

Instance-wise Supervision-level Optimization in Active Learning

Shinnosuke Matsuo¹, Riku Togashi², Ryoma Bise¹, Seiichi Uchida¹, Masahiro Nomura²

¹Kyushu University, Japan ²CyberAgent, Inc., Japan

Abstract

Active learning (AL) is a label-efficient machine learning paradigm that focuses on selectively annotating high-value instances to maximize learning efficiency. Its effectiveness can be further enhanced by incorporating weak supervision, which uses rough yet cost-effective annotations instead of exact (i.e., full) but expensive annotations. We introduce a novel AL framework, Instance-wise Supervision-Level Optimization (ISO), which not only selects the instances to annotate but also determines their optimal annotation level within a fixed annotation budget. Its optimization criterion leverages the value-to-cost ratio (VCR) of each instance while ensuring diversity among the selected instances. In classification experiments, ISO consistently outperforms traditional AL methods and surpasses a state-of-the-art AL approach that combines full and weak supervision, achieving higher accuracy at a lower overall cost. This code is available at <https://github.com/matsuo-shinnosuke/ISOAL>.

1. Introduction

Machine learning technology has achieved remarkable success in recent years across various fields, including computer vision, natural language processing, and robotics. However, a major challenge remains: creating labeled datasets for training models often requires significant time and cost, especially in specialized domains.

To address this issue, active learning [1, 3, 6, 8, 11, 16, 18–21, 24] has been widely studied to improve model performance under limited budgets. In active learning, instances (data samples) for annotation are selected iteratively from an unlabeled dataset, followed by annotation and model retraining. This repeated cycle improves the model’s accuracy while progressively minimizing annotation costs. A common approach in active learning is uncertainty-based sampling [3, 5, 8, 9, 18, 22, 24, 25], where instances with high uncertainty, such as low confident of the model’s prediction, are selected for annotation, as these are expected to contribute most effectively to model performance.

Another approach to reduce annotation costs is *weak su-*

pervision [17, 26, 27]. For classification tasks, a weak supervision approach is often designed to just attach rough class labels to individual instances. For example, instead of attaching the exact class label “song sparrow” or “house sparrow” to an image instance, annotators can attach a rough class label “sparrows.” This approach drastically reduces annotation costs (budgets) because the fee paid to a non-expert annotator, who can only assign rough class labels rather than exact ones via a crowdsourcing service, is lower than that of an expert annotator. Hereafter, we call rough classes as superclasses; therefore, “sparrows” is a superclass of the exact class “song sparrow.”

In general, weak supervision approaches still need instances with *full supervision*, that is, instances with the exact class labels. This is simply because it is impossible to determine the class boundary between, “sparrows” is a superclass of “song sparrow” without the instances with these exact class labels. Consequently, we need to balance weak supervision and full supervision under a pre-specified total annotation budget. In other words, we need to optimize the *supervision levels* of individual instances. Since the total annotation budget is pre-specified as a constraint, the number of instances with full supervision will be limited. More formally, this supervision-level optimization task can be seen as an optimal resource allocation problem with the constraint on available resources.

This paper proposes a label-efficient machine learning approach, called Instance-wise Supervision-level Optimization (ISO) for active learning, as shown in Fig. 1. Roughly speaking, the proposed approach is a sophisticated hybrid of weak supervision and active learning. Different from the state-of-the-art hybrid approach [22] (reviewed in Section 2), our approach has a novel mechanism that automatically determines the optimal supervision levels of individual instances. Specifically, as shown in Fig. 1, all unlabeled instances at a certain round are automatically divided into three categories: instances for full supervision (with exact class annotation), instances for weak supervision (with superclass annotation), and instances that remain unlabeled. The instances selected in the first two categories are considered to be selected for active learning; annotators, therefore, attach appropriate labels to these instances according

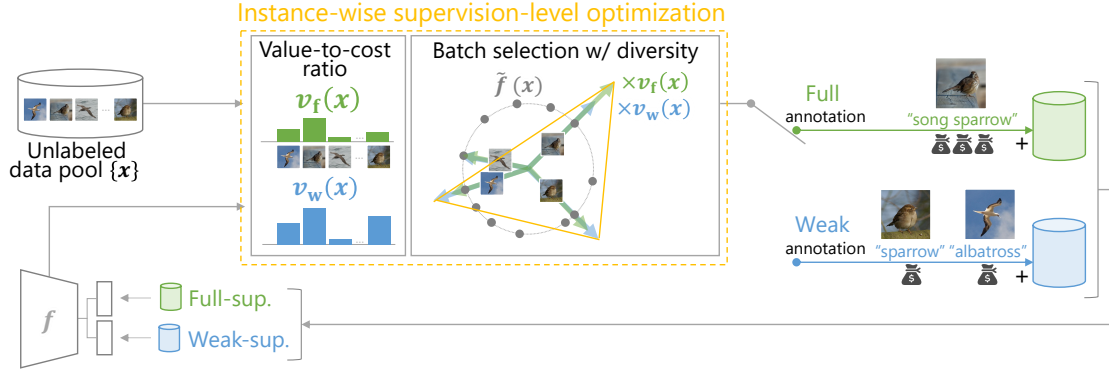


Figure 1. Overview of Instance-wise Supervision-level Optimization (ISO) for active learning. The proposed approach combines weak and full supervision within an active learning framework to maximize annotation efficiency under a fixed budget. First, each instance x in the unlabeled data pool is evaluated for its value-to-cost ratio ($v_f(x)$ for full supervision and $v_w(x)$ for weak supervision). Based on these values and data diversity, instances are selected within the budget constraints for weak supervision (with superclass labels) or full supervision (with exact class labels). This adaptive allocation of supervision-level enables optimal use of resources within the budget constraints in active learning.

to their supervision levels.

In our experiments, the proposed method consistently outperformed existing approaches across multiple datasets, demonstrating a more efficient use of the annotation budget. Our approach achieved the same accuracy as the comparative methods while using only three-fifths of the budget by incorporating weak labels when the cost of weak label annotation is half that of fully supervised annotation. Furthermore, in comparison with the state-of-the-art method that combines active learning and multiple supervision, our approach achieved significantly better cost efficiency by leveraging instance-specific information.

The main contributions of this paper are summarized as follows:

- We introduce a novel active learning framework that optimizes the supervision level for each instance, allowing for dynamic allocation of weak and full annotations based on budget constraints.
- We propose an efficient algorithm that assigns the optimal supervision level to each instance by considering value-to-cost ratios and data diversity, enhancing the cost-effectiveness of the annotation process.
- We validate the effectiveness of our approach on classification tasks, using superclasses as weak supervision, and demonstrate its superiority over conventional AL methods and the state-of-the-art AL method that incorporates both full and weak supervision.

2. Related work

Active learning (AL) [1, 16, 20] is a technique aimed at maximizing model performance under a limited annotation budget by strategically selecting informative instances for annotation. A widely adopted approach in active learning is

uncertainty-based sampling [3, 5, 8, 9, 18, 22, 24, 25]. For example, [18] selects instances with small margins, prioritizing instances for which the model has low confidence. Another common approach focuses on maximizing diversity within selected instances [6, 15, 19, 21], which aims to cover a broad range of the data distribution. For instance, [19] selects a representative core set of instances to maximize diversity. Moreover, there are methods that combine uncertainty and diversity-based strategies to leverage the advantages of both approaches [3, 11]. Additionally, some AL methods [4, 12] address the setting where different annotators have varying accuracy levels. These methods aim to jointly optimize both the selection of data for annotation and the choice of annotators to minimize label noise. However, all of these approaches assume a single supervision-level and do not incorporate multiple supervision-levels.

Another approach to reduce annotation costs is weakly supervised learning (WSL) [17, 26, 27], where less detailed or approximate labels instead of fully supervised labels are utilized. In many real-world applications, obtaining fully supervised labels can be expensive and time-consuming. To address this, WSL methods use weak supervision, which is typically easier and cheaper to obtain. For example, in classification tasks, weak supervision can be provided by superclasses such as categories, and in segmentation tasks, it can be provided by bounding boxes, allowing faster labeling with lower costs.

A hybrid approach between AL and WSL, where the proportion of full and weak supervision is adaptively determined, has been proposed in [22]. Although this approach aims to optimize budget use by setting an ideal balance of weak and full supervision, it applies random sampling after setting the proportion *without* instance-level considera-

Table 1. Comparison of conventional active learning methods. “Uncertainty” (Unc.) refers to whether a method incorporates uncertainty-based instance selection, such as selecting instances with minimal margins. “Diversity” (Div.) indicates if the method considers data diversity to maximize the representativeness of selected instances. “Multiple supervision” (Mul.) refers to the capability to utilize multiple levels of supervision, such as both weak and full supervision. “Instance-wise” (Ins.) represents the ability to choose supervision for each instance adaptively.

	Margin, etc. [8, 18, 24]	Coreset, etc. [6, 19, 21]	BADGE, etc. [3, 11]	APFWA [22]	Ours
Unc.	✓	×	✓	×	✓
Div.	×	✓	✓	×	✓
Mul.	×	×	×	✓	✓
Ins.	✓	✓	✓	×	✓

tions. To our knowledge, no method exists that dynamically assigns full and weak supervision for each instance.

Table 1 summarizes a comparison between popular active learning methods and our instance-wise supervision-level optimization. Each method is evaluated on its ability to incorporate uncertainty, diversity, multiple levels of supervision, and instance-wise adaptability, as defined in the caption. Most uncertainty-based and diversity-based AL methods do not consider multiple levels of supervision, focusing instead on selecting instances based on a single type of annotation. Although APFWA incorporates multiple levels of supervision, it lacks instance-wise optimization. In contrast, our approach is unique in that it optimizes the supervision level on an instance-wise basis, simultaneously considering data uncertainty, diversity, and the annotation cost associated with multiple levels of supervision.

3. Problem setting

Let \mathcal{D}_u represent the unlabeled data pool, and let T denote the number of rounds in the active learning process, where B is the available budget in each round. In this setting, a set of instances is selected to be annotated with either full supervision or weak supervision, depending on the required level of detail and the associated cost. Specifically, we define C_f as the annotation cost per instance for full supervision, where precise labels are assigned, and C_w as the annotation cost per instance for weak supervision, which involves less precise, lower-cost labeling. These annotation costs, C_f and C_w , are provided in advance.

The process of batch selection and training is repeated for T rounds, allowing the model to progressively improve by iteratively selecting and annotating the most informative instances within each budget-constrained round. Once an instance is annotated, it is moved from the unlabeled pool \mathcal{D}_u to either the fully supervised data pool \mathcal{D}_f or the weakly supervised data pool \mathcal{D}_w , depending on the supervi-

Algorithm 1 ISO: Instance-wise Supervision-level Optimization

```

1: Inputs: Round  $T$ , budget of each round  $B$ , annotation
   cost  $C_f, C_w$ , unlabeled data pool  $\mathcal{D}_u$ , feature extractor
    $f$ , classification head  $h_f, h_w$ .
2: Outputs: trained  $f^T, h_f^T$ .
3: for round  $t = 1, \dots, T$  do
4:   if  $t$  is 1 then
5:     Sample randomly instances from  $\mathcal{D}_u$  within
     budget  $B$  to create an initial fully supervised
     dataset  $\mathcal{D}_f$  and weakly supervised dataset  $\mathcal{D}_w$ 
     with an equal number of instances.
6:   else if  $t > 1$  then
7:     Compute the model improvements,  $M_f^t, M_w^t$ ,
     for each supervision using  $\mathcal{D}_f$  and  $\mathcal{D}_w$  by Eq. (3)
     and (4).
8:     for  $x$  in  $\mathcal{D}_u$  do
9:       1. Compute the uncertainties,  $u_f^t(x), u_w^t(x)$ ,
        for each supervision.
10:      2. Compute the VCRs,  $v_f^t(x), v_w^t(x)$ , for
        each supervision by Eq. (1) and (2).
11:      3. Compute the normalized features  $\tilde{f}^t(x)$ .
12:    end for
13:    Select the batch  $\mathcal{D}_f^+, \mathcal{D}_w^+$  using the VCRs,
     $v_f^t(x), v_w^t(x)$ , and the normalized features
     $\tilde{f}^t(x)$  by Sec. 4.3.
14:    end if
15:     $\mathcal{D}_f \leftarrow \mathcal{D}_f \cup \mathcal{D}_f^+, \mathcal{D}_u \leftarrow \mathcal{D}_u \setminus \mathcal{D}_f^+$ 
16:     $\mathcal{D}_w \leftarrow \mathcal{D}_w \cup \mathcal{D}_w^+, \mathcal{D}_u \leftarrow \mathcal{D}_u \setminus \mathcal{D}_w^+$ 
17:    Train  $f^t, h_f^t$ , and  $h_w^t$  using  $\mathcal{D}_f$  and  $\mathcal{D}_w$ .
18: end for

```

sion level chosen. These annotated instances are then used to train a classifier.

The ultimate goal in this setting is to obtain the classifier’s parameters to classify the exact class of input data. To train the classification network, both fully supervised data and weakly supervised data contribute to training the feature extractor in the network. Therefore, balancing fully and weakly supervised data in a cost-effective manner is crucial for optimizing the learning process within budget constraints.

4. Proposed method: Instance-wise Supervision-level Optimization (ISO)

4.1. Overview

We propose an efficient algorithm to address the instance-wise supervision-level optimization problem within a fixed budget. Our approach comprises two primary components, as shown in Fig. 1. The first component calculates the value-to-cost ratio (VCR) for each instance-supervision

level pair, which indicates the expected improvement in model accuracy per unit of cost achieved through annotation for each instance. The VCR for each instance is determined by both the value of the supervision level and its uncertainty. This allows us to evaluate the cost-effectiveness of assigning different supervision levels to each instance individually. The second component performs batch selection within the budget B , optimizing the set of instances based on their VCRs and diversity. This two-part approach enables our approach to maximize the value of annotations within a limited budget by balancing cost-efficiency and diversity.

4.2. Value-to-cost ratio (VCR) estimation

In this section, we explain the calculation of the value-to-cost ratio (VCR) for each instance-supervision pair. An overview of the algorithm is provided in Algorithm 1. For an instance x in the unlabeled pool \mathcal{D}_u , the VCRs $v_f^t(x)$ for full supervision and $v_w^t(x)$ for weak supervision at round t are determined by the uncertainty of the instance, the value, and cost for the two different-level labels, which are defined as:

$$v_f^t(x) = \frac{M_f^t u_f^t(x)}{C_f} \quad (1)$$

$$v_w^t(x) = \frac{M_w^t u_w^t(x)}{C_w} \quad (2)$$

where C_f, C_w represent the annotation costs, and M_f^t, M_w^t are the expected model improvement, where $*_f, *_w$ correspond to full supervision and weak supervision, respectively. $u_f^t(x), u_w^t(x)$ represents the uncertainty for each instance x in the unlabeled pool \mathcal{D}_u at round t , calculated based on model predictions. This formulation allows us to evaluate the cost-effectiveness of assigning either full or weak supervision to each instance, providing a basis for the batch selection strategy within a fixed budget.

The expected model improvements M_f^t for full supervision and M_w^t for weak supervision (i.e., the values for each supervision) are estimated by evaluating the improvement in model performance as the training data is incrementally increased. Consider the t -th annotation round, where fully supervised data \mathcal{D}_f^t and weakly supervised data \mathcal{D}_w^t are available at that time. To evaluate accuracy improvement, we divide the labeled data into K subsets and incrementally add training data. For fully supervised data, the weakly supervised data is kept fixed while subsets of the fully supervised data are added. The model performance $m_{f,k}$ represents the evaluation using the first k subsets of fully supervised training data, with $m_{f,1}$ being the performance when using just one subset. Similarly, we compute $m_{w,k}$ for the weakly supervised data, where the fully supervised data is fixed, and subsets of the weakly supervised data are incrementally added. Here, model performance is evaluated us-

ing a small validation set randomly sampled within the budget of a single round. Then, a weighted average is applied for the improvement at each step, $m_{k+1}^t - m_k^t$ to smooth the sequence of improvements. The reason for the weighted approach is to focus on the latest model behavior. To obtain the improvement per data point, each improvement value is normalized by dividing by $|\mathcal{D}|/K$, where $|\mathcal{D}|$ represents the number of instances. That is, M_f^t and M_w^t are expressed as follows:

$$M_f^t = \frac{\sum_{k=1}^{K-1} k \cdot (m_{f,k+1}^t - m_{f,k}^t)}{\sum_{k=1}^{K-1} k} \cdot \frac{K}{|\mathcal{D}_f|} \quad (3)$$

$$M_w^t = \frac{\sum_{k=1}^{K-1} k \cdot (m_{w,k+1}^t - m_{w,k}^t)}{\sum_{k=1}^{K-1} k} \cdot \frac{K}{|\mathcal{D}_w|} \quad (4)$$

The uncertainty $u_f^t(x)$ and $u_w^t(x)$, corresponding to full and weak supervision, can be flexibly defined based on the specific task and dataset. For example, in the classification task, the entropy of the class prediction probabilities is adopted. If different functions are used to calculate uncertainty for full and weak supervision, there is a risk that the scales of $u_f^t(x)$ and $u_w^t(x)$ may differ significantly. To address this, we normalize each uncertainty value by converting it to a percentile score. In this transformation, the instance with the highest uncertainty is assigned a score of 1, the instance with the lowest uncertainty receives a score of 0, and the instance ranked at the 10th percentile in terms of uncertainty is assigned a score of 0.9. We use percentiles rather than dividing by the maximum uncertainty value because extreme outliers could result in an overly large maximum, causing most values to be close to zero. This percentile-based normalization ensures that the distribution of uncertainty values remains well-scaled, preserving relative differences across instances without distortion due to outliers.

4.3. Valuable and diverse batch selection using VCR

In this section, we describe the batch selection using the value-to-cost ratio (VCR) defined in the previous section. The most straightforward approach for batch selection is to greedily select instances with the highest VCR values. However, as noted in previous studies (e.g. [3, 19]), this approach can lead to selecting similar instances that have high VCR values, resulting in a loss of data diversity within the batch.

To address this, we propose a batch selection algorithm to select a diverse and valuable batch. In other words, we select a group of instances with high VCR values and diverse features. Specifically, we represent each instance as a vector using its VCR and feature vector. The vector differs depending on the type of label: $v_f(x)\tilde{f}(x)$ for full labels

and $v_w(\mathbf{x})\tilde{f}(\mathbf{x})$ for weak labels. Here, $\tilde{f}(\mathbf{x})$ is a normalized feature vector (with a magnitude of 1) output by the feature extractor f . The optimal set of instances is selected based on the vectors.

We aim to select a set of instances that maximizes the area enclosed by the vertices of the selected vectors, as shown in the center of Figure 1. Here, the optimization parameter is which vectors (instances) to select and their label types. When the selected instances have high diversity and value, the area enclosed by the selected vectors becomes larger. We solve this optimization problem subject to budget constraints (i.e., the sum of the costs of the optimized supervision levels for instances is less than B). The optimization problem of selecting vectors and determining their supervision levels can be reformulated as a selection problem, where the goal is to choose the optimal set of vectors from all possible vectors, including both those with weak labels $\{v_w(\mathbf{x})\tilde{f}(\mathbf{x})|\mathbf{x} \in \mathcal{D}_u\}$ and those with full labels $\{v_f(\mathbf{x})\tilde{f}(\mathbf{x})|\mathbf{x} \in \mathcal{D}_u\}$. This allows the problem of choosing the type of label to be integrated into the vector selection process. This optimization problem can be solved by maximizing the determinant of a set of vectors subject to budget constraints.

To perform the batch selection that maximizes the determinant of a matrix, one approach is to sample from a k -Determinantal Point Process (k -DPP) [13]. This method selects a batch containing k vectors with a probability proportional to the determinant of their Gram matrix. However, this differs from the batch selection in our approach in that we choose a set of vectors within a budget B rather than choosing k vectors. Therefore, we propose a new selection method, which is inspired by the work [3] using the k -means++ seeding algorithm [2] as an approximation algorithm for k -DPP.

In our batch selection method, the vector of the batch is sampled sequentially with a probability proportional to the square of the distance to the vector closest to the one selected so far until the budget B is exhausted. Specifically, if the selected vector corresponds to full supervision (i.e., a vector in the set $\{v_f(\mathbf{x})\tilde{f}(\mathbf{x})|\mathbf{x} \in \mathcal{D}_u\}$, the cost C_f is deducted from the budget, and it is added to the fully supervised data. If it corresponds to weak supervision (i.e., a vector in the set $\{v_w(\mathbf{x})\tilde{f}(\mathbf{x})|\mathbf{x} \in \mathcal{D}_u\}$), the cost C_w is deducted from the budget, and it is added to the weakly supervised data. This process is repeated until the budget B is exhausted. By doing this, we can obtain a diverse and valuable batch.

5. Experiments

5.1. Dataset

To demonstrate the effectiveness of the proposed method, we utilized two datasets: CIFAR100 and Caltech-UCSD

Birds-200-2011 (CUB200). The task is a multi-class classification problem, where each class (full label) belongs to a superclass, and the superclass label is provided as a weak label.

CIFAR-100 is a dataset of natural images containing 100 classes that span a wide range of everyday objects and animals. Each class in CIFAR-100 is associated with a pre-defined superclass, grouping similar classes together under broader categories, resulting in 20 superclasses in total. The dataset consists of 60,000 images in total, with 50,000 images for training and 10,000 images for testing.

CUB200 [23] consists of 200 classes of bird species, with each class representing a specific bird. Following previous study [14], we defined superclasses based on the suffixes in the class names, resulting in 70 unique superclasses. The dataset contains 11,788 images, with 5,994 images designated for training and 5,794 images for testing.

5.2. Comparison methods

We compared our approach with eight conventional AL methods, ranging from standard single-supervision AL methods to the state-of-the-art AL method incorporating multiple supervision, as follows:

- **Random**: The simplest baseline that samples instances randomly from the unlabeled pool.
- **Margin [18]**: A method that selects instances in ascending order of the difference between the highest and second-highest predicted class probabilities, prioritizing instances with minimal margins.
- **MaxConf [24]**: An approach that selects instances in order of the lowest maximum predicted class probability, focusing on instances with the least confidence.
- **Entropy [24]**: A method that selects instances according to the highest entropy in the predicted class probabilities, targeting instances with the greatest uncertainty.
- **Coreset [19]**: A diversity-based sampling method that selects instances farthest from the labeled data using a greedy algorithm to enhance diversity.
- **ALBL (Active Learning by Learning) [9]**: A multi-armed bandit approach that selects between AL methods at each round. In this experiment, we utilized MaxConf and Coreset.
- **BADGE (Batch Active learning by Diverse Gradient Embeddings) [3]**: A method that considers both predictive uncertainty and diversity by selecting instances with diverse and high-magnitude gradients in the gradient space.
- **APFWA (Adaptive Proportion of Full and Weak Annotations) [22]**: A state-of-the-art AL method that incorporates both full and weak supervision, which dynamically determines the proportion of full and weak annotations within each batch.

Since AL with both full and weak label annotation is a novel task, there is only one existing method for this task. Instead, we compare our approach with standard single-supervision AL methods as well as the state-of-the-art AL method. This comparison highlights the effectiveness of AL using weak labels. Additionally, a comparison with APFWA demonstrates the advantages of our approach for AL with both fully and weakly supervised labels. Furthermore, in Sec. 6.3, we also compare our approach with a simple combination of weak supervision and conventional AL methods to ensure a fair evaluation.

These methods offer a range of approaches, from simple random sampling to advanced techniques that balance uncertainty and diversity, as well as methods leveraging weak supervision, enabling a comprehensive evaluation of our proposed approach.

5.3. Implementation details

The network consists of a shared feature extractor f and two classification heads, h_f for full supervision and h_w for weak supervision. The encoder of ResNet18 [7] is used for the feature extractor f , and a single linear layer is used for each classification head. For the network trained using both fully and weakly labeled data, we employ a two-stage training approach in which the feature extractor and each classification head are trained first with weakly supervised data and then with fully supervised data. Since both fully supervised and weakly supervised learning are classification tasks in this experiment, we used the cross-entropy loss. The learning rate was set to 0.001, and the optimizer was Adam [10].

The settings for active learning were as follows. The number of rounds T was set to 5. The budget B of each round was set to 1,000 for CIFAR100 and to 500 for CUB200 since CUB200 had approximately 5,000 in training data, which was less than that of CIFAR100. The validation data with full supervision were prepared from training data, 1,000 for CIFAR100 and 500 for CUB200, which are the same number and small amount as in budget B . The cost of full supervision C_f was set to 1, and the cost of weak supervision was set to $\frac{1}{2}$. In addition, in Sec. 6.2, we conducted the experiments when changing C_w to $\frac{1}{4}$ and $\frac{1}{8}$ to investigate the impact of the difference on the cost of weak supervision.

The number of partitions K for estimating the expected model improvements M_f^t and M_w^t was set to 5. As the performance metric of $m_{f,k}^t, m_{w,k}^t$, we used the classification accuracy. To enhance stability, we repeated the estimation process three times with different random seeds and calculated the average. For uncertainty measurement, we adopted the simplest approach, Margin [18], to evaluate instance uncertainty.

6. Experimental results

6.1. Comparison with conventional active learning methods

The comparison results with conventional AL methods on CIFAR100 and CUB200 are shown in Fig. 2(a), (b), respectively. The graphs show the classification accuracy [%] for each round t .

In Figure 2(a), the results on the CIFAR100 dataset demonstrate that our approach archives higher accuracy across all rounds compared to conventional methods, including random sampling, uncertainty-based AL method (Margin, MaxConf, Entropy), diversity-based AL method (Coreset), and other AL methods (ALBL, BADGE), which use only full supervision. This demonstrates that our approach effectively leverages weak supervision, achieving higher performance within the same budget constraints. For example, at round $t = 5$, our approach achieved over 10% higher accuracy compared to conventional AL methods that use only full supervision. Furthermore, while these methods require five rounds (with a budget of 5000) to reach an accuracy of approximately 30%, our approach achieved the same or even higher accuracy in just three rounds, using only three-fifths of the budget.

Our approach also surpasses APFWA, another AL approach that incorporates weak supervision by optimizing the proportions of weak and full supervision within the batch but does not focus on selecting instances. This result suggests that our approach more effectively selects cost-effective instances by dynamically considering instance-wise uncertainty and diversity, thereby maximizing the learning impact of weak supervision.

Similarly, Figure 2(b) presents the results on the CUB200 dataset, where our approach also outperforms other AL methods across all rounds. The consistent improvement observed across both datasets indicates that our approach can dynamically optimize annotation strategies according to the dataset characteristics, achieving high adaptability.

6.2. Analysis for weak supervision costs

To demonstrate that our approach works at any cost of weak supervision and that the cheaper the cost of weak supervision, the greater the effect, we conducted experiments when changing the cost of weak supervision. Figure 3 shows the classification accuracy [%] on CIFAR100 and CUB200 at different costs of weak supervision. In this experiment, the cost of full supervision C_f is fixed at 1, meaning that C_w represents the relative cost to full supervision. It is important to note that the cost C_w is defined before active learning begins and is not a hyperparameter of the proposed method. Imagine a scenario where the fee paid to a non-expert annotator, who can only assign rough class labels (weak la-

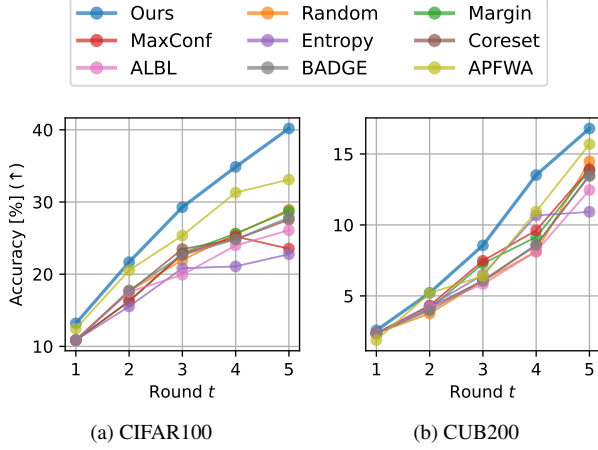


Figure 2. Comparison of our approach with the conventional active learning methods. The figures show the classification accuracy [%] (\uparrow) for each round t on CIFAR100 and CUB200. The cost of full supervision C_f is 1, and that of weak supervision is $\frac{1}{2}$.

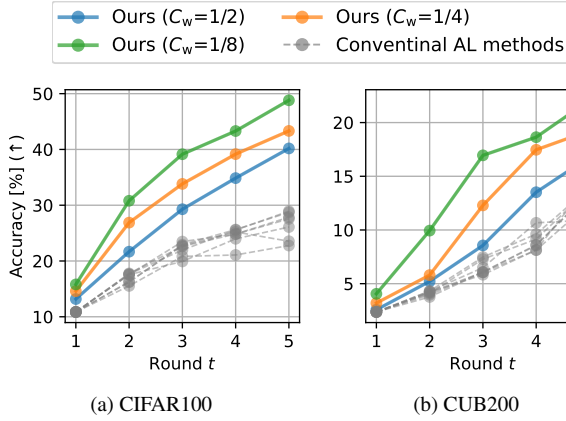


Figure 3. Classification accuracy [%] (\uparrow) of the proposed method on CIFAR100 and CUB200 when the cost of weak supervision, C_w is $\frac{1}{2}$, $\frac{1}{4}$, and $\frac{1}{8}$. Note that the cost C_w is a value set by the user before active learning begins and is not a hyperparameter of the proposed method. This figure illustrates the effect of different weak supervision costs. For reference, we also include the results of conventional active learning methods (gray lines).

bels), is lower than that of an expert annotator. The cost C_w varies for each task or dataset, corresponding to the rewards offered by crowdsourcing, etc.

In previous sections, we set the cost of weak supervision to $C_w = \frac{1}{2}$. Here, we investigate the effects of our approach when further reducing C_w to $\frac{1}{4}$ and $\frac{1}{8}$. As shown in Figure 3, for both datasets, the classification accuracy improves as the cost of weak supervision decreases. This improvement is due to the fact that the lower cost of weak supervision allows more data to be annotated within the same budget, thereby enhancing the training process. In other words, the cheaper the cost of weak supervision, the more impactful weak supervision and our approach become.

6.3. Comparison with combined active learning and multiple supervision baseline

In addition to comparing our approach with conventional active learning approaches, we conduct experiments to evaluate our approach against a straightforward combination of weak supervision and conventional active learning techniques. This baseline approach applies a fixed proportion of weak supervision within the budget, utilizing conventional active learning strategies within each supervision type.¹

Specifically, in this combined method, weak supervision is allocated at a constant rate—20%, 40%, 60%, and 80% in the batch. Within each full and weak supervision, conventional AL methods are employed to select data. In this experiment, the active learning method for each full or weak supervision employed was sampling based on margin, which is used for the uncertainty of the proposed method.

Figure 4 shows the comparison between our approach and baseline methods that combine active learning with weak supervision on the CIFAR100 dataset under different costs for weak supervision ($C_w = \frac{1}{2}, \frac{1}{4}$, and $\frac{1}{8}$). In each case, the baseline methods use fixed proportions of full and weak supervision within each batch, denoted as $F(x\%) + W(y\%)$, where x and y represent the percentages of full and weak supervision, respectively.

Across all weak supervision costs, our approach consistently achieves high accuracy. An important point to note is that the optimal proportion of full and weak supervision is unknown beforehand and varies depending on the cost of weak supervision. For example, when the cost of weak supervision C_w is $\frac{1}{2}$, the configuration with 60% full supervision and 40% weak supervision achieves the best accuracy among the baseline methods. However, when the cost C_w is reduced to $\frac{1}{8}$, a configuration with 20% full supervision and 80% weak supervision performs best. In contrast, our approach dynamically constructs optimal batches by optimizing both the supervision level and instance selection based on the cost-effectiveness obtained from the trained model at each round.

6.4. Ablation study

We conducted ablation studies to demonstrate the effectiveness of instance-wise value (i.e., uncertainty) and the batch selection with diversity on CIFAR100 and CUB200, as shown in Fig. 5.

The first ablation variant, referred to as “Ours w/o uncertainty,” removes the instance-wise uncertainty component from the VCR calculation. In this variant, the VCR is determined solely based on the annotation cost and expected model improvement without taking into account the

¹APFWA [22] is a method that dynamically determines the proportion of full and weak annotations. It must use random instance sampling to estimate the optimal proportion, and as a result, it cannot be combined with the conventional instance-wise AL methods (e.g., margin-based sampling).

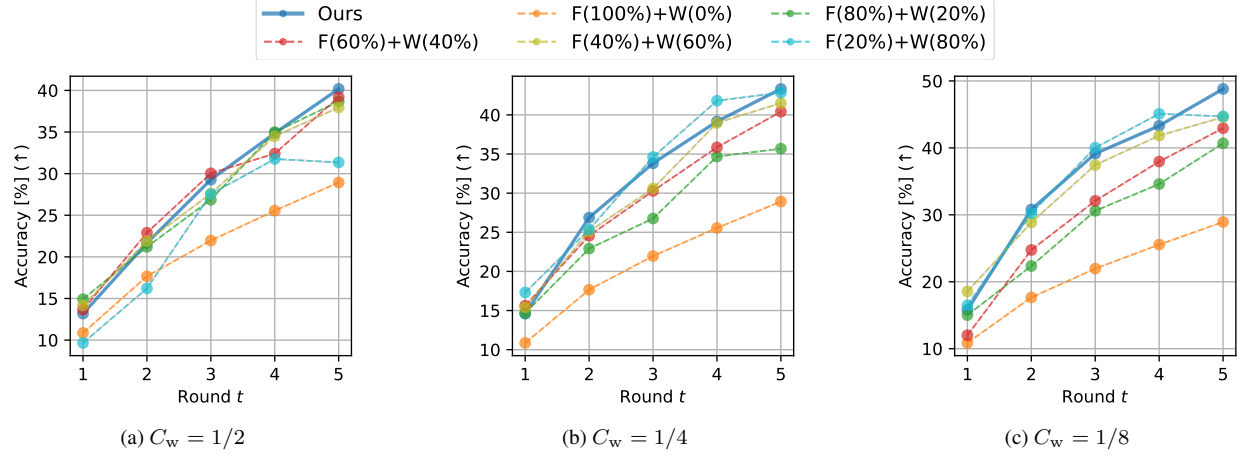


Figure 4. Comparison with combined active learning and multiple supervision baseline. The classification accuracy [%] (↑) on CIFAR100 with fixed proportions of full and weak supervision within each batch. Each batch was sampled with fixed proportions, and conventional active learning methods were applied separately for full and weak supervision.




Figure 5. Ablation study comparing the proposed method without instance-wise uncertainty $u_t(\mathbf{x})$, $u_w(\mathbf{x})$ (Ours w/o uncertainty) and without batch selection considering diversity (Ours w/o diversity). The figures show the classification accuracy [%] (↑) for each round t on CIFAR100 and CUB200.

uncertainty of each instance. Consequently, the supervision level is not optimized for each instance individually. The results show that our approach outperforms this variant on both datasets. This effect is particularly prominent in the CUB200 dataset, demonstrating the benefit of optimizing “instance-wise” supervision-level.

The second ablation variant, called ‘Ours w/o diversity,’ involves a greedy approach to batch selection, where instances are selected in descending order of VCR without considering data diversity. Our approach, which incorporates diversity into the batch selection process, consistently outperforms this variant on both datasets. The impact of diversity-aware selection is especially evident in the CIFAR100 dataset, demonstrating that considering diversity in batch selection significantly enhances model performance.

7. Conclusion, limitation, and future work

In this paper, we proposed a novel active learning (AL) approach, called Instance-wise Supervision-level Optimization (ISO), which optimizes the supervision level (either full or weak annotation) for each instance to maximize annotation efficiency within a budget. By evaluating the uncertainty, diversity, and value-to-cost ratio (VCR) of each instance for full and weak supervision, our approach automatically determines the instances to be annotated and their optimal annotation level. Experiments on classification tasks demonstrated that ISO outperforms not only standard AL methods but also state-of-the-art AL methods with weak supervision, while effectively leveraging instance-specific information to improve both accuracy and cost-efficiency.

A current limitation of this study is that the effectiveness of our ISO has been validated only in classification experiments. The concept of ISO is applicable to various tasks beyond classification. For example, ISO can be applied to object segmentation tasks, where precise yet costly object boundaries are assumed for full supervision, while simpler rectangular bounding boxes are used for weak supervision. Another current limitation is the assumption that supervision levels are only two. We can expand to deal with more supervision levels by introducing class levels other than exact classes and superclasses. In future work, we aim to demonstrate the full versatility of our approach by exploring these extensions across a broader range of tasks, reinforcing its potential as a scalable, cost-effective solution for real-world applications.

Acknowledgements: This work was supported by JSPS KAKENHI Grant Number JP23KJ1723, and JST, ACT-X Grant Number JPMJAX23CR, Japan.

References

- [1] Charu C Aggarwal, Xiangnan Kong, Quanquan Gu, Jiawei Han, and S Yu Philip. Active learning: A survey. In *Data classification*, pages 599–634. Chapman and Hall/CRC, 2014. [1](#), [2](#)
- [2] D ARTHUR. k-means++: the advantages of careful seeding. In *ACM-SIAM symposium on Discrete algorithms*, pages 1027–1035, 2007. [5](#)
- [3] Jordan T Ash, Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. Deep batch active learning by diverse, uncertain gradient lower bounds. In *International Conference on Learning Representations*. [1](#), [2](#), [3](#), [4](#), [5](#)
- [4] Shayok Chakraborty. Asking the right questions to the right users: Active learning with imperfect oracles. In *Association for the Advancement of Artificial Intelligence*, pages 3365–3372, 2020. [2](#)
- [5] Yarin Gal, Riashat Islam, and Zoubin Ghahramani. Deep bayesian active learning with image data. In *International Conference on Machine Learning*, pages 1183–1192, 2017. [1](#), [2](#)
- [6] Guy Hachohen, Avihu Dekel, and Daphna Weinshall. Active learning on a budget: Opposite strategies suit high and low budgets. In *International Conference on Machine Learning*, pages 8175–8195, 2022. [1](#), [2](#), [3](#)
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Computer Vision and Pattern Recognition*, pages 770–778, 2016. [6](#)
- [8] Neil Houlsby, Ferenc Huszár, Zoubin Ghahramani, and Máté Lengyel. Bayesian active learning for classification and preference learning. *arXiv preprint arXiv:1112.5745*, 2011. [1](#), [2](#), [3](#)
- [9] Wei-Ning Hsu and Hsuan-Tien Lin. Active learning by learning. In *Association for the Advancement of Artificial Intelligence*, 2015. [1](#), [2](#), [5](#)
- [10] Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [6](#)
- [11] Andreas Kirsch, Joost Van Amersfoort, and Yarin Gal. Batchbald: Efficient and diverse batch acquisition for deep bayesian active learning. *Neural Information Processing Systems*, 32, 2019. [1](#), [2](#), [3](#)
- [12] Dani Kiyasseh, Tingting Zhu, and David A Clifton. Soqal: Selective oracle questioning for consistency based active learning of cardiac signals. In *International Conference on Machine Learning*, pages 11302–11340, 2022. [2](#)
- [13] Alex Kulesza and Ben Taskar. k-dpps: Fixed-size determinantal point processes. In *International Conference on Machine Learning*, pages 1193–1200, 2011. [5](#)
- [14] Chao Lu and Yuexian Zou. Using coarse label constraint for fine-grained visual classification. In *Conference on Multimedia Modeling*, 2018. [5](#)
- [15] Andrew Kachites McCallum, Kamal Nigam, et al. Employing em and pool-based active learning for text classification. In *International Conference on Machine Learning*, pages 350–358, 1998. [2](#)
- [16] Pengzhen Ren, Yun Xiao, Xiaojun Chang, Po-Yao Huang, Zhihui Li, Brij B Gupta, Xiaojiang Chen, and Xin Wang. A survey of deep active learning. *ACM computing surveys (CSUR)*, 54(9):1–40, 2021. [1](#), [2](#)
- [17] Zeyu Ren, Shuihua Wang, and Yudong Zhang. Weakly supervised machine learning. *CAAI Transactions on Intelligence Technology*, 8(3):549–580, 2023. [1](#), [2](#)
- [18] Dan Roth and Kevin Small. Margin-based active learning for structured output spaces. In *European Conference on Machine Learning*, pages 413–424, 2006. [1](#), [2](#), [3](#), [5](#), [6](#)
- [19] Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. In *International Conference on Learning Representations*, 2018. [2](#), [3](#), [4](#), [5](#)
- [20] Burr Settles. Active learning literature survey. 2009. [2](#)
- [21] Samarth Sinha, Sayna Ebrahimi, and Trevor Darrell. Variational adversarial active learning. In *International Conference on Computer Vision*, pages 5972–5981, 2019. [1](#), [2](#), [3](#)
- [22] Javier Gamazo Tejero, Martin S Zinkernagel, Sebastian Wolf, Raphael Sznitman, and Pablo Márquez-Neila. Full or weak annotations? an adaptive strategy for budget-constrained annotation campaigns. In *Conference on Computer Vision and Pattern Recognition*, pages 11381–11391, 2023. [1](#), [2](#), [3](#), [5](#), [7](#)
- [23] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. Technical report, 2011. [5](#)
- [24] Dan Wang and Yi Shang. A new active labeling method for deep learning. In *International Joint Conference on Neural Networks*, pages 112–119, 2014. [1](#), [2](#), [3](#), [5](#)
- [25] Donggeun Yoo and In So Kweon. Learning loss for active learning. In *Computer Vision and Pattern Recognition*, pages 93–102, 2019. [1](#), [2](#)
- [26] Man Zhang, Yong Zhou, Jiaqi Zhao, Yiyun Man, Bing Liu, and Rui Yao. A survey of semi-and weakly supervised semantic segmentation of images. *Artificial Intelligence Review*, 53:4259–4288, 2020. [1](#), [2](#)
- [27] Zhi-Hua Zhou. A brief introduction to weakly supervised learning. *National science review*, 5(1):44–53, 2018. [1](#), [2](#)