

Human-Corrected Labels Learning: Enhancing Labels Quality via Human Correction of VLMs Discrepancies

Zhongnian Li^{1,2}, Lan Chen¹, Yixin Xu¹, Shi Xu¹, Xinzheng Xu^{1,3*}

¹School of Computer Science and Technology / School of Artificial Intelligence, China University of Mining and Technology, Xuzhou, China

²Mine Digitization Engineering Research Center of the Ministry of Education, China University of Mining and Technology, Xuzhou, China

³State Key Lab. for Novel Software Technology, Nanjing University, Nanjing, China
{zhongnianli, chenlan, yixinxu, shixu, xxzheng}@cumt.edu.cn

Abstract

Vision-Language Models (VLMs), with their powerful content generation capabilities, have been successfully applied to data annotation processes. However, the VLM-generated labels exhibit dual limitations: low quality (i.e., label noise) and absence of error correction mechanisms. To enhance label quality, we propose Human-Corrected Labels (HCLs), a novel setting that efficient human correction for VLM-generated noisy labels. As shown in Figure 1(b), HCL strategically deploys human correction only for instances with VLM discrepancies, achieving both higher-quality annotations and reduced labor costs. Specifically, we theoretically derive a risk-consistent estimator that incorporates both human-corrected labels and VLM predictions to train classifiers. Besides, we further propose a conditional probability method to estimate the label distribution using a combination of VLM outputs and model predictions. Extensive experiments demonstrate that our approach achieves superior classification performance and is robust to label noise, validating the effectiveness of HCL in practical weak supervision scenarios.

Code — <https://github.com/Lilianach24/HCL>

Extended version — <http://arxiv.org/abs/2511.09063>

Introduction

Recently, pre-trained VLMs (Radford et al. 2021; Li et al. 2022; Liu et al. 2023) have demonstrated remarkable zero-shot and few-shot recognition capabilities, making them increasingly popular for data annotation tasks. By aligning images and textual labels in a shared embedding space, VLMs can assign labels to images without task-specific training or human correction. This method offers a scalable and cost-effective alternative to traditional annotation procedures, which typically require large amounts of labor and domain expertise. Compared to classical weakly supervised learning techniques (Yan and Guo 2023; Chen, Yao, and Tang 2023; Chen et al. 2022), such as semi-supervised learning (Van Engelen and Hoos 2020; Cao, Brbic, and Leskovec 2021; Wei et al. 2024), partial-label learning (Feng et al. 2020b; Xia

et al. 2023; Li et al. 2024), and noisy-label learning (Menon et al. 2015; Ghosh, Kumar, and Sastry 2017; Han et al. 2020), VLM-based annotation eliminates the need for initial labeled data, significantly reducing annotation cost and effort (Menghini, Delworth, and Bach 2023; Li et al. 2025), and enabling rapid dataset construction in new domains.

However, the VLM-generated labels exhibit dual limitations: low quality (i.e., label noise) and absence of error correction mechanisms. Most notably, existing approaches typically rely on a single VLM to generate labels (Kabra et al. 2024), which can lead to biased or incorrect predictions. As illustrated in Figure 1(a), the VLM may predict “Crocodile head” as the label with the highest confidence, while the ground-truth label is “Gramophone” (Wang et al. 2022; Menghini, Delworth, and Bach 2023). Moreover, current methods lack mechanisms to correct such errors, making it difficult to ensure label reliability. These limitations highlight the need for new labels strategies that combine the efficiency of VLMs with mechanisms for error detection and correction.

To address these issues, we propose a setting called Human-Corrected Labels (HCLs). HCLs strategically deploys human correction only for instances with VLMs discrepancies, thereby achieving both higher-quality annotations and reduced labor costs. Specifically, during annotation, our HCL method automatically accepts labels when multiple VLMs yield consistent predictions, invokes human correction solely for samples where their predictions are inconsistent. As illustrated in Figure 1 (b), when annotating an image of a “Gramophone”, correction is only required if there exists a discrepancy among the VLMs. This targeted human correction drastically reduces annotation costs while ensuring high annotation quality. To the best of our knowledge, HCL is the first setting that leverages VLMs discrepancies to correct labels.

In this paper, we develop a risk-consistent learning method that unifies VLM predictions and human-corrected labels in a principled way. Specifically, we derive a risk-consistent estimator that incorporates both human corrected labels and VLMs predictions, enabling effective learning under label uncertainty. Furthermore, we introduce a conditional probability formulation to estimate the label distribution using a combination of VLM outputs and model

*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

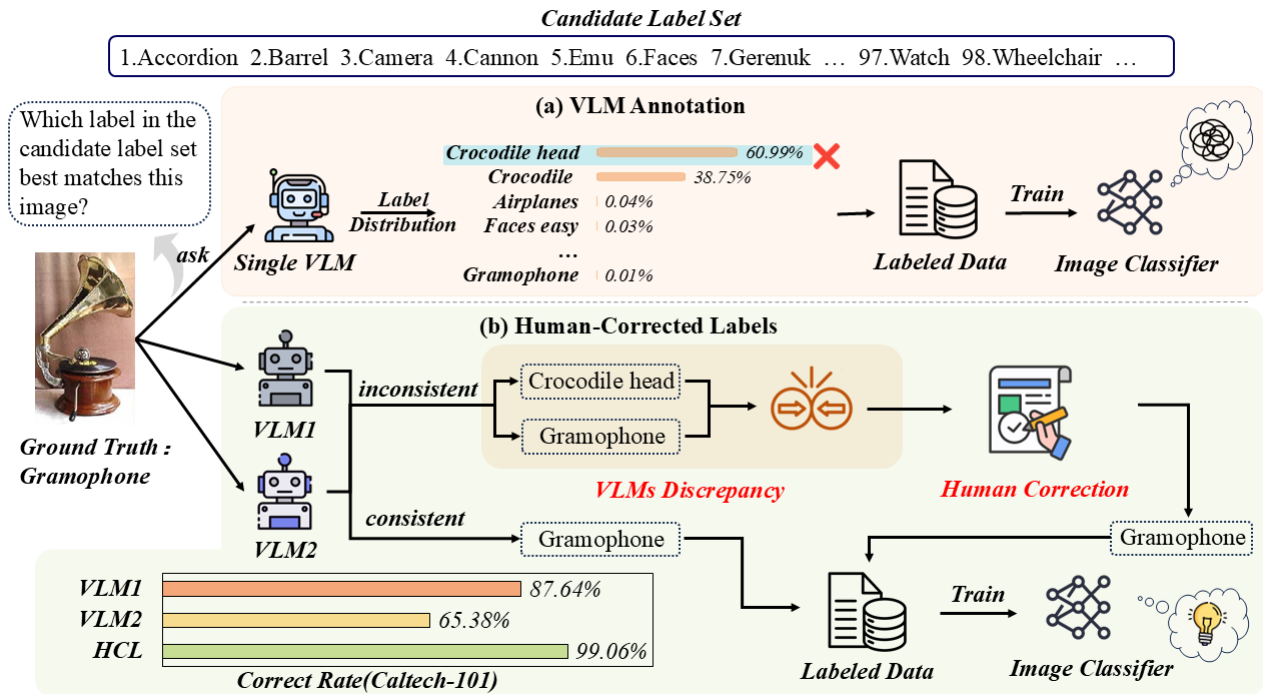


Figure 1: A comparison between traditional VLM annotation and Human-Corrected Labels (HCLs). The zero-shot results depicted are obtained using CLIP with ViT-L/14. The example images and categories are taken from the Caltech-101 dataset. HCL deploys human correction only for instances with VLMs discrepancies, achieving both higher-quality annotations and reduced labor costs.

predictions, enhancing label robustness under weak supervision. Extensive experiments on various datasets clearly demonstrate the effectiveness of the proposed HCL learning method.

Our main contributions are summarized as follows:

- We propose a novel weakly supervised annotation setting, i.e., Human-Corrected Labels, which combines multiple VLMs with selective human correction. This setting enables high-quality, scalable annotation with minimal human effort.
- We propose a risk-consistent method that models the conditional label distribution under the HCLs setting, leveraging both VLM predictions and human corrections.
- Extensive experiments on standard benchmarks demonstrate that our method significantly improves classification performance and robustness compared to existing weakly supervised baselines.

Related Work

Weakly Supervised Learning

Weakly supervised learning aims to train predictive models using data with limited or imperfect annotations, and encompasses three main methods: semi-supervised learning (Chen et al. 2023), partial-label learning (Chapel, Alaya, and Gasso 2020; Wang et al. 2023), and noisy label learning

(Xia et al. 2022; Gong et al. 2022). Semi-supervised learning leverages a small amount of labeled data and abundant unlabeled data to improve performance. Key strategies include entropy minimization (Grandvalet and Bengio 2004; Lee et al. 2013), consistency regularization (Sajjadi, Javanmardi, and Tasdizen 2016; Tarvainen and Valpola 2017; Miyato et al. 2018), and holistic training methods (Berthelot et al. 2019; Sohn et al. 2020; Sosea and Caragea 2023). Partial-label learning assumes each instance is associated with a set of candidate labels, only one of which is correct. Two main strategies are commonly used: average-based disambiguation (Zhang and Yu 2015), and identification-based approaches that treat the ground-truth as a latent variable (Wang, Li, and Zhang 2019; Feng and An 2019). Noisy label learning addresses the challenge of corrupted supervision. Techniques include sample selection (Patel and Sastry 2023), label correction (Albert et al. 2023), and noise-aware regularization (Liu et al. 2020; Sun et al. 2021; Li, Xiong, and Hoi 2021).

While prior work has advanced weakly supervised learning, it typically assumes fixed supervision and overlooks the potential of combining VLMs with human correction. In contrast, we propose the Human-Corrected Label method, which can achieve the collaboration between VLMs and humans.

Dataset	Train Samples	Classes	CLIP	Qwen	VLM-CR	CCP	HCLs
CIFAR100	50,000	100	75.66	45.82	44.02	92.97	96.90
Tiny-ImageNet	100,000	200	72.95	38.72	35.88	94.59	98.06
Caltech-101	6,074	102	87.64	65.38	64.36	98.54	99.06
Food-101	75,750	101	89.77	48.87	49.26	96.48	98.27
EuroSAT	18,900	10	43.21	39.42	26.43	81.82	95.20
DTD	2,800	47	52.27	56.65	48.04	81.33	91.03

Table 1: Annotation accuracy and consistency analysis under the HCLs setting across six benchmarks. Columns show number of training samples and classes, accuracy of CLIP and Qwen labels, VLMs Consistency Rate (VLM CR), Correctness in Consistent Predictions (CCP), and final annotation accuracy after discrepancy-based human correction.

Model Discrepancies

The model discrepancy, defined as the divergence between predictions of different models, has been widely used in weakly supervised learning to estimate label uncertainty and guide pseudo-label handling. Early disagreement-based methods such as co-training (Wang and Zhou 2017) and tri-training (Huang, He, and Huang 2021) leverage multiple classifiers (i.e., often on different feature views or model initializations) to label unlabeled data based on inter-model agreement. More recent methods, such as Co-Teaching+ (Yu et al. 2019), combine model discrepancies with small-loss selection to filter noisy supervision. Dynamic Mutual Training (Feng et al. 2022) further exploits discrepancy for dynamic loss re-weighting, reducing the impact of erroneous VLM-generated labels. In semi-supervised object detection, CrossRectify (Ma et al. 2023) uses discrepancies across two detectors to cross-correct VLM-generated labels and improve detection quality.

However, these methods primarily leverage model discrepancy to improve training robustness, rather than to guide data annotation. Consequently, we utilize discrepancies between multiple pre-trained VLMs (e.g., CLIP (Radford et al. 2021), Qwen (Wang et al. 2024)) to identify inconsistent predictions and selectively invoke human correction.

Methods

In this section, we introduce a detailed description of Human-Corrected Labels. In addition, we propose a risk-consistent estimator to explore and leverage the underlying correlations between the probability distribution of VLM-generated labels and HCLs.

Human-Corrected Labels

The predictions produced by VLMs exhibit dual limitations: low quality (i.e., label noise) and the absence of error correction mechanisms. Specifically, VLM predictions may contain incorrect or ambiguous labels, leading to noisy supervision that can degrade the performance of models trained under weak supervision. As shown in Figure 1(a), VLMs may annotate the instance as “crocodile head” with the highest confidence under a zero-shot setting, which can result in noisy labels.

In this paper, we propose a new setting: Human-Corrected Labels (HCLs), where labels are generated by VLMs, and

annotators only correct instances with inconsistent VLM predictions. For example, as illustrated in Figure 1(b), when the candidate label set includes {“accordion”, “barrel”, “camera”, “cannon”, “emu”, “faces”, “gerenuk”, ... “watch”, “wheelchair”, ...}, for an image of a “gramophone”, the annotator only needs to provide the correct label if the VLMs produce inconsistent predictions; otherwise, the predictions are accepted directly. Compared to traditional manual annotation, HCL leverages VLM knowledge for high-quality labels while cutting annotation costs.

Effectiveness of HCLs. To evaluate the effectiveness of HCLs in improving label quality, we annotate six standard classification benchmarks using CLIP (ViT-L/14) (Radford et al. 2021) and Qwen2.5-VL (7B-Instruct) (Wang et al. 2024), and report results in Table 1. We compute the consistency rate between VLMs and the correctness rate within the consistent predictions. The results indicate that although the agreement between the two VLMs varies across datasets, the labels on which they agree are highly reliable, demonstrating that model consensus is a strong indicator of correctness. Based on this, HCLs performs human correction only on samples with VLM discrepancies. As shown in the “HCLs” column, the resulting labels achieve consistently high accuracy, validating that targeted human correction guided by model discrepancy yields reliable supervision with minimal human effort.

Problem Formalization

Ordinary Labels. In multi-class classification, let $\mathcal{X} \in \mathbb{R}^d$ denote the feature space and $\mathcal{Y} = \{1, \dots, k\}$ denote label space, where d is the feature dimension and $k > 2$ is the number of classes. We assume that $\{(x_i, y_i)\}_{i=1}^N$ are sampled independently from an unknown distribution with density $p(x, y)$. The goal of ordinary multi-class classification is to learn a classifier $f(x) : \mathcal{X} \rightarrow \mathcal{Y}$ that minimizes the classification risk under a multi-class loss $\mathcal{L}(f(x), y)$:

$$R(f) = \mathbb{E}_{x \sim p(x)} \sum_{i=1}^k P(y = i | x) \mathcal{L}[f(x), i] \quad (1)$$

where x denotes an input image, $p(x)$ denotes the density of x , y denotes its ground truth label, and $P(y = i | x)$ denotes the probability that the true label of x is class i .

HCLs. To address the limitations of traditional VLM-generated labels, we propose using multiple VLMs to collaboratively annotate data, leveraging their complementary strengths while detecting potential label noise through discrepancy signals. Given a training set with HCL annotations, denoted as $\mathcal{D}_{HCL} = \{(x_j, Y_j, s_j)\}_{j=1}^N$, where x_j is the input, and $s_j \in \{0, 1\}$ indicates whether the VLM predictions are consistent (i.e., $s_j = 0$) or inconsistent (i.e., $s_j = 1$). The HCL label Y_j is defined as:

$$Y_j = \begin{cases} \hat{y}_j \text{ or } \tilde{y}_j, & \hat{y}_j = \tilde{y}_j \text{ and } s_j = 0 \\ y_j, & \hat{y}_j \neq \tilde{y}_j \text{ and } s_j = 1 \end{cases}, \quad (2)$$

where \hat{y}_j and \tilde{y}_j represent the predicted labels from two different VLMs, respectively.

Our goal is to develop a classifier $f(x)$ that effectively utilizes the complementary pseudo-label signals from multiple VLMs while leveraging the discrepancy indicator to mitigate the impact of noisy labels, thus enhancing the reliability and scalability of VLM-based supervision for multi-class image classification.

Risk-Consistent Estimator

Based on the proposed HCL setup, we present a risk-consistent learning method (Feng et al. 2020b,a; Xu et al. 2022). To rigorously connect the ground-truth label and HCL labels, we introduce the following definition.

Definition 1. (HCL Condition.) Given a sample $(x, Y, s) \in \mathcal{D}_{HCL}$, $\forall i, j \in \mathcal{Y}$, we define the conditional probability of HCLs to satisfy:

$$P(y = i, Y = j \neq i, s = 1 | x) = 0. \quad (3)$$

This definition implies that for samples with VLM discrepancies ($\hat{y}_i \neq \tilde{y}_i$, i.e., $s_j = 1$), the ground-truth label y must be different from the label predicted by at least one VLM (i.e., \hat{y} or \tilde{y}). For these instances, we rely on human correction to provide the true label y as a trustworthy supervision label. Although the true conditional distribution $P(y = i | x)$ in Eq. 1 is unknown, we approximate it under the HCL probabilities on consistent predictions (i.e., $s = 0$) as:

$$\{P(y = i, Y, s = 0 | x)\}_{i=1}^k, \quad m = 1, \dots, k. \quad (4)$$

Lemma 2. Under the HCL Definition 1, the conditional probabilities $P(y = i | x)$ can be expressed as:

$$P(y = i | x) = P(Y = i, s = 1 | x) + \sum_{m=1}^k \left[P(y = i | Y = m, s = 0, x) P(Y = m, s = 0 | x) \right]. \quad (5)$$

This decomposition expresses the total probability of class i given instance x as the sum of two components: the probability of human correction, and a weighted conditional probability by the marginal probability. This formulation enables the model to express the unknown true label distribution using a combination of corrected labels from inconsistent samples and soft VLM-generated labels from consistent samples. The proof is provided in the Appendix A.1.

Theorem 3. To address the HCL learning problem, according to the Definition 1 and Lemma 2, the classification risk $R(f)$ in Eq. 1 can be reformulated as:

$$R_{HCL}(f) = \mathbb{E}_{(x, Y, s=1)} \mathcal{L}[f(x), i] + \mathbb{E}_{(x, Y, s=0)} \sum_{i=1}^k P(y = i | Y, s = 0, x) \mathcal{L}[f(x), i], \quad (6)$$

where $P(y = i | Y, s = 0, x)$ denotes the conditional probability derived from consistent VLM predictions, enabling the classifier f to learn from a conditional label distribution. In contrast, for samples with discrepancies (i.e., $s = 1$), the classifier is trained using human-corrected labels to ensure reliable supervision. The proof is provided in the Appendix A.2.

Practical Implementation

Loss Function. To effectively leverage VLM-generated labels while mitigating noise under weak supervision, we design a tailored loss function aligned with our VLM-consistency filtering setting. Specifically, we adopt a modified MSE-like loss, formulated explicitly to enhance discriminability and training stability. Let the classifier output be $f(x) \in \mathbb{R}^k$. We use the loss as:

$$\mathcal{L}(f(x), i) = \frac{1}{k} \sum_{j=1}^k [1 - f_j(x)(2\delta_{ij} - 1)]^2, \quad (7)$$

where $f_j(x)$ denotes the j -th classifier, and δ_{ij} is the Kronecker delta (i.e., $\delta_{ij} = 1$ if $i = j$, and $\delta_{ij} = 0$, otherwise). Specifically, the correct class (where $i = j$) is assigned 1 to emphasize its importance, while incorrect classes (where $i \neq j$) are assigned -1 to suppress their activation. The use of squared error penalizes both under-activation for the correct class and over-activation for incorrect classes, encouraging sharper class separation.

Empirical Approximation. We arrange the HCLs dataset as $\mathcal{D}_c = \mathcal{D}_H \cup \mathcal{D}_V$, where \mathcal{D}_H and \mathcal{D}_V denote the set of VLM-inconsistent data and VLM-consistent data respectively. Then, the empirical approximation of the classification risk in Eq. 6 is given by:

$$\hat{R}_{HCL}(f) = \frac{1}{|\mathcal{D}_H|} \sum_{x_j \in \mathcal{D}_H} \mathcal{L}(f(x_j), y_j) + \frac{1}{|\mathcal{D}_V|} \sum_{x_j \in \mathcal{D}_V} \sum_{i=1}^k \left[\hat{P}(y_j = i | Y_j, s_j = 0, x_j) \mathcal{L}(f(x_j), i) \right], \quad (8)$$

where $\hat{P}(y_j = i | Y_j, s_j = 0, x_j)$ is the estimated conditional label distribution for sample x_j .

This empirical approximation enables effective learning under weak supervision while controlling VLM-generated label noise, thereby supporting scalable training on large-scale unlabeled datasets with minimal human correction.

	CIFAR100	Tiny-ImageNet	Caltech-101	Food-101	EuroSAT	DTD
Supervised Labels Learning						
FSL-C (Radford et al. 2021)	83.12	81.13	95.77	93.14	95.96	76.02
Weakly Supervised Learning Methods						
Zero-shot CLIP (Radford et al. 2021)	73.90	72.29	88.45	88.21	47.35	51.98
HL-C (Radford et al. 2021)	74.04	72.50	57.54	73.90	95.48	60.90
VL-C (Radford et al. 2021)	66.25	-	42.02	73.66	-	45.54
Only-C (Radford et al. 2021)	77.11	-	50.44	92.48	-	52.22
Only-Q (Wang et al. 2024)	49.77	-	35.53	54.90	11.64	55.35
DIRK (Wu, Wang, and Zhang 2024)	77.26	52.91	63.66	90.12	44.67	67.46
PaPi (Xia et al. 2023)	79.67	25.33	57.66	90.73	29.01	63.62
HCL (Ours)	82.48	79.87	94.68	93.07	95.53	70.76

Table 2: Accuracy (%) of all compared methods on six datasets. The highest value for each dataset is highlighted in bold. Results below 10% are omitted.

Conditional Probability Estimation. Minimizing the empirical risk \hat{R}_{HCL} , as introduced in Theorem 3, requires estimating the conditional probability distribution $P(y = i | Y, s = 0, x)$ for VLM-consistent samples. To leverage both the prior knowledge from VLMs and the model predictions, we estimate this distribution via a weighted distribution of two sources. The model-based conditional distribution P_{model} is obtained by applying the softmax function to the classifier output logits $f(x) \in \mathbb{R}^k$:

$$P_{model}(y = i | x) = \frac{e^{f_i(x)}}{\sum_{l=1}^k e^{f_l(x)}}. \quad (9)$$

The external distribution P_{CLIP} is derived from the CLIP model via image-text alignment. Specifically, for an image x , we compute its image embedding $g_I(x)$ and compare it with a set of pre-encoded text embeddings $Q_T = [q_1, \dots, q_k]$, where each q_i corresponds to class i . The resulting probability distribution is computed using cosine similarity followed by temperature-scaled softmax:

$$P_{CLIP}(y = i | x) = \frac{e^{\tau \cdot \cos(g_I(x), q_i)}}{\sum_{l=1}^k e^{\tau \cdot \cos(g_I(x), q_l)}}. \quad (10)$$

where $\cos(\cdot, \cdot)$ denotes cosine similarity and $\tau > 0$ is a temperature parameter used to adjust confidence sharpness. Empirically, we set $\tau = 100$ to produce confident yet stable predictions.

To form the final estimated conditional distribution $\hat{P}(y = i | Y, s = 0, x)$, we compute a convex combination of P_{CLIP} and P_{model} , weighted by a mixing coefficient $\lambda \in [0, 1]$:

$$\hat{P}(y = i | Y, s = 0, x) = \lambda P_{CLIP} + (1 - \lambda) P_{model}. \quad (11)$$

This interpolation enables dynamic balancing between external VLM-generated labels and the evolving predictions of the model, which is particularly beneficial under weak supervision.

Model. We employ the CLIP model with ViT-L/14 (Radford et al. 2021) as the vision backbone for feature extraction, followed by a trainable linear classification head.

Specifically, we utilize the open-sourced CLIP implementation in FP32 precision, where the image encoder is frozen to act as a feature extractor during the entire training process. The extracted visual features are of dimension 768, which are then fed into a learnable linear layer for classification. This design allows us to leverage the strong representation capabilities of pre-trained VLMs while maintaining computational efficiency.

Experiments

Experimental Setup

Datasets. To comprehensively evaluate the effectiveness of our proposed method, we conduct experiments on six diverse multi-class image classification datasets, spanning both coarse-grained (i.e., CIFAR-100 (Krizhevsky, Hinton et al. 2009), Tiny-ImageNet (Le and Yang 2015), and Caltech-101 (Fei-Fei, Fergus, and Perona 2004)) and fine-grained (i.e., Food-101 (Bossard, Guillaumin, and Van Gool 2014), EuroSAT (Helber et al. 2019), and DTD (Cimpoi et al. 2014)) classification tasks across different domains. All images are resized to 224×224 before being fed into the models. Detailed descriptions of the datasets are provided in the Appendix A.3.

Compared Methods. We evaluate the proposed HCL method against representative baselines, including state-of-the-art weakly supervised methods and VLM-based linear probing under varying supervision settings:

- DIRK (Wu, Wang, and Zhang 2024) and PaPi (Xia et al. 2023): Partial-label learning methods using VLM predictions to construct candidate label sets. For consistent samples (i.e., $s_j = 0$), only the CLIP label is used; for inconsistent ones (i.e., $s_j = 1$), the set includes predictions from CLIP, Qwen, and the human-corrected label.
- FSL-C (Radford et al. 2021): Fully supervised method trained on all labeled training data using ground-truth.
- HL-C (Radford et al. 2021): Training only on samples with VLM discrepancies (i.e., $s_j = 1$), using human-corrected labels; consistent samples are discarded.

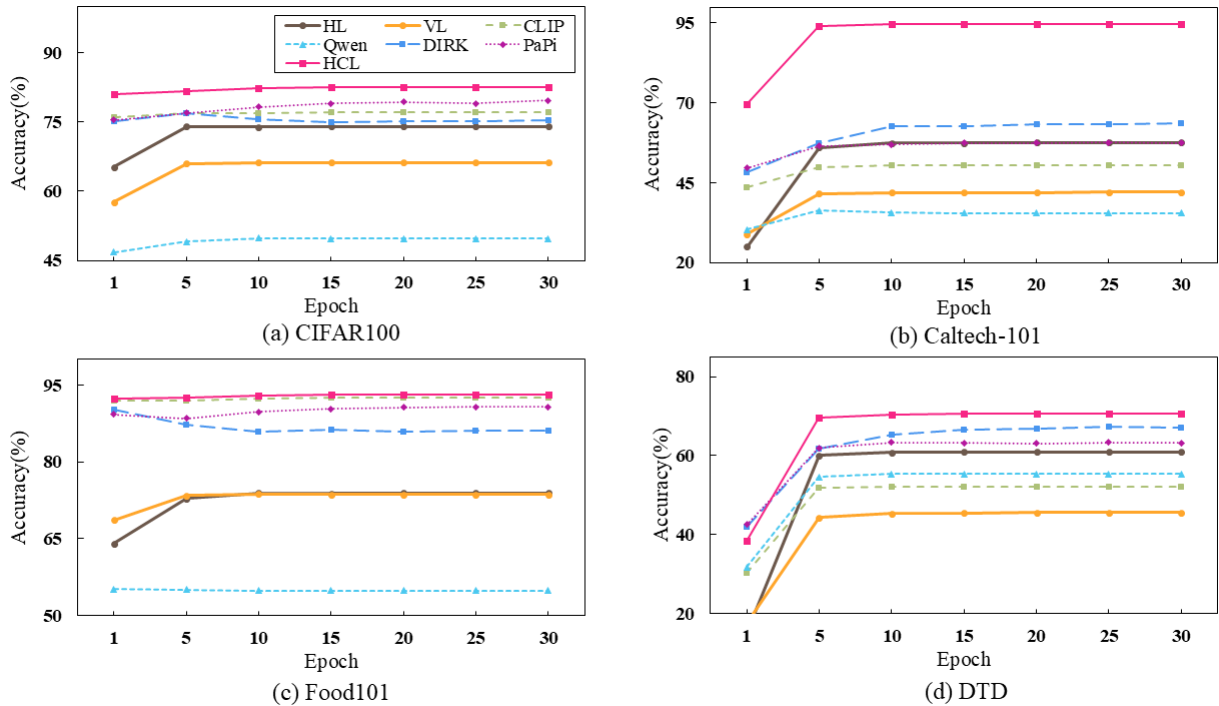


Figure 2: Classification accuracy over training epochs (1–30) for each method on (a) CIFAR100, (b) Caltech-101, (c) Food-101, and (d) DTD.

- VL-C (Radford et al. 2021): Weak supervision method using VLM predictions on consistent samples only (i.e., $s_j = 0$).
- Only-C (Radford et al. 2021) and Only-Q (Wang et al. 2024): Training solely with CLIP or Qwen predictions, without considering consistency.

These baselines enable a thorough comparison under unified settings, highlighting the effectiveness of our hybrid supervision approach guided by VLM discrepancies.

Implementation Details. All experiments are implemented in PyTorch and conducted on a single NVIDIA RTX 4090 GPU. We fix the random seed to 42 to ensure reproducibility. All models are trained for 30 epochs with a batch size of 64 using the AdamW optimizer (Loshchilov and Hutter 2019) with a learning rate of 5×10^{-4} and a weight decay of 1×10^{-4} . To ensure stable convergence, we adopt a step-wise learning rate scheduler that decays the learning rate by a factor of 0.1 every 5 epochs. This lightweight and consistent training setup ensures scalability and efficiency across all datasets and methods. Unless otherwise specified, the hyperparameter λ is typically set to 1.0 by default.

Comparison with State-of-the-Art

To evaluate the effectiveness of the proposed HCL method, we conduct experiments on six benchmark datasets under weakly supervised methods. Label distributions are generated by CLIP with ViT-L/14 and used within HCL to train a linear classifier with minimal human correction. All ex-

periments are trained with identical configurations for fair comparison, and results are summarized in Table 2.

For clarity, results with accuracy below 10% are omitted. In these cases, methods relying solely on VLM-generated labels (i.e., VL-C, Only-C, Only-Q) perform poorly on challenging datasets such as Tiny-ImageNet and EuroSAT. This indicates that raw VLM predictions, without human correction or filtering, are unreliable under domain shifts.

Across all datasets, HCL consistently outperforms existing weakly supervised baselines. On challenging datasets like Tiny-ImageNet and EuroSAT, it achieves substantial improvements over DIRK and PaPi, illustrating the effectiveness of combining VLM priors with targeted human correction. Moreover, HCL approaches the performance of fully supervised models, narrowing the gap between weak and full supervision while greatly reducing annotation cost.

In summary, HCL leverages VLM priors and selectively applies human correction, providing reliable supervision at low cost. By using label distributions as richer training signals, HCL captures finer category boundaries, making it a scalable and practical approach for real-world weakly supervised image classification.

Comparison of Training Cost

To assess the efficiency of HCL, we compare training convergence across four datasets: CIFAR100, Caltech-101, Food-101, and DTD. Figure 2 shows accuracy versus training epochs for HCL, CLIP linear probing, Only-Q probing, DIRK, and PaPi.




Figure 3: Experimental results on the influence of the conditional probability hyperparameter λ . Experiments are performed on human corrected HCLs data.

HCL achieves faster convergence and higher final accuracy than all baselines. It stabilizes within 10 epochs, whereas methods such as DIRK and PaPi require more iterations and converge to lower performance. These results indicate that HCL enhances both training efficiency and accuracy, providing a computationally efficient solution for large-scale weakly supervised learning.

Influence of The Hyperparameter λ

We examine how performance varies with λ on Caltech-101, EuroSAT and DTD. Figure 3 shows that as λ increases, introducing conditional probability distribution does not improve HCL performance on Caltech-101 and EuroSAT; instead, accuracy declines. This is expected, as CLIP already provides highly reliable supervision here, and additional self-predictions may introduce noise. However, on DTD, distribution yields slight but consistent improvements over standard HCL. These results indicate that on domain-shifted or texture-centric datasets, where CLIP supervision may have domain bias or uncertainty, incorporating adaptive model predictions refines and corrects noisy or biased VLM-generated labels, enhancing supervision quality.

Impact of Various VLMs

To investigate how the choice and number of VLMs affect the consistency-based filtering in HCL, we conduct research by varying the VLM combinations used to assess prediction agreement. Specifically, we consider three configurations: (1) Case 1, the default setup using CLIP and Qwen, (2) Case 2, an alternative using CLIP and LLaVA (Liu et al. 2023), and (3) Case 3, an extended configuration incorporating CLIP, Qwen, and LLaVA. In two-VLM settings, we report the proportion of consistent predictions; in the three-VLM setting, we report the proportion of fully inconsistent cases (i.e., $s_j = 1$), where all models disagree. For $s_j = 0$, we adopt the majority-agreed label as supervision.

As shown in Table 3, accuracy remains stable across CIFAR100 and DTD, indicating that HCLs is generally robust

	CIFAR100	Caltech-101	DTD
Case 1	82.48/44.02	94.68/64.36	70.76/48.04
Case 2	82.42/46.68	93.33/38.39	69.46/33.19
Case 3	82.42/33.02	74.16/28.73	70.70/36.61

Table 3: HCLs classification accuracy (left) and consistent probability or full-model discrepancy rate (right) under different VLM configurations.

	CIFAR100/VLM	Caltech-101/VLM	DTD/VLM
Prompt A	82.48/45.82	94.68/65.38	70.76/56.65
Prompt B	82.48/44.69	94.83/43.00	71.06/57.66
Prompt C	82.40/43.57	94.24/38.91	71.35/60.86

Table 4: HCLs classification accuracy (left) and corresponding VLM-generated label accuracy (right) under different prompts across datasets.

to the choice of VLMs. However, on Caltech-101, performance declines under Case 3, likely due to the relatively low proportion of fully inconsistent samples, meaning many training samples rely on partially consistent predictions from two VLMs. If the VLMs themselves make incorrect predictions, this increases the likelihood of introducing noisy labels. This suggests that while HCLs supports flexible VLM integration, more models do not always yield better results.

Impact of Various Prompts

To evaluate the sensitivity of the HCLs setting to prompt design when generating VLM-generated labels, we study the use of different prompt variants. Specifically, we compare three types: (1) Prompt A and Prompt B, same content in different languages, and (2) Prompt C, a simplified English version with concise phrasing. Detailed prompt descriptions are shown in Appendix A.4.

As shown in Table 4, prompt choice affects pseudo-label accuracy, but HCLs maintains stable final performance across datasets. This robustness stems from its use of model discrepancies and human correction, allowing tolerance to prompt-induced noise. These results suggest HCLs reduces the need for extensive prompt tuning when using general-purpose VLMs.

Conclusion

In this paper, we investigate a novel weakly supervised learning setting, Human-Corrected Labels (HCLs), which strategically incorporates minimal human corrections to improve the quality of noisy VLM-generated labels. We formulate a risk-consistent learning method that leverages both VLM predictions and human correction by estimating conditional label distributions, enabling effective learning under label noise. Extensive experiments across standard and real-world datasets demonstrate that our method achieves high classification accuracy while substantially reducing annotation cost, highlighting the practicality and robustness of the proposed HCLs.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (No.62306320, 61976217), the Open Project Program of State Key Lab. for Novel Software Technology (No. KFKT2024B32), and the Natural Science Foundation of Jiangsu Province (No. BK20231063).

References

- Albert, P.; Arazo, E.; Krishna, T.; O'Connor, N. E.; and McGuinness, K. 2023. Is Your Noise Correction Noisy? PLS: Robustness to Label Noise with Two Stage Detection. In *IEEE Winter Conference on Applications of Computer Vision*, 118–127.
- Berthelot, D.; Carlini, N.; Goodfellow, I.; Papernot, N.; Oliver, A.; and Raffel, C. A. 2019. MixMatch: A Holistic Approach to Semi-supervised Learning. *Advances in Neural Information Processing Systems*, 32.
- Bossard, L.; Guillaumin, M.; and Van Gool, L. 2014. Food-101—Mining Discriminative Components with Random Forests. In *European Conference on Computer Vision*, 446–461.
- Cao, K.; Brbic, M.; and Leskovec, J. 2021. Open-World Semi-Supervised Learning. arXiv:2102.03526.
- Chapel, L.; Alaya, M. Z.; and Gasso, G. 2020. Partial Optimal Transport With Applications on Positive-Unlabeled Learning. *Advances in Neural Information Processing Systems*, 33: 2903–2913.
- Chen, H.; Tao, R.; Fan, Y.; Wang, Y.; Wang, J.; Schiele, B.; Xie, X.; Raj, B.; and Savvides, M. 2023. SoftMatch: Addressing the Quantity-Quality Trade-off in Semi-supervised Learning. arXiv:2301.10921.
- Chen, T.; Yao, Y.; and Tang, J. 2023. Multi-Granularity Denoising and Bidirectional Alignment for Weakly Supervised Semantic Segmentation. *IEEE Transactions on Image Processing*, 32: 2960–2971.
- Chen, T.; Yao, Y.; Zhang, L.; Wang, Q.; Xie, G.-S.; and Shen, F. 2022. Saliency Guided Inter-and Intra-Class Relation Constraints for Weakly Supervised Semantic Segmentation. *IEEE Transactions on Multimedia*, 25: 1727–1737.
- Cimpoi, M.; Maji, S.; Kokkinos, I.; Mohamed, S.; and Vedaldi, A. 2014. Describing Textures in the Wild. In *IEEE Conference on Computer Vision and Pattern Recognition*, 3606–3613.
- Fei-Fei, L.; Fergus, R.; and Perona, P. 2004. Learning Generative Visual Models from Few Training Examples: An Incremental Bayesian Approach Tested on 101 Object Categories. In *IEEE Conference on Computer Vision and Pattern Recognition*, 178–178.
- Feng, L.; and An, B. 2019. Partial Label Learning with Self-Guided Retraining. In *AAAI Conference on Artificial Intelligence*, volume 33, 3542–3549.
- Feng, L.; Kaneko, T.; Han, B.; Niu, G.; An, B.; and Sugiyama, M. 2020a. Learning with Multiple Complementary Labels. In *International Conference on Machine Learning*, 3072–3081.
- Feng, L.; Lv, J.; Han, B.; Xu, M.; Niu, G.; Geng, X.; An, B.; and Sugiyama, M. 2020b. Provably Consistent Partial-Label Learning. *Advances in Neural Information Processing Systems*, 33: 10948–10960.
- Feng, Z.; Zhou, Q.; Gu, Q.; Tan, X.; Cheng, G.; Lu, X.; Shi, J.; and Ma, L. 2022. DMT: Dynamic Mutual Training for Semi-Supervised Learning. *Pattern Recognition*, 130: 108777.
- Ghosh, A.; Kumar, H.; and Sastry, P. S. 2017. Robust loss functions under label noise for deep neural networks. In *AAAI Conference on Artificial Intelligence*, volume 31.
- Gong, C.; Ding, Y.; Han, B.; Niu, G.; Yang, J.; You, J.; Tao, D.; and Sugiyama, M. 2022. Class-wise Denoising for Robust Learning Under Label Noise. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3): 2835–2848.
- Grandvalet, Y.; and Bengio, Y. 2004. Semi-supervised Learning by Entropy Minimization. *Advances in Neural Information Processing Systems*, 17.
- Han, B.; Niu, G.; Yu, X.; Yao, Q.; Xu, M.; Tsang, I.; and Sugiyama, M. 2020. Sigua: Forgetting may make learning with noisy labels more robust. In *International Conference on Machine Learning*, 4006–4016.
- Helber, P.; Bischke, B.; Dengel, A.; and Borth, D. 2019. EuroSAT: A Novel Dataset and Deep Learning Benchmark for Land Use and Land Cover Classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7): 2217–2226.
- Huang, Q.; He, Y.; and Huang, Z. 2021. A Novel Maximum Mean Discrepancy-Based Semi-Supervised Learning Algorithm. *Mathematics*, 10(1): 39.
- Kabra, R.; Matthey, L.; Lerchner, A.; and Mitra, N. J. 2024. Leveraging VLM-Based Pipelines to Annotate 3D Objects. In *International Conference on Machine Learning*.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning Multiple Layers of Features from Tiny Images. Technical report, University of Toronto.
- Le, Y.; and Yang, X. 2015. Tiny ImageNet Visual Recognition Challenge. *CS231N Course Project Reports*, 7(7): 3.
- Lee, D.-H.; et al. 2013. Pseudo-Label: The Simple and Efficient Semi-supervised Learning Method for Deep Neural Networks. In *International Conference on Machine Learning*, volume 3, 896.
- Li, J.; Li, D.; Xiong, C.; and Hoi, S. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, 12888–12900.
- Li, J.; Xiong, C.; and Hoi, S. C. 2021. Learning from Noisy Data with Robust Representation Learning. In *IEEE International Conference on Computer Vision*, 9485–9494.
- Li, Z.; Wei, M.; Ying, P.; Sun, T.; and Xu, X. 2024. Learning from Concealed Labels. In *ACM International Conference on Multimedia*, 7220–7228.
- Li, Z.; Xu, J.; Ying, P.; Wei, M.; and Xu, X. 2025. Learning from True-False Labels via Multi-modal Prompt Retrieving. In *International Conference on Machine Learning*.

- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023. Visual instruction tuning. *Advances in Neural Information Processing Systems*, 36: 34892–34916.
- Liu, S.; Niles-Weed, J.; Razavian, N.; and Fernandez-Granda, C. 2020. Early-Learning Regularization Prevents Memorization of Noisy Labels. *Advances in Neural Information Processing Systems*, 33: 20331–20342.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*.
- Ma, C.; Pan, X.; Ye, Q.; Tang, F.; Dong, W.; and Xu, C. 2023. CrossRectify: Leveraging Disagreement for Semi-Supervised Object Detection. *Pattern Recognition*, 137: 109280.
- Menghini, C.; Delworth, A.; and Bach, S. 2023. Enhancing clip with clip: Exploring pseudolabeling for limited-label prompt tuning. *Advances in Neural Information Processing Systems*, 36: 60984–61007.
- Menon, A.; Van Rooyen, B.; Ong, C. S.; and Williamson, B. 2015. Learning from corrupted binary labels via class-probability estimation. In *International Conference on Machine Learning*, 125–134.
- Miyato, T.; Maeda, S.-i.; Koyama, M.; and Ishii, S. 2018. Virtual Adversarial Training: A Regularization Method for Supervised and Semi-supervised Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(8): 1979–1993.
- Patel, D.; and Sastry, P. 2023. Adaptive Sample Selection for Robust Learning Under Label Noise. In *IEEE Winter Conference on Applications of Computer Vision*, 3932–3942.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, 8748–8763.
- Sajjadi, M.; Javanmardi, M.; and Tasdizen, T. 2016. Regularization with Stochastic Transformations and Perturbations for Deep Semi-supervised Learning. *Advances in Neural Information Processing Systems*, 29.
- Sohn, K.; Berthelot, D.; Carlini, N.; Zhang, Z.; Zhang, H.; Raffel, C. A.; Cubuk, E. D.; Kurakin, A.; and Li, C.-L. 2020. FixMatch: Simplifying Semi-Supervised Learning with Consistency and Confidence. *Advances in Neural Information Processing Systems*, 33: 596–608.
- Sosea, T.; and Caragea, C. 2023. MarginMatch: Improving Semi-Supervised Learning with Pseudo-Margins. In *IEEE Conference on Computer Vision and Pattern Recognition*, 15773–15782.
- Sun, Z.; Liu, H.; Wang, Q.; Zhou, T.; Wu, Q.; and Tang, Z. 2021. Co-LDL: A Co-Training-Based Label Distribution Learning Method for Tackling Label Noise. *IEEE Transactions on Multimedia*, 24: 1093–1104.
- Tarvainen, A.; and Valpola, H. 2017. Mean Teachers are Better Role Models: Weight-averaged Consistency Targets Improve Semi-supervised Deep Learning Results. *Advances in Neural Information Processing Systems*, 30.
- Van Engelen, J. E.; and Hoos, H. H. 2020. A Survey on Semi-Supervised Learning. *Machine Learning*, 109(2): 373–440.
- Wang, D.-B.; Li, L.; and Zhang, M.-L. 2019. Adaptive Graph Guided Disambiguation for Partial Label Learning. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 83–91.
- Wang, P.; Bai, S.; Tan, S.; Wang, S.; Fan, Z.; Bai, J.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Fan, Y.; Dang, K.; Du, M.; Ren, X.; Men, R.; Liu, D.; Zhou, C.; Zhou, J.; and Lin, J. 2024. Qwen2-VL: Enhancing Vision-Language Model’s Perception of the World at Any Resolution. arXiv:2409.12191.
- Wang, W.; and Zhou, Z.-H. 2017. Theoretical Foundation of Co-Training and Disagreement-Based Algorithms. arXiv:1708.04403.
- Wang, X.; Wu, Z.; Lian, L.; and Yu, S. X. 2022. Debaised learning from naturally imbalanced pseudo-labels. In *IEEE Conference on Computer Vision and Pattern Recognition*, 14647–14657.
- Wang, Y.; Chen, H.; Heng, Q.; Hou, W.; Fan, Y.; Wu, Z.; Wang, J.; Savvides, M.; Shinozaki, T.; Raj, B.; Schiele, B.; and Xie, X. 2023. FreeMatch: Self-adaptive Thresholding for Semi-supervised Learning. In *International Conference on Learning Representations*.
- Wei, M.; Li, Z.; Zhou, Y.; and Xu, X. 2024. Learning From Reduced Labels for Long-Tailed Data. In *International Conference on Multimedia Retrieval*, 111–119.
- Wu, D.-D.; Wang, D.-B.; and Zhang, M.-L. 2024. Distilling Reliable Knowledge for Instance-Dependent Partial Label Learning. In *AAAI Conference on Artificial Intelligence*, volume 38, 15888–15896.
- Xia, S.; Lv, J.; Xu, N.; Niu, G.; and Geng, X. 2023. Towards Effective Visual Representations for Partial-Label Learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 15589–15598.
- Xia, X.; Han, B.; Wang, N.; Deng, J.; Li, J.; Mao, Y.; and Liu, T. 2022. Extended \$T\$T: Learning with Mixed Closed-Set and Open-Set Noisy Labels. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3): 3047–3058.
- Xu, N.; Qiao, C.; Lv, J.; Geng, X.; and Zhang, M.-L. 2022. One Positive Label is Sufficient: Single-Positive Multi-Label Learning with Label Enhancement. *Advances in Neural Information Processing Systems*, 35: 21765–21776.
- Yan, Y.; and Guo, Y. 2023. Mutual Partial Label Learning With Competitive Label Noise. In *International Conference on Learning Representations*.
- Yu, X.; Han, B.; Yao, J.; Niu, G.; Tsang, I.; and Sugiyama, M. 2019. How Does Disagreement Help Generalization Against Label Corruption? In *International Conference on Machine Learning*, 7164–7173.
- Zhang, M.-L.; and Yu, F. 2015. Solving the Partial Label Learning Problem: An Instance-Based Approach. In *International Joint Conference on Artificial Intelligence*, 4048–4054.