I.J. Intelligent Systems and Applications, **2026, 3, 81-91**
Published Online on June 8, 2026 by MECS Press (http://www.mecs-press.org/)
DOI: 10.5815/ijisa.2026.03.05

A Hybrid Active and Semi-Supervised Learning Framework for Classification with Minimal Labeled Data

Kostiantyn O. Minkov*
Department of Mathematical Problems of Control and Cybernetics, Chernivtsi National University, Chernivtsi, Ukraine
E-mail: minkov.kostiantyn@chnu.edu.ua
ORCID iD: https://orcid.org/0009-0007-3400-050X
*Corresponding author

Igor V. Malyk
Department of Mathematical Problems of Control and Cybernetics, Chernivtsi National University, Chernivtsi, Ukraine
E-mail: i.malyk@chnu.edu.ua
ORCID iD: https://orcid.org/0000-0002-1291-9167

Received: 26 January 2026; Revised: 17 March 2026; Accepted: 25 April 2026; Published: 08 June 2026

Abstract: Modern machine learning models typically require large amounts of precisely labeled data to perform effectively. However, obtaining such labels is time-consuming and costly, especially in specialized domains such as medical image analysis and document classification, where unlabeled data is abundant but expert annotation is scarce. This paper addresses the problem of learning from very few labeled examples by jointly leveraging weak supervision, active learning (AL), and semi-supervised learning (SSL). A hybrid framework is proposed in which a small set of informative samples is actively selected for manual annotation using an entropy-based acquisition function combined with weak label disagreement scoring, while a large pool of unlabeled or weakly labeled data is exploited through SSL based on the FixMatch algorithm. The approach iteratively corrects noisy labels and refines the model with minimal human involvement. The framework is evaluated using a ResNet-18 classifier on the CIFAR-10 benchmark dataset and is compared against two baselines: pure active learning and pure semi-supervised learning. Each method is run independently across three random seeds at the key active learning rounds, and accuracy is reported as mean ± standard deviation. Across three independent seeds, the hybrid framework consistently leads both baselines at intermediate labelling budgets, with the largest absolute gap at Round 15 (+1.27 percentage points over pure active learning, +1.35 percentage points over pure SSL). The framework also offers a clear *label-efficiency* advantage: at Round 15, with $|D_L| = 6500$ labels, the hybrid method already reaches 0.6792 ± 0.0097 test accuracy – exceeding the accuracies that pure active learning (0.6730 ± 0.0139) and pure SSL (0.6687 ± 0.0056) attain only at Round 20 with $|D_L| = 7000$. By Round 20 all three methods saturate near a common data ceiling, indicating that the integrated use of weak supervision, active learning, and consistency-based SSL is most valuable when the annotation budget is genuinely constrained.

Index Terms: Active Learning, Semi-Supervised Learning, Neural Networks, Classification, Machine Learning, Data Analysis, Model, Accuracy.

1. Introduction

Each paradigm, however, has a blind spot. Weak supervision provides cheap labels but suffers from noise that is hard to correct without additional supervision. Classical AL queries an oracle for informative samples yet trains only on the growing labeled subset, discarding the bulk of the unlabeled pool. SSL, conversely, learns from all available data but assumes a fixed labeled set and provides no principled way to decide *which* examples should be labeled first. The paradigms are therefore complementary: AL determines *which* data points to label, and SSL determines *how* to learn from the rest.

This work is open access and licensed under the Creative Commons CC BY 4.0 License.

Recent work has begun to integrate these paradigms [9, 10, 17]. For instance, Song et al. [10] show that adding active sample selection to MixMatch yields clear accuracy gains under a fixed labeling budget on image classification benchmarks. However, most existing combinations either focus on AL+SSL without weak supervision or treat weak labels as a static initialization rather than as a signal to be iteratively corrected. A unified framework that simultaneously

(i) refines noisy weak labels through targeted oracle queries and (ii) exploits the remaining unlabeled and weakly labeled pool via consistency-based SSL remains underexplored, particularly under tight annotation budgets.

This paper addresses that gap. The primary contribution is a hybrid architecture that couples an entropy-and-disagreement-based active acquisition strategy with FixMatch-style semi-supervised refinement, enabling iterative correction of weak labels while extracting full value from the unlabeled pool with minimal human involvement. The empirical study, conducted across three random seeds at the key active learning rounds, shows that this combination yields a consistent label-efficiency advantage at intermediate annotation budgets.

2. Background and Related Work

2.1. Active Learning

Active learning is a method in which the learning algorithm can interactively query an oracle (e.g., a human annotator) to obtain labels for carefully selected unlabeled instances. The core idea is to choose data points that the current model considers most informative or uncertain, thereby maximizing the impact of each labeled instance [11]. Common query strategies include uncertainty sampling (selecting examples the model is least confident about), query-by-committee (where multiple models disagree most), and expected model change or error reduction. By labeling only the most infor- mative samples, active learning aims to achieve a target accuracy with significantly fewer labeled examples than random sampling.

In contemporary research, the ratio of labeled, noisy, and unlabeled data varies substantially depending on task specifics. In classical active learning setups, researchers often begin with a very small initial pool (e.g., $1-5\%$ of the total data) and use uncertainty-based or query-by-committee sampling to iteratively grow the labeled fraction to $10-20\%$ [11]. In weak supervision scenarios, the majority of the training set ($80-90\%$) may consist of noisy or heuristically generated labels, handled via sample reweighting or meta-learning algorithms [6]. For effective operation under extreme imbalance, hybrid algorithms are employed. For instance, Variational Adversarial Active Learning (VAAL) [9] enables stable training when labeled data constitutes less than 10%. Meanwhile, integrating AL strategies with semi-supervised algorithms such as MixMatch [10] demonstrates that even with a starting labeling budget below 1% (e.g., 250 labels for CIFAR-10), the model can successfully utilize the remaining 99% of unlabeled data through consistency regularization and pseudo-labeling.

However, standard active learning methods typically train the model only on the growing set of labeled data, discarding information from unlabeled examples. This means that once a subset of points is labeled and added to the training set, the remaining unlabeled examples do not directly contribute to model training. Integrating these methods with semi-supervised approaches can address this limitation by allowing the model to learn from *all* data points.

2.2. Semi-Supervised Learning

Semi-supervised learning techniques leverage unlabeled data alongside a small set of labeled data to improve learning. Methods such as pseudo-labeling (where model predictions on unlabeled data are treated as labels when confidence is sufficiently high), consistency regularization (encouraging stable model predictions under input perturbations), and graph-based label propagation are widely used [7, 8]. These methods effectively increase the training data volume without additional manual labels by assuming that the model can extract useful signals from unlabeled examples. In recent years, advanced SSL methods (e.g., FixMatch, mean teacher models) have shown that, under the right settings, a model can achieve performance close to fully supervised learning even when only a small fraction of data is labeled. However, SSL alone does not specify which data should be labeled first; this is where active learning can help.

2.3. Weak Supervision

Weak supervision provides a framework for generating approximate labels when true labels are limited. This includes methods such as distant supervision, heuristic rules, or crowd-sourced labels from non-experts [4, 5]. The resulting labels are noisy and potentially biased, but can be used to train an initial model. One challenge of weak supervision is managing noise and errors in weak labels. Approaches such as loss function correction, sample reweighting, or label smoothing are used to reduce the impact of label noise; recent comprehensive surveys catalogue dozens of such techniques and the benchmark datasets used to evaluate them [23]. Another approach is meta-learning applied to label correction: for example, Zheng et al. [12] propose a meta-learning framework (MLC) that learns to correct noisy labels by treating the correction procedure as a meta-level task, yielding significant improvements over naive training on weak labels.

2.4. Combining Active Learning with Weak and Semi-Supervised Learning

Only recently have researchers begun to unify these paradigms. The combination can take several forms.

A. Active Learning with Semi-Supervised Learning

Several hybrid methods combine SSL with active label acquisition, but they differ substantially in how the two components interact. Sinha et al. [9] propose Variational Adversarial Active Learning (VAAL), in which a variational autoencoder learns a latent representation of all data and an adversarial discriminator selects unlabeled samples whose latent code differs most from the labeled set; the unlabeled pool contributes to training implicitly through the VAE encoder rather than through an explicit SSL loss. Song et al. [10] take a more direct route, layering an active selection step on top of the MixMatch SSL algorithm and showing clear gains on CIFAR-10 with very few labels. Other methods alternate SSL training phases with active query phases: Ma et al. [13] integrate active learning into a teacher–student semi-supervised segmentation framework by automatically refining pseudo-labels for easy cases and querying labels only for the most challenging regions, illustrating that even unlabeled data not selected for manual annotation can contribute through pseudo-label refinement. Jimenez et al. [17] extend this combination to histopathology image segmentation with missing annotations, treating unannotated regions as a partially labeled pool exploited via consistency regularization. More recent work has continued to refine this paradigm: Roda and Geva [22] couple a convolutional autoencoder with contrastive learning so that the latent space simultaneously serves SSL representation learning and AL query selection, while Huang et al. [21] propose a hierarchical AL/SSL framework (HSSAL) that combines weak-to-strong self-training with progressive clustering-based acquisition and demonstrates strong label efficiency on remote-sensing benchmarks.

B. Active Learning for Weak Label Refinement

When initial labels are provided via weak supervision, active learning can be used to iteratively refine these labels. In this paradigm, an initial labeling model provides probabilistic labels and the system actively queries the oracle about data points that will most improve the labeling model. By focusing on the most uncertain points, a small number of queries can substantially improve the overall label quality.

C. Critical Comparison of Hybrid Approaches

To clarify how the proposed framework relates to existing work, hybrid AL/SSL methods are compared along three principal dimensions: (i) the criterion used to select samples for oracle annotation, (ii) the mechanism by which unlabeled data contributes to training, and (iii) whether weak (noisy) labels are explicitly handled or the non-oracle pool is treated as truly unlabeled. Table 1 summarizes representative methods along these dimensions.

Table 1. Critical comparison of hybrid AL/SSL methods

Method	Selection strategy	Unlabeled / weak-label strategy	Key limitation
VAAL [9]	Adversarial dis-Crimination in VAE latent space (representative- ness)	Implicit via VAE encoder; weak labels not used	No task-specific uncertainty signal; no weak-label handling
Song et al. [10]	Uncertainty/margin on Mix- Match features	MixMatch (consistency + label sharpening); weak labels not used	Treats non-oracle data as label-free; no noisy-label refinement
Ma et al. [13]	Region-level uncertainty (seg- mentation)	Teacher–student pseudo-labels; weak labels not used	Domain-specific (segmentation); no global label correction
Jimenez et al. [17]	Uncertainty over missing annotations	Consistency regularization; missing annotations treated as unlabeled	Missing-annotation regime, not noisy-label regime
MLC [12]	None (no active component)	Weak labels cor-Rected via meta- learning	No active query; no SSL exploitation of unlabeled data
Proposed	**Entropy + weak-label disagreement**	**FixMatch consistency + iterative weak-label refinement**	–

This combination of gaps – AL+SSL methods that ignore weak labels, and weak-label correction methods that ignore active selection – motivates the unified framework proposed in this paper, which integrates all three components: entropy-and-disagreement-based active selection, FixMatch-style consistency on the unlabeled pool, and iterative refinement of the weak-label set as new oracle labels arrive.

3. Methodology

The proposed framework is a hybrid active and semi-supervised learning architecture designed to address classification tasks using a minimal amount of human-labeled data. The workflow consists of the following key components. The overall system diagram, illustrating the active query and semi-supervised learning cycles, is presented in Fig. 1.

3.1. Neural Network Architecture

The classifier is based on a deep neural network (NN) – a parametric model composed of stacked layers of differentiable transformations whose weights are learned end-to-end via gradient-based optimization. The specific architecture used in this work is ResNet-18 (18-layer Residual Network) [14], chosen for its ability to learn effectively from medium- sized images (as in CIFAR-10) through the use of residual blocks, which prevent the vanishing gradient problem in deep networks. The network consists of an initial convolutional layer, a series of four groups of residual blocks for hierarchical feature extraction, a global average pooling layer, and a final fully connected layer. The network output

is a logit vector of dimension K, which is converted via a Softmax layer into a probability distribution $p(k|x;\ \theta)$ for each class. The parameters θ (network weights) are initialized using the He method and optimized during training to minimize the combined loss function.

3.2. Initial Weak Labeling

An initial training dataset $D = D_L \cup D_U$ is assumed, where D_L is a small set of instances with true labels (in the experiments of this work it constitutes 10%), and D_U is a substantially larger set of unlabeled instances for which weak (noisy) labels are available or generated. Formally, let the input space be X and the label space be $Y = \{1, \ldots, K\}$. The labeled set is defined as $D_L = \{(x_i, y_i)\}_{i=1}^{N_L}$, where $y_i \in Y$ denotes the true label, and the unlabeled pool as $D_U = \{(x_i)\}_{i=1}^{N_U}$, where $N_U \gg N_L$.

Access to weak labeling functions [4] may also be available, providing noisy labels for instances in D_U . Let $f_w: X \rightarrow Y$ represent a weak labeling function. The initial weakly labeled set D_W is constructed as:

$$D_W = \{(x_i, \tilde{y}_i) | x_i \in D_U, \tilde{y}_i = f_w(x_i)\} \tag{1}$$

where $\tilde{y}_i$ represents the noisy or weak label. If explicit weak labeling functions are not provided, D_W can be obtained by training an initial model on D_L and labeling the unlabeled data with this model to serve as pseudo-labels.

3.3. Base Model Training and Objective Function

An initial classification model M_0 is trained using the weakly labeled data D_W. To train the parametrized model $M_0(\theta)$, a weighted cross-entropy loss function is minimized, defined in Eq. (3). To mitigate the effect of label noise in D_W, label smoothing is applied – a technique popularized in Inception architectures [15]. Let $p(y|x;\ \theta)$ be the predicted probability distribution. The smoothed target distribution $q(k|x)$ for class k is defined as:

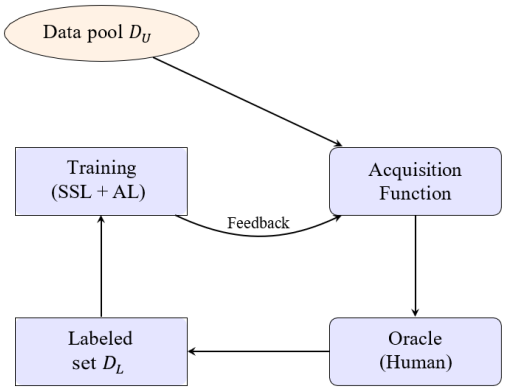


Fig. 1. Diagram of the hybrid active and semi-supervised learning framework.

$$q(k|x) = (1 - \varepsilon)\delta_{k,\bar{y}} + \frac{\varepsilon}{K} \tag{2}$$

where ε is the smoothing parameter and $\delta_{k,\bar{y}}$ is the Kronecker delta.

The following objective function was used for model training, minimizing the weighted cross-entropy between the predicted and smoothed distributions:

$$\mathcal{L}_{base}(\theta) = \frac{1}{|D_W|} \sum_{(x,\bar{y}) \in D_W} w(x) \sum_{i=1}^{K} q(k|x) \log\big(p(k|x;\theta)\big) \tag{3}$$

In Eq. (3), $w(x)$ is a per-sample confidence weight that down-weights instances whose weak labels are likely to be erroneous, and $q(\cdot\ |x)$ is the smoothed target distribution defined above. The remaining symbols follow standard cross-entropy convention: θ are the network parameters, L is the number of classes, $p(k|x;\ \theta)$ is the predicted class probability, and $|D_W|$ is the size of the weakly labeled set. The result is an initial model M_0 that captures some patterns from the data but is likely suboptimal due to label noise.

3.4. Active Query Selection

The framework then enters the active learning loop, corresponding to the Acquisition Function block in Fig. 1. At each iteration, the current model M_i is used to assess the pool for informative samples. An acquisition function $a(x;\ M_i)$ is defined to score instances by their expected utility upon labeling. Since weak labels are available, the strategy targets points where the model is likely incorrect or uncertain, and the weak label may be incorrect. A classic entropy-based sampling strategy [11] combined with weak label disagreement is used. Intuitively, the score should be high for samples

that are both uncertain (high predictive entropy) and disagree with the weak prior, focusing oracle effort where the model and the weak label conflict. The acquisition function is defined as:

$$a(x; M_i) = \mathcal{H}(p(y|x; \theta_i)) + \lambda_{dis} 1(argmax_k\{p(k|x; \theta_k) \neq \tilde{y}_{weak}\}) \quad (4)$$

where $\mathcal{H}(p(y|x; \theta_i)) = -\sum_{k=1}^{K} p_k \log(p_k)$ denotes the Shannon entropy of the predictive distribution, and $1(\cdot)$ is the indicator function. The query batch $Q_i \subset D_U$ is then formed by selecting the b samples with the highest acquisition values, since the score is computed independently per sample.

These instances are sent to a human annotator (the *Oracle* block in Fig. 1) to obtain true labels. The resulting labels are added to the labeled set D_L, and the weak label set is updated for these points (weak labels are replaced by true labels).

3.5. Semi-Supervised Label Refinement

After the queries, a semi-supervised learning phase is carried out to update the model using all available data, corresponding to the *Training (SSL + AL)* block in Fig. 1. The expanded labeled set D_L and the remaining unlabeled set are used. The model training objective at iteration $i + 1$ becomes a joint optimization task combining supervised losses on high-quality labels and unsupervised consistency losses on the remaining data, building on state-of-the-art semi- supervised learning approaches [7]. The total loss $\mathcal{L}_{total}$ is defined as:

$$\mathcal{L}_{total}(\theta) = \mathcal{L}_s(D_L) + \lambda_u \mathcal{L}_u(D_U \cup D_W) \quad (5)$$

where $\mathcal{L}_s(D_L)$ is the standard cross-entropy loss on the expanded labeled set D_L, computed as:

$$\mathcal{L}_s(D_L) = -\frac{1}{|D_L|}\sum_{(x,\bar{y})\in D_L} \log\big(p(y|x; \theta)\big) \quad (6)$$

For the unsupervised term $\mathcal{L}_u$, pseudo-labeling with a hard confidence threshold is used, as implemented in the FixMatch algorithm [7]. The intuition is to trust the model only on samples it is confident about: when the maximum predicted probability exceeds threshold τ, the argmax is treated as a hard pseudo-label and the model is trained against it with standard cross-entropy; samples below the threshold are ignored, which prevents propagation of low-confidence errors. Formally, let $\hat{y} = argmax_k\big(p(k|x; \theta)\big)$ denote the pseudo-label; then

$$\mathcal{L}_u(D_U) = -\frac{1}{N_U}\sum_{x\in D_U} 1\left(\max_k p(k|x; \theta) \geq \tau\right) \log\big(p(\hat{y}|x; \theta)\big). \quad (7)$$

Additionally, the model uses the remaining weak labels in D_W for regularization. If the weak label of an unlabeled sample agrees with the model's prediction, this reinforces the prediction; if it contradicts it, the model may flag this as uncertainty. By applying a strategy similar to WARM (Weakly-supervised Active Refinement Model), the weak label probabilities can be iteratively adjusted based on the model's improved understanding. The logic of this approach is that the initial noisy labels are treated as probabilistic estimates. By actively querying the expert (oracle) for true labels on the most uncertain data points, the system uses this new information not only to update the classifier weights, but also to globally refine and correct the remaining weak labels in pool D_W. This allows substantial improvement of the training set quality with a minimal number of human queries.

3.6. Iteration

The active query selection and semi-supervised label refinement steps are repeated for several rounds or until the labeling budget is exhausted, closing the feedback loop shown in Fig. 1. Each iteration yields a progressively better model with a small increment of labeled data. By integrating semi-supervised learning, the model extracts maximum benefit from the unlabeled examples between queries.

3.7. Computational Complexity

Acquisition step. Computing the acquisition score $a(x; M_i)$ for every sample in the unlabeled pool requires one forward pass per sample, giving a cost of $\mathcal{O}(|D_U| \cdot C_{fwd})$, where C_{fwd} is the cost of a single ResNet-18 inference. Selecting the top b samples from the resulting scores can be performed with a partial sort in $\mathcal{O}(|D_U| \log(b))$, which is negligible relative to the forward passes for typical pool sizes.

Oracle annotation. Exactly b samples are queried per round, independently of the size of D_U. This step represents the human in-the-loop bottleneck that the framework explicitly aims to minimize and contributes a constant cost per round.

Training step. The supervised loss $\mathcal{L}_s$ is evaluated on the labeled set D_L at a cost of $\mathcal{O}\big(|D_L| \cdot C_{step}\big)$ per epoch, where C_{step} denotes the cost of a single forward and backward pass. The unsupervised FixMatch loss $\mathcal{L}_u$ requires both a weakly and a strongly augmented forward pass per unlabeled sample, contributing $\mathcal{O}\big(|D_U| \cdot 2C_{step}\big)$

per epoch. Therefore, the SSL component dominates the training cost.$|D_U| \gg |D_L|$

Comparison with baselines. Per round:

- *Pure SSL* incurs only the training cost (no acquisition step) and is therefore the cheapest method per round, but receives no targeted feedback about which samples are most informative.
- *Baseline AL* incurs the acquisition cost plus supervised training restricted to D_L. Because $|D_L|$ is small, its per- round training cost is low, but the unlabeled pool is left unused.
- The *proposed hybrid* combines acquisition with full SSL training on $D_L \cup D_U$ and is therefore the most expensive per round, by a factor that asymptotically equals the ratio $|D_U|/|D_L|$ amplified by the FixMatch dual-augmentation overhead.

This additional per-round cost is the price of exploiting the unlabeled pool. As shown in Section 4.2, the hybrid framework reaches the asymptotic accuracy of the baselines with approximately 500 fewer oracle queries; when expert annotation is the dominant cost of the system, the amortized cost per saved query is favorable for the hybrid approach, since the additional per-round compute is spent once while each saved annotation represents a recurring reduction in expert effort.

Empirical training time. The experiments reported in Section 4.2 were run on a single Apple M1 Pro processor without dedicated GPU acceleration. Under this configuration, one training epoch takes approximately 20 minutes for the proposed Hybrid method, which is the most expensive of the three compared methods because it performs the FixMatch consistency loss on the unlabeled pool in addition to supervised training on D_L. Wall-clock measurements naturally depend on hardware and implementation; on dedicated GPU accelerators the per-epoch cost is expected to drop by an order of magnitude or more, while the relative ordering across the three compared methods follows the asymptotic analysis above.

Scalability to large and high-dimensional data. The asymptotic analysis above identifies three scaling axes that determine how the framework extends beyond the CIFAR-10-scale regime studied here.

Large unlabeled pool ($|D_U|$ large). The acquisition step requires one forward pass per pool sample and thus dominates per-round cost when $|D_U|$ is in the millions. The standard mitigation is stochastic subsampling: a random subset $\widetilde{D}_U \subset D_U$ of fixed size is scored at each round, and the top-b samples are selected from $\widetilde{D}_U$ rather than from D_U . This reduces acquisition cost to $\mathcal{O}\left(|D_U|C_{fwd}\right)$ at the price of slightly suboptimal selections; theoretical and empirical analyses indicate that the loss is small for sufficiently large $\left|\widetilde{D}_U\right|$ [11]. The training step, dominated by the FixMatch consistency loss on D_U, can be controlled by limiting the number of unlabeled samples per training batch independently of $|D_U|$, since the SSL signal is computed mini-batch-wise.

High-dimensional inputs. For inputs of higher resolution or modality (e.g., 3D medical volumes, satellite imagery), both C_{fwd} and C_{step} grow with the per-sample compute of the backbone. The framework itself is architecture-agnostic: ResNet-18 can be replaced by any classifier producing a probability vector $p(k|x;\theta)$, including more parameter-efficient backbones (e.g., MobileNet variants) or larger backbones combined with mixed-precision training. The FixMatch dual-augmentation overhead, which doubles the forward-pass cost on the unlabeled set, is the most sensitive component; in memory-constrained settings the unlabeled batch size or the augmentation strength can be reduced before the supervised pipeline is touched.

Many classes (K large). The Shannon entropy in the acquisition score and the label-smoothing target both scale as $\mathcal{O}(K)$ per sample, so even for K in the thousands the cost remains negligible relative to a single forward pass. This axis is therefore not a practical bottleneck.

In summary, the cost structure of the framework scales gracefully with $|D_U|$ and K, and depends on the choice of backbone for high-dimensional inputs. A full empirical study at industrial scale (millions of unlabeled samples, high-resolution inputs, hundreds of classes) is beyond the scope of this paper and is left as future work.

4. Experiments and Results

The proposed framework is evaluated on the CIFAR-10 image classification benchmark, with a multi-seed experimental protocol described in detail below.

4.1. Dataset and Setup

The CIFAR-10 image classification dataset was used [16]. A weak supervision scenario is simulated starting with 5,000 training image samples with true labels. The remaining samples are treated as unlabeled or weakly labeled. To ensure correct validation and testing, the full dataset was split into three independent subsets. The data split was performed according to the following scheme:

- Training set: Constitutes 80% of the total data. This portion is used for directly training the model parameters (neural network weights). It includes both the initial small labeled set (D_L) and the large unlabeled data pool

(D_U).

- Validation set: Constitutes 4% of the data. This set is used for hyperparameter tuning (such as ε, λ_{dis}, τ) and for selecting the best model during training (model checkpointing).
- Test set: Constitutes 16% of the data. This set is used exclusively for the final evaluation of algorithm performance after all training phases are complete.

The ResNet-18 architecture is initialized with random weights using the He method – the standard choice for ReLU-based deep networks [14], since it preserves activation variance across layers and mitigates the vanishing/exploding gradient problem in deep residual architectures. Optimization is performed using Stochastic Gradient Descent (SGD) with momentum 0.9 and an initial learning rate of 0.03. These values follow the FixMatch protocol for CIFAR-10 [7], where they were tuned for ResNet-style backbones with consistency-based SSL losses; SGD with momentum 0.9 remains the canonical optimizer for ResNet training [14]. Reusing this configuration is a deliberate methodological choice: it ensures that the SSL component of the proposed framework operates in its native regime, and that accuracy differences observed across the three compared methods (Pure SSL, Baseline AL, Proposed Hybrid) reflect the learning paradigm rather than optimizer-level tuning. The remaining method-specific hyperparameters – the label-smoothing factor ϵ, the disagreement weight λ_{dis}, the unsupervised loss weight λ_u, and the FixMatch confidence threshold τ – were selected on the validation set.

Table 2. Test accuracy on CIFAR-10 over 20 active learning rounds. Values at key rounds (marked †) are reported as mean ± standard deviation over three independent seeds {42, 123, 2024}; intermediate rounds (1–4) are reported as single-seed values from the trend run (seed 42) and are shown for trajectory visualization only

Round	Labeled	Pure SSL	Baseline AL	Proposed Hybrid
0†	5000	0.4051 ± 0.0193	0.3948 ± 0.0198	0.4051 ± 0.0193
1	5100	0.4729	0.4515	0.4671
2	5200	0.5165	0.5108	0.5092
3	5300	0.5681	0.5397	0.5654
4	5400	0.5936	0.5844	0.5813
5†	5500	0.6093 ± 0.0028	0.6009 ± 0.0058	**0.6183 ± 0.0169**
10†	6000	0.6428 ± 0.0026	0.6461 ± 0.0146	**0.6563 ± 0.0100**
15†	6500	0.6657 ± 0.0107	0.6665 ± 0.0058	**0.6792 ± 0.0097**
20†	7000	0.6687 ± 0.0056	0.6730 ± 0.0139	0.6770 ± 0.0078

Three main methods are compared:

- Pure SSL (FixMatch): a semi-supervised learning method without active sample selection.
- Baseline AL: a standard active learning method (entropy-based sampling).
- Proposed Hybrid (AL + SSL): the proposed approach combining active selection with SSL refinement.

Statistical Methodology

To assess robustness, each of the three methods is run independently on three random seeds, {42, 123, 2024}, with all sources of stochasticity (random number generators in NumPy and PyTorch, CUDA kernels) seeded explicitly and cudnn.deterministic enabled. A full 20-round trajectory is collected on the trend seed (seed 42), and additional full runs on the remaining two seeds are evaluated at the key rounds {0, 5, 10, 15, 20}. For these key rounds, accuracy is reported as mean ± sample standard deviation (Bessel-corrected, ddof = 1) across the three seeds. Pairwise comparisons between methods at the same round are performed using a paired *t*-test on the per-seed accuracies, with $N = 3$ pairs and df = 2; *p*-values reported below should be interpreted accordingly, since with only three seeds the test has limited statistical power.

4.2. Results

An extended experiment of 20 active learning rounds was conducted, during which 100 samples were selected per round for manual annotation. The test accuracy dynamics for the three methods are presented numerically in Table 2 and visualized in Fig. 2.

The intermediate-budget regime (Rounds 5–15) is where the proposed framework shows its most consistent behaviour. At Round 5, the hybrid method outperforms Baseline AL on all three seeds (mean gap $+1.74$ pp) and outperforms Pure SSL on all three seeds ($+0.91$ pp); at Round 10 the gap is $+1.01$ pp over Baseline AL and $+1.35$ pp over Pure SSL; at Round 15 the gap is $+1.27$ pp over Baseline AL and +1.35 pp over Pure SSL, with the latter approaching significance (p = 0.07, paired t-test). With $N = 3$ seeds the differences do not reach the conventional $\alpha = 0.05$ threshold, but the direction of the effect is consistent across seeds and rounds, as can be seen in the trajectories of Fig. 2.

A practically more meaningful way to interpret the trajectory is in terms of label efficiency rather than asymptotic accuracy. At Round 15, with $|D_L| = 6500$ labels, the hybrid framework reaches a test accuracy of 0.6792 ± 0.0097,

which is already higher than the asymptotic accuracies attained by Baseline AL (0.6730 ± 0.0139) and Pure SSL (0.6687 ± 0.0056) at Round 20 with $|D_L| = 7000$. In other words, the hybrid framework matches or exceeds the converged accuracy of both baselines using approximately 500 fewer expert annotations. When expert labelling dominates the cost of the system, this label-efficiency gain is the operationally relevant quantity, even when the final-round accuracy differences are within statistical noise.

By Round 20 all three methods saturate near a common data ceiling ($0.6687 - 0.6770$) and the pairwise differences fall within the standard deviations: Hybrid versus Baseline AL has $p = 0.45$, Hybrid versus Pure SSL has $p = 0.34$. This convergence is also evident in the right-hand portion of Fig. 2, where all three trajectories enter overlapping shaded bands. The interpretation is not that the methods are equivalent; it is that the CIFAR-10 task at this annotation budget is no longer label-limited. Studying the same combination on more challenging or more weakly labelled problems is a natural direction for future work.

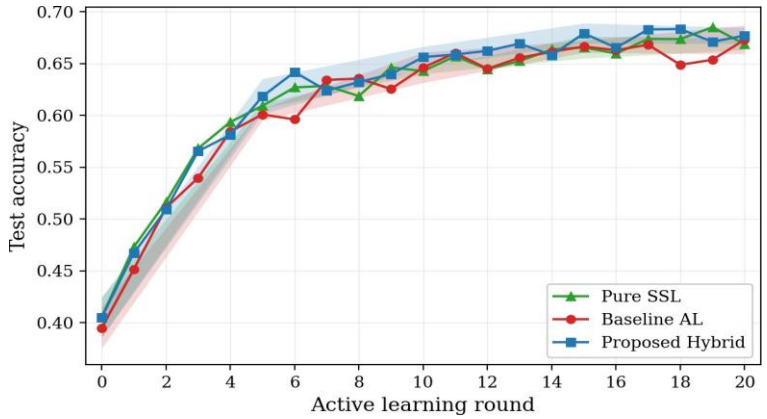


Fig. 2. Test accuracy on CIFAR-10 across 20 active learning rounds for the three compared methods. Solid lines show single-seed trajectories (seed 42); shaded bands at the key rounds {0, 5, 10, 15, 20} denote ±1 standard deviation across three independent seeds

5. Conclusions

The experimental study on CIFAR-10, repeated across three independent seeds at the key rounds, supports two main empirical claims. First, at intermediate labelling budgets (Rounds 5–15) the hybrid framework consistently leads both baselines on average, with the most uniform signal at Round 5 (+1.74 pp over Baseline AL across all three seeds) and the largest absolute gap at Round 15 (+1.27 pp over Baseline AL, +1.35 pp over Pure SSL). With $N = 3$ seeds these differences do not reach the conventional $\alpha = 0.05$ threshold of statistical significance, but the direction of the effect is reproducible across seeds. Second, and more practically meaningful, the framework offers a clear label-efficiency advantage: at Round 15 the hybrid method already reaches an accuracy (0.6792 ± 0.0097) that the two baselines only attain at Round 20 (0.6730 ± 0.0139 for Baseline AL, 0.6687 ± 0.0056 for Pure SSL). At fixed final accuracy, the hybrid framework therefore reduces the required expert annotation budget by approximately 500 labels.

By Round 20 all three methods saturate near a common data ceiling and the remaining differences fall within statistical noise ($p = 0.45$ for Hybrid versus Baseline AL, $p = 0.34$ for Hybrid versus Pure SSL, paired t-test), as visible in the converging trajectories of Fig. 2. The interpretation is that CIFAR-10 at this budget is no longer label-limited rather than that the methods are equivalent. The hybrid framework also incurs a higher per-round computational cost than the baselines (Section 3.7), but this overhead is asymptotically bounded and is offset, when expert annotation is the dominant cost, by the label-efficiency gain reported above.

Two limitations follow directly from this analysis and motivate future work. First, statistical power with $N = 3$ seeds is insufficient to detect single-percentage-point effects; replication with more seeds, additional acquisition strategies, or stronger SSL backbones is needed to establish significance at the $\alpha = 0.05$ level. Second, the saturation observed at Round 20 indicates that the proposed framework should be evaluated on more challenging or more weakly labelled tasks – such as medical image analysis or document classification with realistic noisy weak labels – where the data ceiling is harder to reach and the operational value of label efficiency is correspondingly higher.

Generalization to other datasets. The primary objective of this work was to develop the mathematical foundation and the algorithmic structure of the proposed hybrid framework. A single, widely studied benchmark (CIFAR-10) was chosen for the empirical evaluation precisely because it has been used by many of the methods to which the proposed approach is compared, which makes a like-for-like comparison possible. Generalization to other settings is being explored as the next stage of this work: in particular, evaluation on CIFAR-100 [16] (more classes, finer-grained categories and a less saturated regime than CIFAR-10) and on a real-world weakly supervised text dataset (such as IMDB-style movie review classification) is planned as part of follow-up work. These domains exercise different properties of the framework that are not stressed by CIFAR-10 – namely, larger K in the acquisition score and label-smoothing computation, and more realistic, non-uniform weak-label noise patterns – and will allow the label-efficiency claim reported here to be tested under

operating conditions closer to those motivating the framework in the first place (medical and document classification with limited expert annotation).

Limitations and ethical considerations. Two limitations of this study, in addition to those already discussed, deserve explicit mention. First, the experimental study uses a single, class-balanced benchmark with synthetic weak labels; the obtained quantitative results should not be extrapolated directly to imbalanced or domain-specific datasets without further validation. Second, although the present setting – a balanced classification benchmark with comparably sized classes – minimizes the risk of *bias propagation from weak labels*, this risk does become substantive once the framework is applied to imbalanced datasets, datasets with sensitive demographic attributes, or private data. In such settings, weak labels typically encode the biases of the heuristic or crowd-sourcing source from which they were generated, and both the active-acquisition step and the FixMatch self-training step can amplify these biases if no countermeasures are applied. Practical deployment of the framework in those settings should therefore include auxiliary monitoring of (i) per-class and per-group accuracy, not only aggregate accuracy; (ii) the demographic or class composition of the oracle-queried subset across rounds; and (iii) standard fairness and privacy safeguards appropriate to the data domain. These considerations fall outside the scope of the methodological contribution of this paper but are essential for any application of the framework to imbalanced or sensitive data, and are part of the planned future evaluation on non-CIFAR datasets.

All the Declarations and Statements

Author Contributions Statement

Kostiantyn O. Minkov: Conceptualization, Methodology, Software Implementation, Data Curation, Model Training, Validation, Formal Analysis, Visualization, and Writing – Original Draft: proposed the research idea, implemented the hybrid framework, ran the experiments, analyzed the results, and drafted the manuscript.
Igor V. Malyk: Conceptualization, Methodology, Supervision, Project Administration, and Writing – Review and Editing: supervised the research, refined the methodology, and reviewed and edited the manuscript.

All authors have read and agreed to the published version of the manuscript.

Conflict of Interest Statement

The authors declare no conflicts of interest.

Funding Declaration

None.

Data Availability Statement

This study analyzed a publicly available dataset. The CIFAR-10 dataset can be found here: https://www.cs.toronto.edu/~kriz/cifar.html, accessed on 25 March 2026.

Ethical Declarations

This study did not involve any human participants or animals; therefore, ethical approval was not required.

Acknowledgment

The authors wish to thank the colleagues of the Department of Mathematical Problems of Control and Cybernetics at Chernivtsi National University for valuable discussions.

Declaration of Generative AI in Scholarly Writing

During the preparation of this work the authors used a generative AI-based language tool in order to improve the language and readability of the manuscript. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Abbreviations

The following abbreviations are used in this manuscript:

AL - Active Learning
SSL NN - Semi-Supervised Learning Neural Network
ResNet-18 - 18-layer Residual Network [14]
CIFAR-10 - Canadian Institute for Advanced Research image dataset, 10-class subset [16]
SGD - Stochastic Gradient Descent
FixMatch - SSL algorithm based on consistency regular- ization with confidence-thresholded pseudo- labels [7]

MixMatch - SSL algorithm combining consistency regu- larization with label sharpening
VAAL WARM - Variational Adversarial Active Learning [9] Weakly-supervised Active Refinement Model
MLC - Meta Label Correction [12]

Appendix

None

References

[1] S. Budd, E. C. Robinson, and B. Kainz, "A survey on active learning and human-in-the-loop deep learning for medical image analysis," *Medical Image Analysis*, vol. 71, p. 102062, 2021.

[2] A. Siddhant and Z. C. Lipton, "Deep Bayesian Active Learning for Natural Language Processing: Results of a Large-Scale Empirical Study," in Proc. 2018 Conf. Empirical Methods in Natural Language Processing (EMNLP), Brussels, Belgium, pp. 2904–2909, 2018. doi:10.18653/v1/D18-1318.

[3] M. Gao, Z. Zhang, G. Yu, S. O. Arik, L. S. Davis, and T. Pfister, "Consistency-Based Semi-Supervised Active Learning: Towards Minimizing Labeling Cost," in Proc. European Conference on Computer Vision (ECCV), Lecture Notes in Computer Science, pp. 510-526, 2020, doi: 10.1007/978-3-030-58607-2_30.

[4] A. Ratner et al., "Snorkel: Rapid training data creation with weak supervision," in *Proc. VLDB Endowment*, vol. 11, no. 3, pp. 269–282, 2017.

[5] V. C. Raykar et al., "Learning from Crowds," *Journal of Machine Learning Research*, vol. 11, pp. 1297–1322, 2010.

[6] M. Ren et al., "Learning to Reweight Examples for Robust Deep Learning," in *International Conference on Machine Learning (ICML)*, pp. 4334–4343, 2018.

[7] K. Sohn et al., "FixMatch: Simplifying Semi-Supervised Learning with Consistency and Confidence," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, pp. 596–608, 2020.

[8] A. Tarvainen and H. Valpola, "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 30, 2017.

[9] S. Sinha, H. Ebrahimi, and T. Darrell, "Variational Adversarial Active Learning," in *Proc. IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 5972–5981, 2019.

[10] S. Song, D. Berthelot, and A. Rostamizadeh, "Combining MixMatch and Active Learning for Better Accuracy with Fewer Labels," arXiv preprint arXiv:1912.00594, 2019.

[11] B. Settles, "Active Learning Literature Survey," *Computer Sciences Technical Report 1648*, University of Wisconsin– Madison, 2009.

[12] G. Zheng, A. H. Awadallah, and S. Dumais, "Meta Label Correction for Noisy Label Learning," in *Proc. AAAI Conference on Artificial Intelligence*, vol. 35, no. 12, pp. 11053–11061, 2021.

[13] W. Ma, O. Karakus, and P. L. Rosin, "Integrating Semi-Supervised and Active Learning for Semantic Segmentation," *arXiv preprint arXiv:2501.19227*, 2025.

[14] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016.

[15] C. Szegedy et al., "Rethinking the Inception Architecture for Computer Vision," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pp. 2818–2826, 2016.

[16] A. Krizhevsky, "Learning Multiple Layers of Features from Tiny Images," *Technical Report*, University of Toronto, 2009.

[17] J. Jimenez et al., "Computational Evaluation of the Combination of Semi-Supervised and Active Learning for Histopathology Image Segmentation with Missing Annotations," in *Proc. ICCVW*, 2023.

[18] D. Levkivskyi, V. Vysotska, L. Chyrun, Y. Ushenko, D. Uhryn, and C. Hu, "Agile Methodology of Information Engineering for Semantic Annotations Categorization and Creation in Scientific Articles Based on NLP and Machine Learning Methods," *International Journal of Information Engineering and Electronic Business (IJIEEB)*, vol. 17, no. 2, pp. 1–50, 2025. doi:10.5815/ijieeb.2025.02.01.

[19] V. Vysotska, K. Przystupa, Y. Kulikov, S. Chyrun, Y. Ushenko, Z. Hu, and D. Uhryn, "Recognizing Fakes, Propaganda and Disinformation in Ukrainian Content Based on NLP and Machine-Learning Technology," *In- ternational Journal of Computer Network and Information Security (IJCNIS)*, vol. 17, no. 1, pp. 92–127, 2025. doi:10.5815/ijcnis.2025.01.08.

[20] V. Vysotska, K. Przystupa, L. Chyrun, S. Vladov, Y. Ushenko, D. Uhryn, and Z. Hu, "Disinformation, Fakes and Propaganda Identifying Methods in Online Messages Based on NLP and Machine Learning Methods," *In- ternational Journal of Computer Network and Information Security (IJCNIS)*, vol. 16, no. 5, pp. 57–85, 2024. doi:10.5815/ijcnis.2024.05.06.

[21] W. Huang, Z. Xiong, C. Liu, and X. X. Zhu, "Hierarchical Semi-Supervised Active Learning for Remote Sensing," *arXiv preprint arXiv:2511.18058*, 2025.

[22] H. Roda and A. B. Geva, "Semi-supervised active learning using convolutional auto-encoder and contrastive learn- ing," *Frontiers in Artificial Intelligence*, vol. 7, p. 1398844, 2024. doi:10.3389/frai.2024.1398844.

[23] M. Li and C. Zhu, "Noisy Label Processing for Classification: A Survey," *arXiv preprint arXiv:2404.04159*, 2024.

Authors' Profiles



Kostiantyn O. Minkov was born in Ukraine on January 29, 2000. In 2019 he graduated with honors from Chernivtsi Polytechnic College (Chernivtsi, Ukraine). In 2021 he received his Bachelor's degree from Yuriy Fedkovych Chernivtsi National University (Chernivtsi, Ukraine) in the specialty of Computer Science. In 2022 he received his Master's degree from Yuriy Fedkovych Chernivtsi National University (Chernivtsi, Ukraine). Currently, he is a PhD student at Yuriy Fedkovych Chernivtsi National University in the specialty of Software Engineering. His PhD research focuses on the application of Active Learning methods in weakly supervised learning tasks. Since the beginning of his PhD studies, he has been working as an assistant at the Department of Mathematical Problems of Control and Cybernetics (Yuriy Fedkovych Chernivtsi National University, Chernivtsi, Ukraine). Research interests are Active Learning, weakly supervised machine learning, computer vision, and deep learning.

He has published research papers in peer-reviewed journals.



Igor V. Malyk was born in Ukraine on July 19, 1983. In 2005 he graduated from Yuriy Fedkovych Chernivtsi National University (Chernivtsi, Ukraine) in the specialty of statistics. In 2009 he defended his PhD thesis at the Institute of Cybernetics (Kyiv, Ukraine) in the specialty of theoretical foundations of computer science and cybernetics. In 2018 he defended his doctoral thesis at the Institute of Cybernetics (Kyiv, Ukraine) in the specialty of theoretical foundations of computer science and cybernetics.

From 2005 to 2017, he worked as an assistant, and later as an associate professor at the Department of Systems Analysis and Insurance and Financial Mathematics (Yuriy Fedkovych Chernivtsi National University, Chernivtsi, Ukraine). Since September 2017, he has been working as an associate professor at the Department of Mathematical Problems of Control and Cybernetics (Yuriy Fedkovych Chernivtsi National University, Chernivtsi,

Ukraine), and since 2024, he has been a full professor and head of the department. Research interests are self-adaptive algorithms in machine learning, random processes, stochastic differential equations, deep learning, and neural networks.

Prof. Malyk is a member of the CHEST Journal editorial team (statistical reviewer), Multidisciplinary Digital Publishing Institute (MDPI) editorial team, American Mathematical Society, and International Society for Computational Biology (ISCB).

How to cite this paper

Kostiantyn O. Minkov, Igor V. Malyk, "A Hybrid Active and Semi-Supervised Learning Framework for Classification with Minimal Labeled Data", International Journal of Intelligent Systems and Applications(IJISA), Vol.18, No.3, pp.81-91, 2026. DOI:10.5815/ijisa.2026.03.05