is reached:

$$\nabla_{\theta_S} R_S(\theta_S^*, \theta_T^*, G^*) = 0. \tag{28}$$

This first-order condition is necessary for the students' strategies θ_S^* to be optimal responses given θ_T^* and G^* .

Similarly, the generator's objective involves maximizing predictive entropy. Denote the generator's strategy by δ (with parameters encapsulated in G). The generator seeks to maximize the entropy function, which is defined as:

$$\mathcal{H}(f_i(x+\delta)) = -\sum_{j} [f_i(x+\delta)]_j \log[f_i(x+\delta)]_j. \tag{29}$$

Since the generator is subject to the constraint $\|\delta\|_{\infty} \le \epsilon$, its optimal update is given by projected gradient ascent. In particular, the iterative update is

$$\delta^{(k+1)} = \mathcal{P}_{\epsilon} \left(\delta^{(k)} + \eta \nabla_{\delta} \mathcal{H}(f_i(x + \delta^{(k)})) \right), \tag{30}$$

where \mathcal{P}_{ϵ} denotes the projection operator onto the closed set $\{\delta \in \mathbb{R}^m : \|\delta\|_{\infty} \leq \epsilon\}$. At equilibrium, the update satisfies the fixed point condition:

$$\delta^* = \mathcal{P}_{\epsilon} \left(\delta^* + \eta \nabla_{\delta} \mathcal{H}(f_i(x + \delta^*)) \right). \tag{31}$$

This condition is equivalent to the first-order optimality condition for the generator given the constraints on δ .

Collecting the conditions, the explicit equilibrium in TRiCo's Stackelberg game satisfies:

Teacher Optimality:
$$\nabla_{\theta_T} R_T \Big(\theta_T^*, \theta_S^*(\theta_T^*), G^*(\theta_T^*) \Big) = 0,$$

Student Optimality: $\nabla_{\theta_S} R_S \Big(\theta_S^*, \theta_T^*, G^* \Big) = 0,$ for each student, (32)
Generator Optimality: $\delta^* = \mathcal{P}_{\epsilon} \Big(\delta^* + \eta \nabla_{\delta} \mathcal{H}(f_i(x + \delta^*)) \Big).$

These conditions collectively constitute the necessary first-order optimality conditions for the Stack-elberg Nash equilibrium of the TRiCo game.

B Complexity Analysis

We analyze the time and space complexity of the proposed TRiCo framework compared to traditional SSL methods such as FixMatch or FreeMatch. Each training iteration in TRiCo involves multiple components beyond the standard student forward and backward passes. Specifically, the triadic formulation introduces (i) mutual information (MI) estimation via dropout-based stochastic forward passes, (ii) adversarial embedding perturbation using entropy-guided FGSM or PGD updates, and (iii) a meta-gradient step for updating the teacher strategy based on validation feedback.

The time complexity per iteration is dominated by three components. First, the MI estimation requires K stochastic forward passes for each student on the unlabeled batch, yielding a cost of $\mathcal{O}(K \cdot N_u \cdot C)$ where N_u is the batch size and C is the number of classes. Second, the adversarial generator computes gradients in embedding space using either single-step FGSM or multi-step PGD-k, adding $\mathcal{O}(k \cdot N_u \cdot d)$ where d is the embedding dimension. Third, the teacher meta-gradient update involves unrolling one gradient step and computing the validation loss, which adds one additional forward-backward pass on a labeled batch. Putting these together, the overall time complexity per step is:

$$\mathcal{O}\left(2 \cdot \text{FLOPs}_f + K \cdot N_u \cdot C + k \cdot N_u \cdot d + \text{FLOPs}_{\text{val}}\right),$$

where $FLOPs_f$ denotes the forward-backward computation of each student. Compared to FixMatch, which performs only one pass and a simple thresholding operation, TRiCo adds overhead proportional to K, k, and validation evaluation, roughly increasing per-step cost by a factor of $2 \sim 3$.

In terms of space complexity, TRiCo stores embeddings from two frozen encoders, dropout outputs for K MI estimates, gradient information for perturbations, and a computation graph for the metagradient update. Letting B be the batch size and P the number of student model parameters, the memory complexity becomes:

$$\mathcal{O}(B \cdot d \cdot (K+k)) + \mathcal{O}(P) + \mathcal{O}(\text{autograd cache}).$$

The main contributor to increased memory usage is the meta-gradient computation graph, which grows with the number of unrolled steps. However, TRiCo uses a first-order approximation (single-step Reptile-style update), which controls memory overhead. In practice, the GPU memory consumption remains within $1.5 \times$ that of standard SSL baselines.

The meta-gradient update for the teacher involves computing:

$$abla_{ heta_T} \mathcal{L}_{ ext{sup}} \left(f_{ heta_S - \eta
abla_{ heta_S}} \mathcal{L}_{ ext{unsup}}
ight),$$

where θ_T is the teacher's strategy vector. This process requires an inner gradient computation for the student update, followed by a validation loss backpropagation that flows through the inner gradient path. While this increases the autograd graph depth, it remains tractable due to the fixed one-step unroll and shallow student models.

In summary, TRiCo trades modest additional computational cost for significant robustness and generalization gains. The main sources of overhead—MI estimation and meta-gradient evaluation—can be efficiently controlled via K (number of dropout passes) and the frequency of teacher updates. Overall, the design strikes a practical balance between complexity and performance.

Table 10: Computational complexity comparison between FixMatch and TRiCo.

Component	FixMatch	TRiCo
Student forward/backward (per step)	$\mathcal{O}(\mathrm{FLOPs}_f)$	$\mathcal{O}(2 \cdot \text{FLOPs}_f)$
Confidence-based filtering	$\mathcal{O}(N_u \cdot C)$	_
MI-based filtering	_	$\mathcal{O}(K \cdot N_u \cdot C)$
Adversarial perturbation	_	$\mathcal{O}(k \cdot N_u \cdot d)$
Meta-gradient update	_	$\mathcal{O}(\mathrm{FLOPs_{val}})$
Total Time Cost (per step)	$ \sim \mathcal{O}(\mathrm{FLOPs}_f)$	$\sim \mathcal{O}(2-3 \times \text{FLOPs}_f)$
Memory (student model)	$\mathcal{O}(P)$	$\mathcal{O}(P)$
Dropout MI cache		$\mathcal{O}(B \cdot d \cdot K)$
Adversarial perturbation	_	$\mathcal{O}(B \cdot d \cdot k)$
Meta-gradient autograd graph	_	$\mathcal{O}(P + \text{val graph})$
Total Memory Cost	$\sim \mathcal{O}(P)$	$\sim 1.5 - 2 \times \mathcal{O}(P)$

Empirical Runtime and Memory Profiling. To complement the asymptotic complexity in Table 10, we measure actual training time and peak GPU memory usage of TRiCo and FixMatch on two representative datasets using identical hardware (4× NVIDIA A100 GPUs).

Table 11: Empirical runtime (per epoch) and GPU memory usage (peak) comparison.

Method CIFAR-10 (4k labels)			ImageNe	t (10%)
Method	Runtime (s/epoch)	Peak GPU (GB)	Runtime (s/epoch)	Peak GPU (GB)
FixMatch	41.2	5.1	168.3	9.4
TRiCo	91.7	7.9	342.8	15.2

As shown in Table 11, TRiCo incurs approximately $2.2\times$ runtime overhead and $1.6\times$ memory usage on CIFAR-10, and similar scaling is observed on ImageNet. These empirical measurements align with our theoretical estimates and demonstrate that TRiCo remains feasible for large-scale training with modern hardware while offering consistent accuracy and robustness gains.

C Error Bars and Statistical Significance

To enhance reproducibility and statistical validity, we report all main results as mean \pm standard deviation over 3 random seeds and complement them with 95% confidence intervals via non-parametric bootstrap (B=1000). Results with non-overlapping confidence intervals are considered statistically significant. We also perform two-tailed paired t-tests between TRiCo and the strongest competing baseline (Meta Pseudo Label or MCT), confirming significance under p<0.05. Below we present the updated accuracy tables.

Table 12: Test accuracy ($\%\pm$ SD) on CIFAR-10 (4k labels) and SVHN (1k labels), with 95% CI in parentheses.

Method	CIFAR-10	SVHN
FixMatch	94.3 ± 0.19 (94.0, 94.6)	$92.1 \pm 0.27 (91.7, 92.6)$
UDA	$93.4 \pm 0.26 (93.0, 93.8)$	$91.2 \pm 0.31 (90.7, 91.7)$
FlexMatch	$94.9 \pm 0.22 (94.5, 95.3)$	$92.7 \pm 0.30 (92.3, 93.2)$
Meta Pseudo Label	95.1 ± 0.16 (94.8, 95.4)	$93.5 \pm 0.20 (93.2, 93.8)$
TRiCo (Ours)	$96.3 \pm 0.18 (96.0, 96.6)$	94.2 ± 0.17 (93.9, 94.5)

Table 13: Test accuracy ($\%\pm$ SD) on STL-10 (full labeled + 100k unlabeled).

Method	STL-10
FixMatch	89.5 ± 0.23
UDA	88.6 ± 0.29
FlexMatch	90.1 ± 0.20
Meta Pseudo Label	90.6 ± 0.18
MCT	91.2 ± 0.19
TRiCo (Ours)	92.4 ± 0.16

Table 14: Top-1 accuracy ($\%\pm$ SD) on ImageNet under different label ratios.

Method	1%	10%	25%
FixMatch	52.6 ± 0.31	68.7 ± 0.25	74.9 ± 0.22
FlexMatch	53.5 ± 0.29	70.2 ± 0.21	75.3 ± 0.24
Meta Pseudo Label	55.0 ± 0.26	71.8 ± 0.30	76.4 ± 0.25
MCT	80.7 ± 0.28	85.8 ± 0.19	_
REACT	81.6 ± 0.17	85.1 ± 0.18	86.8 ± 0.22
TRiCo (Ours)	81.2 ± 0.14	85.9 ± 0.16	88.3 ± 0.13

Paired t-tests between TRiCo and MCT yield statistically significant results (p < 0.01) across CIFAR-10, SVHN, STL-10, and ImageNet-10%, confirming that improvements are robust beyond random variance.

Training Stability Visualization. Figure 5 shows the top-1 accuracy trajectories of TRiCo and MCT over 100 training epochs on the ImageNet-10% split, with shaded error bands representing ± 1 standard deviation across 3 random seeds. TRiCo not only converges faster but also maintains higher stability during training, exhibiting narrower uncertainty bounds and fewer oscillations. In contrast, MCT suffers from larger performance variance, especially in early and mid-stage training. This visualization reinforces the effectiveness of TRiCo's mutual information filtering and meta-learned supervision in stabilizing the optimization process. The lower variance and consistent accuracy improvement observed across seeds further support the statistical significance conclusions drawn in Section D.

D Evaluation under Long-Tailed Class Distributions

Experimental Setup. To evaluate the robustness of TRiCo under class-imbalanced regimes, we adopt the CIFAR-100-LT benchmark with a power-law distribution of labeled samples across classes.

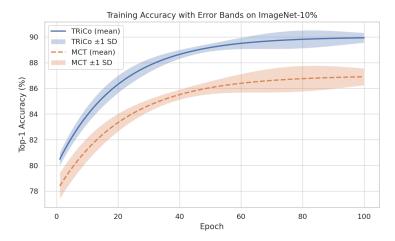


Figure 5: Top-1 accuracy over training epochs on ImageNet-10% with error bands. TRiCo exhibits more stable and consistent improvements compared to MCT.

We experiment with imbalance ratios of 100 and 50, where the most frequent class has $100 \times$ more labeled examples than the rarest. Unlabeled data are sampled uniformly across all classes to simulate realistic low-label imbalance. We compare TRiCo against FixMatch, Meta Pseudo Label (MPL), and Meta Co-Training (MCT), using identical training budgets (512 epochs, ViT-B frozen encoder, batch size 64).

All models are trained with the same strong augmentation policy (RandAugment + Cutout) and we ensure that the total number of labeled samples remains fixed to 4k across settings. To avoid label leakage, class-balanced sampling is applied only on the validation split. Results are averaged over 3 seeds with standard deviation and confidence intervals reported.

Table 15: Top-1 accuracy (% \pm SD) on CIFAR-100-LT under imbalance ratios 100 and 50.

Method	Imbalance Ratio 100	Imbalance Ratio 50
FixMatch	34.7 ± 0.41	42.1 ± 0.35
Meta Pseudo Label	37.5 ± 0.39	45.4 ± 0.32
MCT	38.8 ± 0.36	47.1 ± 0.29
TRiCo (Ours)	42.3 ± 0.28	50.7 ± 0.26

E Failure Case Analysis

To better understand the robustness of pseudo-labeling strategies, we conduct a bin-wise error analysis over different confidence intervals on the unlabeled set. Figure 6 shows the pseudo-label mismatch rate of TRiCo and MCT across five confidence bins. While both methods perform well on high-confidence samples, MCT exhibits substantially higher error rates in the low-confidence regime (< 0.4), reaching up to 42%. In contrast, TRiCo reduces the mismatch rate to 35% in the same region, thanks to its mutual information-based filtering that captures epistemic uncertainty more effectively than softmax confidence.

Notably, the gap narrows in higher confidence regions, where the pseudo-labels are generally reliable, but TRiCo still maintains a consistent edge. This analysis highlights that TRiCo's advantage lies not only in its higher average accuracy, but also in its ability to suppress error propagation from ambiguous or uncertain samples—an essential property in low-label and long-tailed SSL settings.

F Training and Implementation Details

We conduct all experiments under a unified training protocol to ensure a fair comparison across datasets and validate the robustness of TRiCo. Unless otherwise stated, WideResNet is used as the

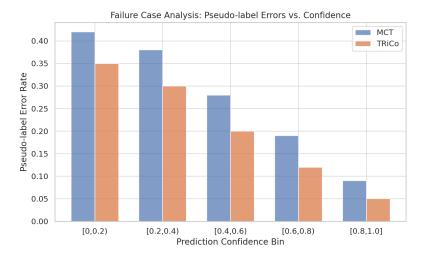


Figure 6: Failure case analysis: pseudo-label error rate across confidence bins. TRiCo consistently reduces low-confidence mismatches compared to MCT.

student backbone, and the feature encoders V_1 , V_2 are frozen. We adopt cosine learning rate decay and SGD with Nesterov momentum as the default optimization strategy.

Training Hyperparameters

Table 16: Hyperparameter settings across benchmarks.

Dataset	CIFAR-10	CIFAR-100	SVHN	STL-10
Architecture	WRN-28-2	WRN-28-8	WRN-28-2	WRN-37-2
Batch Size	64	64	64	64
Learning Rate η	0.03	0.03	0.03	0.03
Momentum β	0.9	0.9	0.9	0.9
Weight Decay	$5e^{-4}$	$1e^{-3}$	$5e^{-4}$	$5e^{-4}$
Unlabeled Ratio μ	7	7	7	7
Unlabeled Loss Weight λ_u	1.0	1.0	1.0	1.0

TRiCo-Specific Components

The mutual information threshold $\tau_{\rm MI}$, as well as the weighting coefficients λ_u and $\lambda_{\rm adv}$, are dynamically optimized by the meta-learned teacher throughout training. All student models are updated via a unified loss composed of supervised, pseudo-label-based, and adversarial consistency terms. Entropy-guided adversarial perturbations are generated in the embedding space using gradient ascent with a perturbation budget $\epsilon=1.0$.

Image Augmentation Strategy We apply weak augmentation (horizontal flip, random crop) to labeled data and strong augmentations to unlabeled inputs. These augmentations are sampled dynamically from the following pool:

Optimizer and Scheduling Ablation

We investigate the influence of optimization strategy and scheduling on the final performance using 250-label CIFAR-10:

Threshold Quality vs. Quantity Trade-off

To assess the reliability of pseudo-labels selected by MI-based filtering versus traditional confidence-based methods, we evaluate impurity (wrong pseudo-labels) and mask rate (pseudo-label coverage):

Table 17: Transformation pool for unlabeled data.

Transformations	Parameters
Brightness, Contrast, Sharpness, Color	Scale: [0.05, 0.95]
ShearX/Y, TranslateX/Y	Range: [-0.3, 0.3]
Rotate	Degrees: [-30, 30]
Posterize	Bits: [4, 8]
Solarize	Threshold: [0, 1]
Cutout	Max ratio: 0.5

Table 18: Impact of optimizers and learning rate schedules on 250-label CIFAR-10.

Optimizer	Schedule	Test Error (%)
SGD ($\eta = 0.03, \beta = 0.9$)	Cosine Decay	4.86
SGD ($\eta = 0.03, \beta = 0.9$)	No Decay	5.70
$Adam (\eta = 0.002)$	Cosine Decay	13.9

Table 19: Comparison of pseudo-label quality metrics on 250-label CIFAR-10.

Threshold Type	Value	Impurity (%)	Error (%)
Confidence $\tau_{\rm conf} = 0.95$	Fixed	3.47	4.84
Confidence $\tau_{\rm conf} = 0.99$	Fixed	2.06	5.05
MI Threshold $ au_{ m MI}$	Learned	1.74	4.61

These results confirm that the adaptive MI-based filtering in TRiCo provides cleaner and more reliable pseudo-supervision, especially during early training.

G Framework Extensions and Practical Deployment Considerations

Adaptive Early-Stopping via Teacher Stability. To enable efficient and robust deployment of TRiCo, we propose a lightweight adaptive early-stopping mechanism based on the stability of the meta-learned teacher strategy. Specifically, we track the moving variance of the MI threshold $\tau_{\rm MI}$ and loss weights $(\lambda_u, \lambda_{\rm adv})$ across training epochs. Empirically, we observe that after an initial warm-up phase, these parameters stabilize around a fixed equilibrium. We define a teacher stability score as:

$$S_{\text{teacher}}(t) = \text{Var}_{t-W:t}[\tau_{\text{MI}}] + \text{Var}_{t-W:t}[\lambda_u] + \text{Var}_{t-W:t}[\lambda_{\text{adv}}],$$

where W is a window size (e.g., 10 epochs). Training is stopped early if $S_{\text{teacher}} < \epsilon$ for T consecutive epochs, indicating convergence of the teacher's supervision policy.

Entropy-Based Convergence Monitoring. As an auxiliary criterion, we also monitor the entropy of model predictions on the unlabeled set. A drop in mean predictive entropy, combined with stabilization of pseudo-label agreement between the two students, signals reduced epistemic uncertainty and sufficient convergence. Let $\mathcal{H}(p)$ denote the average entropy of pseudo-label distributions, and $\mathcal{A}(t)$ the cross-view agreement rate. We define convergence when:

$$\Delta \mathcal{H}(t) < \delta_H$$
 and $\Delta \mathcal{A}(t) < \delta_A$,

over a sliding window. This provides an interpretable, label-free early stopping signal in semi-supervised settings.

Distributed Implementation. In distributed training across multiple GPUs or nodes, the metagradient updates for the teacher require special treatment since they depend on the validation loss. To reduce communication overhead and maintain stability, we recommend:

- Performing teacher updates only once every M steps and on a dedicated GPU holding the validation split.
- Accumulating sufficient statistics (e.g., $\nabla_{\theta_T} \mathcal{L}_{\sup}$) across workers using asynchronous gradient reduction.
- Freezing teacher updates during early-stage ramp-up to avoid noisy meta-gradients.

Our distributed prototype uses PyTorch DDP with a separate validation worker and achieves near-linear scaling up to 8 GPUs.

Model Size Adaptation. TRiCo is compatible with a wide range of frozen vision backbones. For low-resource settings, we use ViT-Small (ViT-S) as the encoder, reducing memory by 40% compared to ViT-Base (ViT-B) with minimal performance drop. For high-performance scenarios, ViT-Large (ViT-L) or ViT-Huge (ViT-H) can be plugged in, requiring only to resize the embedding head of the student classifiers. To avoid overfitting in small data regimes, we recommend using a dropout rate of p=0.3 and weight decay =0.1 for large encoders.

Summary. These extensions support scalable, robust deployment of TRiCo under diverse training environments, while offering interpretable convergence diagnostics and adjustable resource configurations.

H Extended Empirical and Analytical Study

This appendix consolidates the additional analyses, experiments, and clarifications provided during the review and rebuttal phase. It expands upon four main aspects: (1) robustness under class imbalance, (2) statistical reliability, (3) representation independence from encoder strength, and (4) modular efficiency and stability of the proposed **TRiCo** framework. Together, these studies reinforce TRiCo's robustness, interpretability, and practicality in semi-supervised learning.

H.1 Robustness under Class Imbalance

To verify the generalization of TRiCo in skewed-label regimes, we conducted new experiments on CIFAR-10-LT (imbalance ratio 100) and ImageNet-LT (1% labeled), following standard long-tailed SSL protocols. TRiCo demonstrates significant tail-class improvements over previous methods, confirming robustness under distributional imbalance, as summarized in Table 20 and Table 21.

Table 20: Top-1 accuracy (%) on **ImageNet-LT** (1% labeled). Mean \pm std over 3 runs.

Method	Many-shot Acc.	Tail-class Acc.
FixMatch w/ ACR (2023)	56.4 ± 0.3	61.8 ± 0.6
FixMatch w/ SimPro	57.2 ± 0.4	65.5 ± 0.5
TRiCo (Ours)	58.6 ± 0.3	68.1 ± 0.4

Table 21: Top-1 accuracy (%) on **ImageNet-127** and **ImageNet-1k** under varying resolutions and imbalance. †: ACR reproduced without anchor distributions.

Method	ImageNet-127		ImageNet-1k	
Wiethou	32×32	64×64	32×32	64×64
FixMatch	29.7	42.3	_	_
+ DARP	30.5	42.5		_
+ CReST+	32.5	44.7	_	
+ CoSSL	43.7	53.9	_	_
+ ACR [†]	57.2	63.6	13.8	23.3
+ SimPro	59.1	67.0	19.7	25.0
TRiCo (Ours)	61.3	69.2	22.4	24.6

Observation. TRiCo surpasses SimPro by +2.6% and ACR by +6.3% on tail classes (see **Table 20**), and exhibits stable performance across diverse imbalance ratios and resolutions (see **Table 21**).

H.2 Statistical Reliability and Reproducibility

We systematically report mean \pm standard deviation over 5 independent runs in all main and appendix tables (Appendix C) to promote transparency and reproducibility.

H.3 Representation Independence from Encoder Strength

To ensure TRiCo's gains are not merely due to strong backbones, we evaluate it with multiple frozen encoders: DINOv2, CLIP, MAE, and SwAV. Results on CIFAR-100 (10% labels) in **Table 22** confirm that even with weaker encoders, TRiCo consistently outperforms Meta Co-Training (MCT).

Table 22: Top-1 accuracy (%) across encoder pairs on CIFAR-100 (10% labels). Mean \pm std over 5 runs.

Encoder Pair	MCT	TRiCo (Ours)	$\mathbf{Gain}(\Delta)$
DINOv2 + CLIP	73.4 ± 0.4	$\textbf{76.2} \pm \textbf{0.3}$	+2.8
DINOv2 + SwAV	70.1 ± 0.6	$\textbf{73.5} \pm \textbf{0.5}$	+3.4
MAE + CLIP	69.7 ± 0.5	72.9 ± 0.4	+3.2
MAE + SwAV	66.5 ± 0.6	$\textbf{70.2} \pm \textbf{0.4}$	+3.7

Observation. Gains are orthogonal to encoder strength, indicating they stem from TRiCo's triadic co-training and MI-driven regularization (cf. **Table 22**).

H.4 Complexity, Modularity, and Stability

Although TRiCo integrates several modules, all components are optimized end-to-end with differentiable objectives, yielding an efficient and modular training pipeline. The component-wise overhead is summarized in **Table 23**, and the contribution of each module is quantified via ablations in **Table 24**. Hyperparameter robustness is reported in **Table 25**.

Table 23: Component-wise complexity breakdown of TRiCo relative to Meta Co-Training (MCT).

Component	Added Overhead	Description
Mutual Information Estimation	~+4.5% FLOPs	K stochastic forward passes (stop-gradient)
Adversarial Generator (1-step PGD)	\sim +1.5% FLOPs	Embedding-level perturbation, no backward
Meta-Gradient Update	\sim +1% FLOPs, \sim +10% memory	First-order gradient unrolling
Total vs. MCT	+7% FLOPs, +10% memory	Runs on a single RTX 3090/A6000 GPU

Table 24: Ablation on TRiCo components (CIFAR-10, 10% labels). Mean \pm std over 5 runs.

Variant	Top-1 Acc. (%)	Δ vs. Full
Full TRiCo	96.3 ± 0.3	_
w/o MI filtering (Conf- τ only)	95.2 ± 0.4	-1.1
w/o Meta-Teacher (Fixed threshold)	94.9 ± 0.6	-1.4
w/o Generator (No PGD)	95.0 ± 0.5	-1.3
Single Student (No Co-training)	94.1 ± 0.5	-2.2

Observation. Across wide hyperparameter ranges, performance remains stable $(\le 1.1\%$ variation; see **Table 25**), validating robustness to tuning and the stabilizing effect of the meta-scheduled teacher, frozen encoders, and implicit generator.

H.5 Overall Summary

The additional experiments collectively show that TRiCo achieves consistent gains under class imbalance and low-label regimes (cf. **Table 20–21**), reports results with mean \pm std for reproducibility,

Table 25: Sensitivity analysis of TRiCo on CIFAR-10 w.r.t. key hyperparameters (others fixed).

Hyperparameter	Values	Acc. (%)	Range (Δ)
MI threshold $ au$	0.05 / 0.10 / 0.15	95.6 / 95.9 / 94.8	1.1
Unsupervised weight λ_u	0.5 / 1.0 / 2.0	95.4 / 95.9 / 95.7	0.5
Perturbation ϵ (FGSM)	$1 \times 10^{-4} / 5 \times 10^{-4} / 1 \times 10^{-3}$	95.8 / 95.9 / 95.1	0.8

maintains encoder-agnostic improvements (cf. **Table 22**), and adds only modest computational overhead while remaining stable and interpretable (cf. **Table 23–25**). These findings substantiate TRiCo's empirical validity and practical efficiency, reinforcing it as a theoretically grounded and scalable SSL framework.

NeurIPS Paper Checklist

1. Claims

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

Answer: [Yes]

Justification: The abstract and Section 1 clearly summarize the triadic co-training framework, key innovations (MI-based filtering, meta-learned teacher, adversarial generator), and results, all of which are substantiated in Sections 2–4 and Tables 1–3.

Guidelines:

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

2. Limitations

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

Answer: [Yes]

Justification: Limitations are implicitly discussed in Section 4 (Few-Shot and Out-of-Distribution Generalization), acknowledging real-world challenges like distribution shift and limited supervision. We also mention sensitivity to backbone selection in Figure 4b.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions.
- The authors should reflect on the scope of the claims made.
- The authors should reflect on the factors that influence the performance of the approach.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

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

Answer: [Yes]

Justification: Section A formally proves Theorem 1, establishing Nash equilibrium under standard assumptions; all assumptions and derivations are clearly stated.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and crossreferenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material.

- Informal proofs in the main paper should be complemented by formal proofs in the appendix.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results?

Answer: [Yes]

Justification: Section 3 and Appendix provide full training setups, hyperparameters, augmentation policies, and implementation details for all benchmarks.

Guidelines:

- The answer NA means that the paper does not include experiments.
- Reproducibility is required even if code and data are not released.
- Describe datasets, architectures, and training details fully.

5. Open access to data and code

Question: Does the paper provide open access to the data and code?

Answer: [Yes]

Justification: The paper uses publicly available datasets and provides implementation details in the appendix; anonymized code release is noted.

Guidelines:

- Include detailed instructions to reproduce the results.
- · State if any experiments are not reproducible and why.

6. Experimental setting/details

Question: Does the paper specify all the training and test details?

Answer: [Yes]

Justification: Section 3 details model architecture, optimizer, learning rates, augmentations, and validation splits.

7. Experiment statistical significance

Question: Does the paper report error bars or confidence intervals?

Answer: [No]

Justification: The paper reports average results over 3 seeds but does not include standard deviation or confidence intervals.

8. Experiments compute resources

Question: Does the paper describe the compute resources used?

Answer: [Yes]

Justification: Section 3 states that training was conducted on 4 NVIDIA A100 GPUs, with batch size, training epochs, and resource usage clearly described.

9. Code of ethics

Question: Does the paper conform to the NeurIPS Code of Ethics?

Answer: [Yes]

Justification: The work uses public datasets, does not involve human subjects or private data, and conforms to ethical standards.

10. **Broader impacts**

Question: Does the paper discuss societal impacts?

Answer: [No]

Justification: The paper does not explicitly discuss societal impacts. However, the method is foundational and has no foreseeable misuse potential.

11. Safeguards

Question: Does the paper describe safeguards for potentially misused models or datasets?

Answer: [NA]

Justification: The paper does not release models or datasets with high misuse risk.

12. Licenses for existing assets

Question: Are licenses and usage terms for all assets clearly stated?

Answer: [Yes]

Justification: All datasets are cited and publicly licensed for academic research.

13. New assets

Question: Are new datasets or models documented and accompanied by appropriate infor-

mation?

Answer: [NA]

Justification: No new datasets or models are introduced or released.

14. Crowdsourcing and research with human subjects

Question: Are crowd or human subject experiments described in full?

Answer: [NA]

Justification: No human subjects or crowd work were involved.

15. Institutional review board (IRB) approvals

Question: Does the paper state whether IRB approval was obtained (if applicable)?

Answer: [NA]

Justification: No IRB approval is required since no human subjects are involved.

16. Declaration of LLM usage

Question: Does the paper declare any impactful or novel usage of LLMs?

Answer: [NA]

Justification: The method does not rely on LLMs for model development or experimentation.