

Beyond Active Learning: Leveraging the Full Potential of Human Interaction via Auto-Labeling, Human Correction, and Human Verification

Nathan Beck Krishnateja Killamsetty Suraj Kothawade Rishabh Iyer
 University of Texas, Dallas

{nathan.beck, krishnateja.killamsetty, suraj.kothawade, rishabh.iyer}@utdallas.edu

Abstract

Active Learning (AL) is a human-in-the-loop framework to interactively and adaptively label data instances, thereby enabling significant gains in model performance compared to random sampling. AL approaches function by selecting the hardest instances to label, often relying on notions of diversity and uncertainty. However, we believe that these current paradigms of AL do not leverage the full potential of human interaction granted by automated label suggestions. Indeed, we show that for many classification tasks and datasets, most people verifying if an automatically suggested label is correct take $3\times$ to $4\times$ less time than they do changing an incorrect suggestion to the correct label (or labeling from scratch without any suggestion). Utilizing this result, we propose CLARIFIER (aCTive LeARnIng From tIEred haRdness), an Interactive Learning framework that admits more effective use of human interaction by leveraging the reduced cost of verification. By targeting the hard (uncertain) instances with existing AL methods, the intermediate instances with a novel label suggestion scheme using submodular mutual information functions on a per-class basis, and the easy (confident) instances with highest-confidence auto-labeling, CLARIFIER can improve over the performance of existing AL approaches on multiple datasets – particularly on those that have a large number of classes – by almost $1.5\times$ to $2\times$ in terms of relative labeling cost.

1. Introduction

In recent years, machine learning practitioners have employed the use of deep neural networks (DNNs) to achieve state-of-the-art performance on many tasks, such as image classification, object detection, language translation, and so forth. The performance of these DNNs largely owe their success to increasingly large datasets used in training. Unfortunately, procuring large training datasets often requires extensive labeling efforts in the process. Indeed, these labeling efforts can be prohibitively expensive considering the

task. A prime example of this is medical imaging, where the labels are provided by trained radiologists whose services can cost well over 100 USD per hour. As such, the need for label-efficient approaches has become a hot topic for deep learning.

Active learning (AL) is one such method that addresses these labeling costs. AL methods aim to choose a cardinality-constrained subset of the unlabeled dataset whose addition to the labeled dataset after being labeled produces the largest gain in performance. Such methods do so by leveraging information contained within a working model, its training dataset, and an unlabeled dataset. Typically, such methods use this information to select a possibly diverse set of instances whose label predictions are uncertain. Examples of such methods include entropy sampling [16], which chooses unlabeled instances whose class probability distributions have maximal entropy, and BADGE [3], which chooses a diverse set of uncertain instances by applying K-MEANS++ initialization [1] to the space of hypothesized loss gradients of the unlabeled dataset. In general, these selected instances are hard to predict for the current model; thus, their inclusion in the labeled dataset tends to improve model performance after they are labeled.

While AL methods are a label-efficient way to improve model performance, most AL approaches focus on the chosen subset size as the metric of labeling cost. However, recent efforts in industry¹ are placing more emphasis on *automating label assignment of portions of the subset selection process and enabling humans to leverage label suggestions when assigning labels*. Hence, the time that an annotator spends labeling an instance can differ depending on whether he or she can utilize a correct label suggestion or has to fix an incorrect label suggestion. As a result, this necessitates using a new labeling cost metric since the true cost of labeling is dependent not only on the number of selected instances but also on the time spent to label each, which can change depending on the correctness of the label suggestion.

¹See AWS's Amazon SageMaker Automated Labeling Service <https://docs.aws.amazon.com/sagemaker/latest/dg/sms-automated-labeling.html> and <https://thehive.ai>

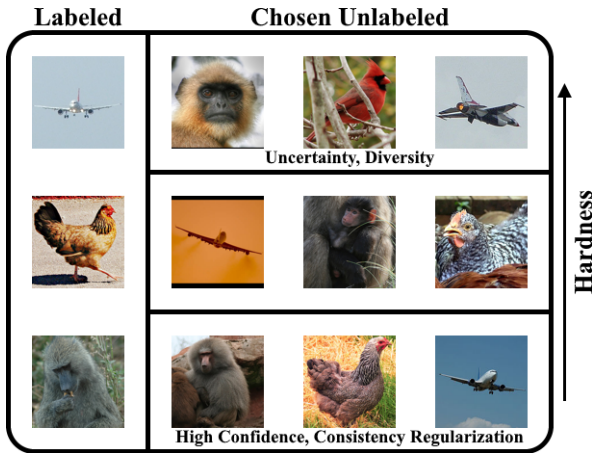


Figure 1. Depiction of the unlabeled instances selected by differing frameworks (images from STL-10 [5]). Active Learning (AL) usually selects harder instances, often maximizing uncertainty, diversity, or both. Semi-Supervised Learning (SSL) usually selects instances whose classes are easy to predict (high confidence) or are otherwise consistent across some augmentation of the input (consistency regularization). In between these two extremes are those instances whose correct label information can impart better model performance than the auto-labeled instances of SSL while being easier to make accurate label suggestions for human labelers than pure AL.

Another related paradigm is the concept of automatically labeling or pseudo-labeling unlabeled data instances, something that is followed heavily in the semi-supervised learning (SSL) literature [17]. While these auto-labeled instances come at no cost, these instances are often those that are easy to label and do not impart as much performance benefit as the harder points selected by AL do. Furthermore, there are those instances that lie within the two label hardness extremes of AL and SSL that have not been highlighted but have desirable benefits when incorporating label suggestions in the annotation process. Indeed, verifying a correctly suggested label is typically a less costly procedure than fixing an incorrectly suggested label, and the instances lying between the two hardness extremes are likely to impart more performance benefit than auto-labeled instances while retaining a significant portion of the suggested label accuracy. We summarize the relationship between all three regimes in Figure 1. In summary, methods that factor in the benefits of all three regimes into their selection and present a meaningful computation of the labeling cost for those selected points are needed to reflect more faithful returns in model performance versus the labeling cost.

1.1. Our Contributions

To confirm the claim that it is easier for a human to verify a suggested label than to correct it, we conduct a human study across four datasets to explore the values of the cost of fixing

an incorrectly suggested label (c_a) and the cost of verifying a correctly suggested label (c_v). We show that $c_a : c_v$ is between 3 to 4, suggesting that it takes $3\times$ to $4\times$ longer to fix incorrect label suggestions (to label from scratch) than to verify correct label suggestions (see Figure 2).

As it is often the case that labels can be suggested to annotators, we offer an improved baseline method for incorporating label suggestions into any AL algorithm (which we call SUGGEST) – we simply suggest the model’s label prediction on each selected AL instance to the annotator. By performing a simple AL experiment with BADGE, we observe that BADGE with SUGGEST is $1.15\times$ to $1.4\times$ more cost-efficient (label-efficient) compared to BADGE without SUGGEST (see Figure 3).

While the improved baseline SUGGEST is a good starting point, it does not provide significant gains for datasets that have many classes. We propose a new framework called CLARIFIER, which utilizes tiered hardness by considering all three of the regimes mentioned in Figure 1 in an effort to most efficiently leverage the true cost of labeling. Our method supplements existing AL methods via the following: 1) The SUGGEST paradigm is used with existing AL methods to select hard instances with label suggestions. 2) The instances of intermediate hardness are chosen via the use of submodular mutual information (SMI) [7], which effectively trades off between typical measures favored by AL methods (such as diversity and uncertainty) and the relevance to the currently labeled instances while also providing a novel label suggestion scheme. Lastly, 3) our method selects the easy instances via common auto-labeling techniques to benefit from auto-labeling approaches. Notably, this framework admits most choices of AL algorithms and auto-labeling algorithms while providing a meaningful way to calculate labeling cost in accordance with recent labeling frameworks.

Using our method, we show an increase of 2% to 5% in test accuracy over the simple AL label suggestion scheme mentioned in our second contribution across CIFAR-100 [12], Caltech-UCSD Birds-200-2011 [19], and Stanford Dogs [8] in Figure 4. Additionally, we show labeling efficiencies of up to $2\times$ with respect to existing AL methods without label suggestions. We then study the label suggestion accuracy of each component of our method in Figure 5, showing that each component targets its intended regime of Figure 1. We conclude by highlighting the importance of each component via an ablation study of our method in Figure 6.

2. Related Work

Our method is positioned in a vast space of existing AL paradigms, lending part of its functionality to recent works in AL. One such work is BADGE [3], which aims to select hard instances via notions of uncertainty and diversity. BADGE selects unlabeled instances by representing each through its

hypothesized loss gradients at the model’s last-layer parameters, wherein the hypothesized loss is calculated using the highest-confidence class as the pseudo-label. By using K-MEANS++ initialization [1] on these representations, BADGE effectively selects a set of instances whose gradients have high magnitudes (and are thus uncertain) and are diverse. BAIT [2], a recent state-of-the-art AL method, employs the use of Fisher information matrices to make informed decisions for batch selection. [2] additionally link BAIT as a generalization of BADGE as its gradient embeddings correspond to rank-one approximations of BAIT’s Fisher information matrices. Other AL methods select hard instances by directly modeling the prediction uncertainty of each unlabeled instance, such as entropy sampling, least-confidence sampling, and margin sampling [16]. While the introduction of these methods are not as recent, [4] show that the inclusion of data augmentation techniques can often bring these simple uncertainty-based techniques to match state-of-the-art AL approaches on many common image classification datasets. In our work, we utilize BADGE [3] and entropy sampling [16] as possible choices of AL methods to be used in our framework amidst a vast space of AL paradigms.

Our chosen manner of selecting the instances of intermediate hardness also utilizes submodular information measures [7], which have recently been used to perform a wide variety of selection tasks in AL. In particular, submodular mutual information (SMI) has been used to perform targeted subset selection via a simple cardinality-constrained optimization method detailed in [13]. To capture a variety of desirable aspects of the selected subset of targeted instances, [11] study the use of a wide variety of SMI instantiations, offering insights to the utility of each in the selection process. [9] utilize many of the SMI instantiations offered in [11] and the hypothesized gradient functionality of BADGE [3] to perform targeted AL in the presence of rare classes. Lastly, [10] apply the use of SMI for selecting rare-slice images to label in the context of object detection. In our work, we utilize SMI [7] to query for instances of intermediate hardness that are related to a labeled instance subset of uniform class as a means of procuring suggested labels (detailed in Section 4).

3. The Benefit of Label Suggestions

To further motivate our framework, we more closely examine the cost of labeling in recent pipelines, where label suggestions are provided to the human annotators. To better capture the cost of labeling in pipelines with label suggestions, we model the labeling cost as $c_v \times n_{correct} + c_a \times (n - n_{correct})$, where c_v is the cost to verify a correct suggested label and c_a is the cost to assign a corrected label when an incorrect suggestion is given. Recent AL works have framed the cost of labeling in terms of the number of selected unlabeled instances; however, the cost of the final assignment

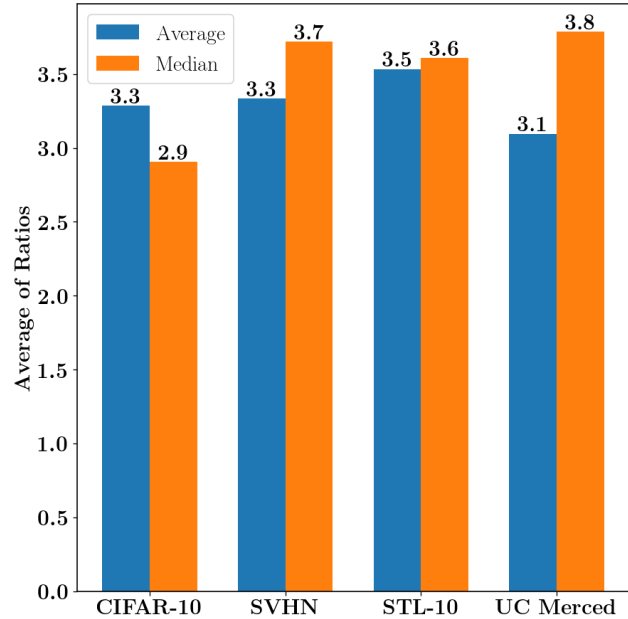


Figure 2. Average of each subject’s average/median ratios for c_a and c_v across each of the four datasets examined in our labeling experiment. We see that most datasets admit a 3 – 4× increase in c_a vs c_v .

of an instance’s label fluctuates depending on the suggested label. Instances with correct label suggestions require less time in assigning final labels (c_v) versus those with incorrect label suggestions (c_a) due to the extra correction effort. Hence, to get the best returns in model performance versus the labeling cost, AL methods must select points that impart the most information to the model while also providing accurate suggested labels.

To better understand the impact of suggested label verification on the model performance gain of these pipelines, further study about the values of c_a and c_v is needed. Indeed, the value of each depends on an annotator’s proficiency in using the labeling pipeline, his or her knowledge of the target domain, and inherent properties of the data itself. To this end, we conduct a labeling experiment to ascertain the ranges of values that c_a and c_v can take for different datasets. In our experiment, voluntary human subjects utilize a labeling tool that iterates over images in the CIFAR-10 [12], STL-10 [5], SVHN [14], and UC Merced Land Use [20] datasets. The subject is asked to label 100 images of each dataset (105 images for UC Merced Land Use) before moving on to the next dataset. During the labeling process, the tool offers label suggestions for each image, where the suggested label has a 50% chance of being the correct ground truth label. The time taken to label each image is recorded, and images that are incorrectly labeled by the subject are discarded. Using the correctly labeled images, the average and median times to

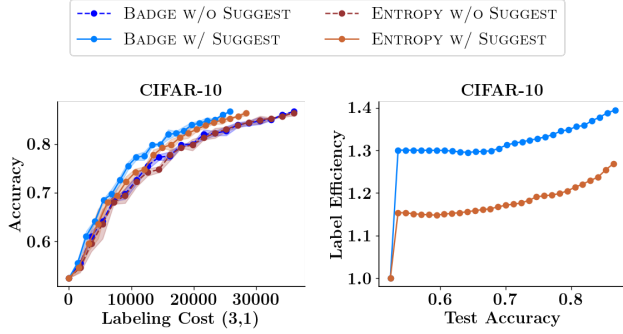


Figure 3. Improvement in test accuracy and labeling efficiency on CIFAR-10 by introducing label suggestions. Here, we see labeling efficiencies between $1.15\times$ and $1.4\times$ versus AL baselines without label suggestions when $c_a = 3, c_v = 1$.

verify correctly suggested image labels are calculated, which we pose as the subject’s possible values of c_v . Likewise, the average and median times to fix incorrectly suggested image labels are calculated, which we pose as the subject’s possible values of c_a . We present these metrics in Figure 2, showing that $c_a : c_v$ often resides in the 3 : 1 to 4 : 1 range. More details about the experiment and the labeling tool used can be found in Appendix B.

3.1. An Improved Baseline: SUGGEST

In order to effectively enable suggestions in existing active learning approaches, we propose SUGGEST, a simple baseline that essentially uses the current model’s label prediction as the suggested label for each of the selected AL instances. To visualize the effect that the SUGGEST paradigm has on existing AL methods, we run an experiment on CIFAR-10 [12] to show the better gains in model performance and labeling efficiency that AL with label suggestions has with respect to AL without label suggestions. Here, we define labeling efficiency as the ratio between the labeling costs needed for two AL methods to reach a desired test accuracy as is done in [4].

Using appropriate choices of c_a and c_v obtained from our labeling experiment, we can now give more faithful estimates of the true cost of labeling in pipelines with automatic label suggestions. Based on the results given in Figure 3, we can conclude that there is a benefit in applying automated label suggestions for leveraging more out of the human labeling effort. While these AL methods improve in their labeling efficiency, however, they do not consider leveraging *more* of the human labeling effort via label suggestions by targeting instances that admit easier label suggestions while retaining some of the hardness properties that they use in their selection. We expand upon this idea to present our CLARIFIER framework in Section 4, comparing against the SUGGEST baseline in Section 5.

4. Our Proposed Framework: CLARIFIER

Here, we present CLARIFIER, a new Interactive Learning framework for leveraging the most out of the human labeling effort via automatic label suggestions. In particular, our framework presents three selection methodologies for selecting instances that span across the degrees of label suggestion hardness as detailed in Figure 1. Notably, our framework utilizes the trade-off between query relevance and uncertainty when selecting instances to label. Instances that are relevant to the labeled set (e.g., are similar to the labeled set) naturally admit confident label predictions but tend to not impart new knowledge for improving model performance. Instances that are *not* relevant to the labeled set are naturally harder to classify correctly (as the prediction is more uncertain) but tend to impart more knowledge for improving model performance. By varying the degree at which query relevance and prediction uncertainty are used, a selection method can effectively target points within one of the hardness tiers mentioned in Figure 1. By choosing instances from all hardness tiers, our method enjoys the benefits that each tier provides in the context of AL with automatic label suggestions, which we detail in this section. We summarize our framework in Algorithm 1.

Algorithm 1 CLARIFIER

- 1: **Input:** Labeled set \mathcal{L} , Unlabeled set \mathcal{U} , Model M , Model parameters θ , Hard budget b_1 , Intermediate budget b_2 , Easy budget b_3
 - 2: **repeat**
 - 3: $X_{hard}, \hat{y}_{hard} \leftarrow \text{AL_Select}(\mathcal{L}, \mathcal{U}, M, \theta, b_1)$
 - 4: $X_{int}, \hat{y}_{int} \leftarrow \text{SMI_Suggest}(\mathcal{L}, \mathcal{U}, M, \theta, b_2)$
 - 5: $X_{easy}, y_{easy} \leftarrow \text{Auto_Label}(\mathcal{L}, \mathcal{U}, M, \theta, b_3)$
 - 6: $y_{hard} \leftarrow \text{Label}(X_{hard}, \hat{y}_{hard})$
 - 7: $y_{int} \leftarrow \text{Label}(X_{int}, \hat{y}_{int})$
 - 8: $X \leftarrow X_{hard} \cup X_{int} \cup X_{easy}$
 - 9: $y \leftarrow y_{hard} \cup y_{int} \cup y_{easy}$
 - 10: $\mathcal{U} \leftarrow \mathcal{U} - X$
 - 11: $\mathcal{L} \leftarrow \mathcal{L} \cup \left[\bigcup_{i=1}^{|X|} (X_i, y_i) \right]$
 - 12: $\theta \leftarrow \text{Train}(\mathcal{L}, M, \theta)$
 - 13: **until** converged
-

Selecting Hard Instances (AL_Select): While choosing instances that are relevant to the labeled set admits more accurate label suggestions, strategies that over-focus on maintaining highly confident label suggestions tend to select instances that are not as informative for model training as other instances in the unlabeled set. For this reason, using existing AL methods for selecting hard instances is favorable for maintaining consistent improvement in model performance. As previously discussed, however, selecting too many hard instances *en masse* tends to neglect the benefit of reduced

labeling cost granted by correct label verification. As such, using existing AL methods to select hard instances can be done in moderation to glean important knowledge from the unlabeled set while being cognizant of the labeling cost. In our experiments, we examine the use of BADGE [3] and entropy sampling [16] as the AL component of our framework.

Selecting Intermediate Instances (SMI_Suggest): As previously mentioned, selecting too many hard instances misses out on the benefit of correct label verification. On the other hand, selecting too many easy instances misses out on the knowledge gained from the hard instances. To gain benefits from both aspects, instances that lie between these two hardness extremes are desired. To target these instances, methods that factor in relevance to the labeled set and notions of uncertainty can be used. To that end, we utilize the recently proposed submodular mutual information [7] to select instances that fall between both hardness extremes and to provide their suggested labels. Submodular mutual information (SMI), denoted $I_f(A; Q)$, measures the overlap in information between two sets of instances A, Q with respect to a submodular information function f . By selecting f to model diversity information (to get instances that are not over-focused on relevance), one can capture both query relevance and diversity/uncertainty information through $I_f(A; Q)$. To select instances from the unlabeled set and to provide their suggested labels, the following cardinality-constrained optimization problem is solved on a per-class basis using the simple greedy algorithm proposed by [13]:

$$\max_{A_c \subseteq \mathcal{U}, |A_c| \leq k} I_f(A_c; \mathcal{L}_c) \quad (1)$$

where \mathcal{L}_c denotes the set of labeled instances belonging to class c . Hence, each selected subset A_c is a set of unlabeled instances that have query relevance to \mathcal{L}_c ; accordingly, we assign a suggested label c for each point in A_c . Each A_c (along with each suggested label) is coalesced to form the set of selected instances of intermediate hardness. For our experiments, we utilize LOGDETMI [7, 9, 11] to capture query relevance and diversity/uncertainty information. More details can be found in Appendix C.

Selecting Easy Instances (Auto_Label): The last component of our framework addresses those instances that are highly relevant to the labeled set. In many cases, there are cases where the model or some other mechanism can make highly confident label predictions for some of the unlabeled instances. The high-confidence of these instances tends to originate from high relevance to the labeled set, so these instances tend to not impart as much knowledge for improving model performance. However, as these instances admit high-confidence predictions, one can choose to automatically

assign the suggested label as the ground-truth label instead of presenting them to human labelers, which may present a waste of labeling effort. Hence, these instances can be added to the labeled set at no additional cost, imparting some additional knowledge for improving the model. Such selection is not without risk, however, as there is a risk of auto-assigning noisy labels; hence, the budget of this component of our framework can be attenuated according to the accuracy of the chosen auto-labeling method. In our experiments, we select the easy instances by choosing those instances in the unlabeled set with the highest-confidence predictions.

5. Experiments

In this section, we conduct experiments to answer the following questions: **1)** Utilizing our method, what is the improvement in accuracy over AL versus the labeling cost discussed in Section 3? **2)** How hard is it to predict the correct label for the instances selected by each component; *e.g.*, is it true that the label hardness presented in Figure 1 is correct? Lastly, **3)** Are the different components of our method necessary? We analyze our method across three datasets – CIFAR-100 [12], Caltech-UCSD Birds-200-2011 [19], and Stanford Dogs [8] – to answer each.

5.1. Experimental Setting

Here, we enumerate various aspects of the experimental setting used in our experiments. Additional reproducibility information is given in Appendix A.

Compute and Datasets: To run our experiments, we utilize a machine with 96 GB of RAM, an NVIDIA RTX A6000 GPU with 48 GB of VRAM, and an Intel Xeon Silver 4208 CPU. To effectively utilize all these components, we implement our code in PyTorch [15]. Using this architecture, we conduct experiments across three datasets: CIFAR-100 [12], Caltech-UCSD Birds-200-2011 [19], and Stanford Dogs [8]. CIFAR-100 is an image classification dataset of 60k images spread across 100 classes. Caltech-UCSD Birds-200-2011 is a fine-grained image classification dataset of approximately 12k bird images, each of which is assigned one of 200 possible species. Lastly, Stanford Dogs is another fine-grained image classification dataset of approximately 21k dog images, each of which is assigned one of 120 possible breeds.

Model Architecture: For each dataset, we utilize the ResNet-18 [6] architecture for image recognition. We reset the model parameters after each round of labeling following the observation that resetting parameters leads to better generalization performance as detailed in [4]. However, both the Caltech-UCSD Birds-200-2011 and Stanford Dogs datasets do not admit enough data to fully train all parameters of the ResNet-18 model without severe underfitting. To rectify this,

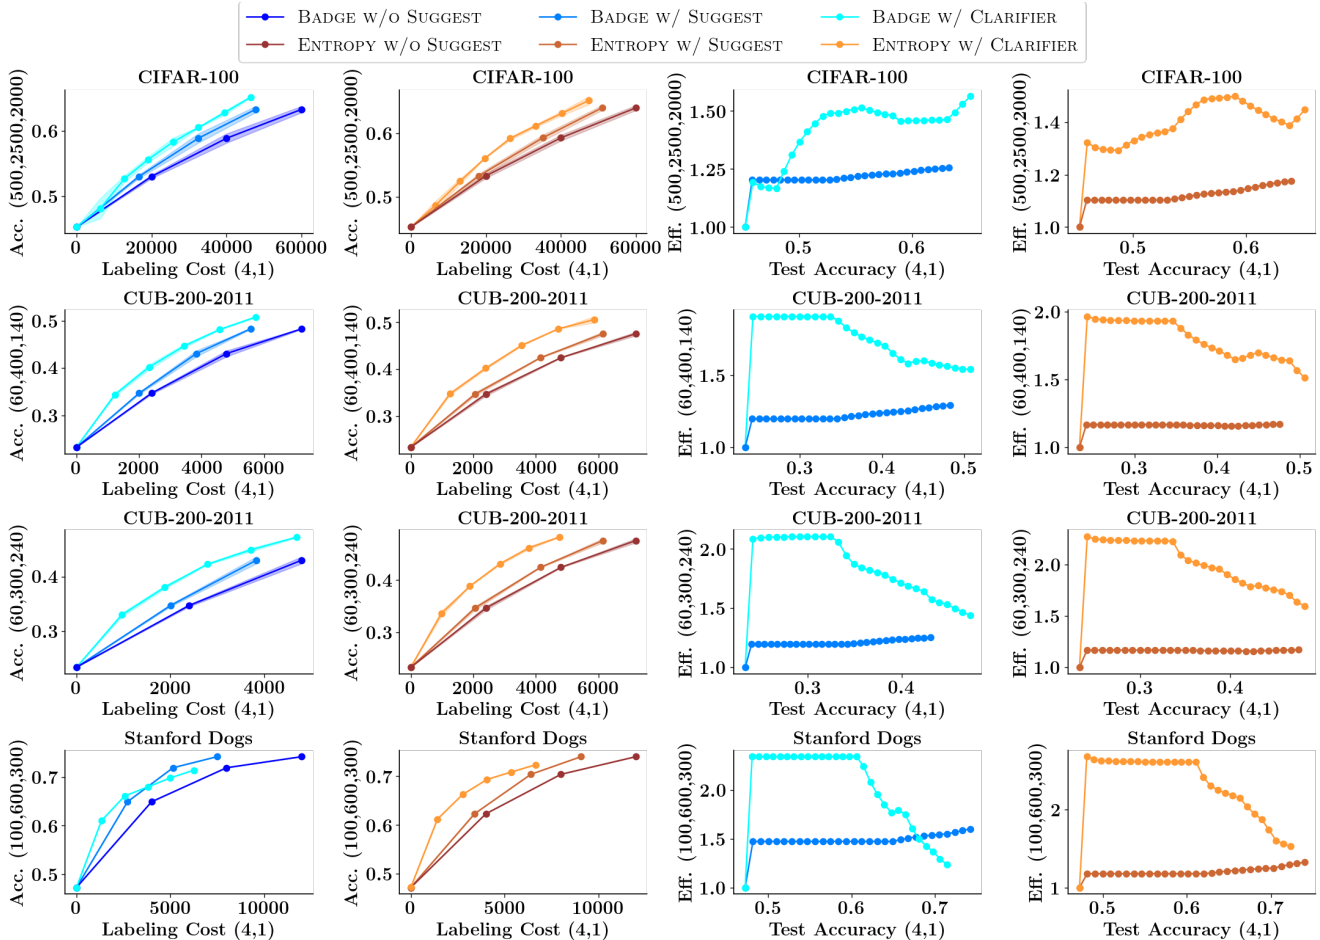


Figure 4. Accuracies on all four datasets. Using CLARIFIER shows improvement over existing AL methods and the SUGGEST baseline discussed in Section 3. The labeling cost is calculated based on c_a and c_v values that are observed in Section 3.

we opt to use pre-trained feature extractor weights obtained from PyTorch’s collection of pre-trained models for these two datasets. These weights are kept frozen throughout each experiment on these two datasets; only the fully connected layer is reset and trained after each selection.

Training: To train each model between labeling rounds, we utilize PyTorch’s SGD optimizer with momentum ($\mu = 0.9$) and weight decay ($\lambda = 5 \times 10^{-4}$). Additionally, we use a cosine-annealing learning rate with an initial setting of $\gamma = 0.01$. T_{max} is set to the maximum number of epochs that we allow in training each model, which we set to 300 for CIFAR-100 and 500 for Caltech-UCSD Birds-200-2011 and Stanford Dogs. Additionally, we stop training the model parameters once training accuracy has reached 99% on CIFAR-100 and 95% on the other datasets. We employ data augmentations during training to leverage better model generalization as is discussed in [4]. On CIFAR-100, we utilize random crops, random horizontal flips ($p = 0.5$), and

normalization for training; we utilize only normalization during testing. For the other datasets, we resize each image to a 256×256 dimension before applying random cropping (to 224×224 images), random horizontal flips ($p = 0.5$), and normalization for training; we resize each image to a 256×256 dimension, apply center cropping (to 224×224 images), and normalization for testing.

Selection Partitioning: While some selection methods are effective at choosing hard (BADGE) or otherwise targeted (SMI) instances, the size of the unlabeled set often imposes prohibitive memory requirements on some systems. Following the strategy used in [9], we randomly partition the unlabeled set into equal-sized chunks and perform selections on each, equally partitioning the selection budget among each chunk. In doing so, these methods that would otherwise not scale effectively to the large unlabeled set can be utilized on reasonably sized portions of the unlabeled set. On CIFAR-100, we partition the unlabeled dataset into 5 partitions when

selecting intermediate instances with SMI. Additionally, we partition the dataset into 5 partitions when using AL to select instances in CIFAR-100. Otherwise, we do not perform any other partitioning when selecting instances to label.

Seed Size and Budgets: To begin each experiment, we choose to start with labeled seed sets of size 10k, 600, and 500 for CIFAR-100, Caltech-UCSD Birds-200-2011, and Stanford Dogs, respectively. We then choose values of b_1 , b_2 , and b_3 as described in Algorithm 1 for each dataset, which are listed in order in parentheses for each plot in Figure 4. For the compared AL baseline in each plot, we utilize a budget equal to the sum of b_1 , b_2 , and b_3 so that the same number of instances are added per labeling round.

5.2. Accuracy Improvement

Here, we present the results of each experiment. To produce error bars, we run each compared method 3 times across CIFAR-100. We run each compared method 2 times for Caltech-UCSD Birds-200-2011 and 1 time for Stanford Dogs. We show improvement in test accuracy versus labeling cost in Figure 4, along with improvement in labeling efficiency (see Section 3). Utilizing the labeling experiment results in Section 3, we assume c_a, c_v values as stated in the parentheses in Figure 4 to calculate the labeling cost. CLARIFIER enjoys accuracy improvements of up to 3% on CIFAR-100 and 5% on Caltech-UCSD Birds-200-2011 and Stanford Dogs. Accordingly, the labeling efficiencies for these datasets range from $1.0\times$ to $2.5\times$ when comparing CLARIFIER to existing AL approaches without label suggestions, demonstrating more efficient use of the human labeling effort (even compared to the same AL approaches with label suggestions given by SUGGEST). We conclude that our method returns more improvement in model performance versus the labeling cost, especially for datasets with large class counts.

5.3. Label Suggestion Hardness

Next, we study the hardness of the suggested label for each of the domains presented in Figure 1. To do so, we present the label suggestion accuracy of each component during each labeling round of Caltech-UCSD Birds-200-2011 and Stanford Dogs in Figure 5. Interestingly, we find that BADGE enjoys better suggestion accuracy over entropy sampling, which most likely occurs since entropy sampling directly selects those instances whose class probabilities are most uncertain. As shown in each column, the label suggestion accuracy for instances selected by AL methods is consistently less than that for instances selected by LOGDETM. Likewise, the label suggestion accuracy for instances selected by LOGDETM is consistently less than that for instances selected by the auto-labeler. This verifies

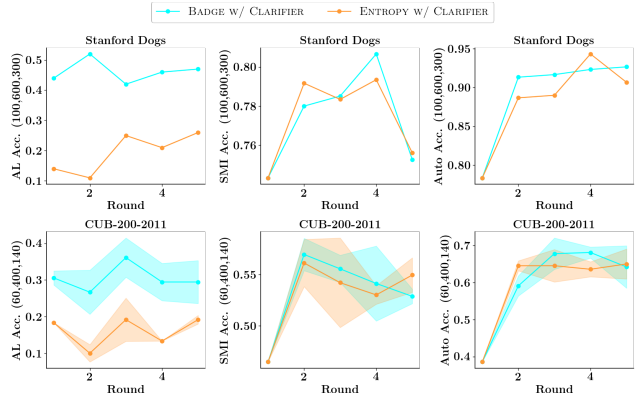


Figure 5. Suggestion accuracies on Stanford Dogs and Caltech-UCSD Birds-200-2011. Across each experiment, the AL methods selects the hardest instances, followed by SMI. Auto-labeling selects the easiest examples to label.

our claim that each component selects instances from its targeted domain as portrayed in Figure 1.

5.4. Ablation

Lastly, we study the effect that each component has on the returns in model performance versus the labeling cost. We present ablated versions of a CIFAR-100 experiment using the ENTROPY+LOGDETM+HIGHEST-CONFIDENCE configuration of CLARIFIER in Figure 6. In each, we examine the effect of setting the selection budget of each component to 0. Note that we already present pure AL methods in Figure 4. Furthermore, we do not present the configuration where only auto-labeled instances are assigned. Such configurations have been more heavily studied in SSL, and since our auto-labeling component makes hard assignments on easy instances, the test accuracy is not likely to improve significantly. Lastly, we present these ablated versions using multiple configurations for c_a and c_v to further analyze the contribution of each component in different cost settings.

Across all configurations, we see that the performance of the ablated version either matches that of the original ENTROPY+LOGDETM+HIGHEST-CONFIDENCE CLARIFIER configuration or is below it. Notably, Figure 6 suggests that the $c_a : c_v$ ratio directly influences the model performance gain versus the labeling cost of each component of CLARIFIER. Namely, we observe that the AL component (which selects the hard instances) is more beneficial in leveraging the labeling cost when the $c_a : c_v$ ratio is small since its removal results in a noticeable performance degradation. When the $c_a : c_v$ ratio is increased, this performance degradation is nullified, suggesting that the AL component becomes less beneficial as the $c_a : c_v$ increases. We observe the opposite behavior with the SMI component (which selects the intermediate instances). Indeed, there is a noticeable

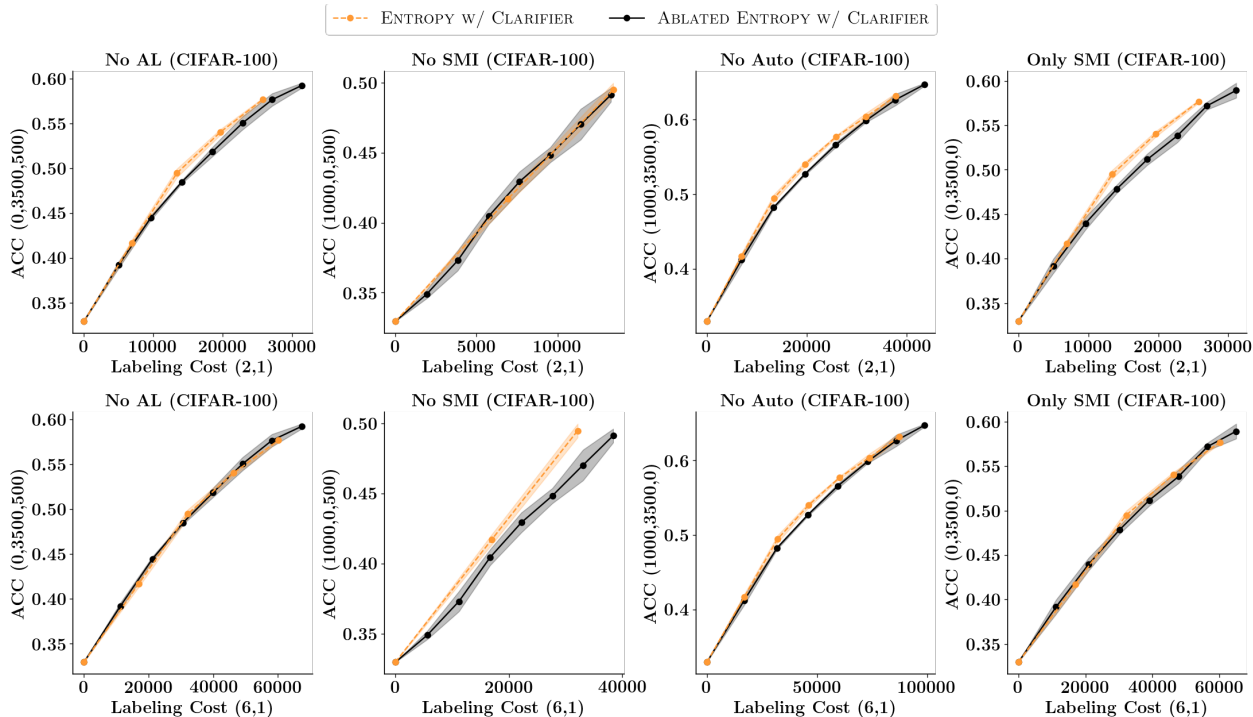


Figure 6. Ablation studies across CIFAR-100 using an initial labeled seed set of 5k instances. Removing a component from CLARIFIER results in poorer test accuracy improvement versus the labeling cost. The AL component is more useful with lower $c_a : c_v$; the SMI component is more useful with higher $c_a : c_v$; the auto-labeling component helps improve accuracy at no additional cost.

performance degradation when removing SMI in the higher $c_a : c_v$ setting presented in Figure 6, but this performance degradation is nullified in the smaller $c_a : c_v$ setting. Lastly, we observe that there is a slight performance degradation when removing the auto-labeling component, which is unaffected by the $c_a : c_v$ ratio. We conclude that each component is necessary to fully leverage the human labeling effort to maximize the gain in model performance across multiple cost scenarios.

6. Discussion

In this work, we more closely examine the need for Interactive Learning frameworks that match the trend of increased use of auto-labeling. Accordingly, we frame this increased use of AL methods with auto-labeling in the context of label prediction hardness. Through a human labeling experiment, we show that verifying correctly suggested labels is less costly than fixing incorrectly suggested labels, highlighting a potential source of improvement in gaining model performance versus the true human labeling cost. By combining this result with the observation of label prediction hardness, we formulate CLARIFIER, a new Interactive Learning framework that targets instances from each hardness tier to leverage better model performance gains out of the labeling effort. By our experiments, we show that this framework

achieves better gains in accuracy over existing AL methods and reinforces our view of label prediction hardness.

As future work, more adaptive choices of each component’s budget can be studied to leverage even more control of the labeling effort as the choice of the budget parameters are not studied heavily, which limits the direct application of CLARIFIER. We also stress that the use of CLARIFIER is more applicable with large class counts and is highly dependent on the human component: those annotators with low $c_a : c_v$ ratios will not see marked benefit in using CLARIFIER. As we view CLARIFIER as a foundational work in this setting, we do not anticipate any clear negative societal impacts.

Acknowledgements

We would like to acknowledge support from NSF Grant Number IIS-2106937, a gift from Google Research, and the Adobe Data Science Research Award. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the NSF, Google Research, Adobe Data Science, or the U.S. Government.

References

- [1] David Arthur and Sergei Vassilvitskii. k-means++: The advantages of careful seeding. Technical report, Stanford, 2006. [1](#), [3](#)
- [2] Jordan T Ash, Surbhi Goel, Akshay Krishnamurthy, and Sham Kakade. Gone fishing: Neural active learning with fisher embeddings. *arXiv preprint arXiv:2106.09675*, 2021. [3](#)
- [3] Jordan T Ash, Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. Deep batch active learning by diverse, uncertain gradient lower bounds. *arXiv preprint arXiv:1906.03671*, 2019. [1](#), [2](#), [3](#), [5](#), [13](#)
- [4] Nathan Beck, Durga Sivasubramanian, Apurva Dani, Ganesh Ramakrishnan, and Rishabh Iyer. Effective evaluation of deep active learning on image classification tasks. *arXiv preprint arXiv:2106.15324*, 2021. [3](#), [4](#), [5](#), [6](#), [10](#), [11](#)
- [5] Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 215–223. JMLR Workshop and Conference Proceedings, 2011. [2](#), [3](#), [11](#)
- [6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [5](#)
- [7] Rishabh Iyer, Ninad Khargoankar, Jeff Bilmes, and Himanshu Asanani. Submodular combinatorial information measures with applications in machine learning. In *Algorithmic Learning Theory*, pages 722–754. PMLR, 2021. [2](#), [3](#), [5](#), [12](#)
- [8] Aditya Khosla, Nityananda Jayadevaprakash, Bangpeng Yao, and Li Fei-Fei. Novel dataset for fine-grained image categorization. In *First Workshop on Fine-Grained Visual Categorization, IEEE Conference on Computer Vision and Pattern Recognition*, Colorado Springs, CO, June 2011. [2](#), [5](#), [11](#)
- [9] Suraj Kothawade, Nathan Beck, Krishnateja Killamsetty, and Rishabh Iyer. Similar: Submodular information measures based active learning in realistic scenarios. *Advances in Neural Information Processing Systems*, 34, 2021. [3](#), [5](#), [6](#), [13](#)
- [10] Suraj Kothawade, Saikat Ghosh, Sumit Shekhar, Yu Xiang, and Rishabh Iyer. Talisman: Targeted active learning for object detection with rare classes and slices using submodular mutual information. *arXiv preprint arXiv:2112.00166*, 2021. [3](#), [13](#)
- [11] Suraj Kothawade, Vishal Kaushal, Ganesh Ramakrishnan, Jeff Bilmes, and Rishabh Iyer. Prism: A rich class of parameterized submodular information measures for guided subset selection. *arXiv preprint arXiv:2103.00128*, 2021. [3](#), [5](#), [13](#)
- [12] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. [2](#), [3](#), [4](#), [5](#), [11](#)
- [13] George L Nemhauser, Laurence A Wolsey, and Marshall L Fisher. An analysis of approximations for maximizing submodular set functions—i. *Mathematical programming*, 14(1):265–294, 1978. [3](#), [5](#), [12](#), [13](#)
- [14] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011. [3](#), [11](#)
- [15] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. [5](#), [11](#)
- [16] Burr Settles. Active learning literature survey. 2009. [1](#), [3](#), [5](#)
- [17] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 596–608. Curran Associates, Inc., 2020. [2](#)
- [18] Tzutalin. Labelimg, 2015. [11](#)
- [19] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011. [2](#), [5](#), [11](#)
- [20] Yi Yang and Shawn Newsam. Bag-of-visual-words and spatial extensions for land-use classification. In *Proceedings of the 18th SIGSPATIAL international conference on advances in geographic information systems*, pages 270–279, 2010. [3](#), [11](#)