

Fusing Conditional Submodular GAN and Programmatic Weak Supervision

Kumar Shubham, Pranav Sastry, Prathosh AP

Indian Institute of Science, Bangalore, India
shubhamkuma3@iisc.ac.in, pranavsastri@iisc.ac.in, prathosh@iisc.ac.in

Abstract

Programmatic Weak Supervision (PWS) and generative models serve as crucial tools that enable researchers to maximize the utility of existing datasets without resorting to laborious data gathering and manual annotation processes. PWS uses various weak supervision techniques to estimate the underlying class labels of data, while generative models primarily concentrate on sampling from the underlying distribution of the given dataset. Although these methods have the potential to complement each other, they have mostly been studied independently. Recently, WSGAN proposed a mechanism to fuse these two models. Their approach utilizes the discrete latent factors of InfoGAN to train the label model and leverages the class-dependent information of the label model to generate images of specific classes. However, the disentangled latent factors learned by InfoGAN might not necessarily be class-specific and could potentially affect the label model’s accuracy. Moreover, prediction made by the label model is often noisy in nature and can have a detrimental impact on the quality of images generated by GAN. In our work, we address these challenges by (i) implementing a noise-aware classifier using the pseudo labels generated by the label model (ii) utilizing the noise-aware classifier’s prediction to train the label model and generate class-conditional images. Additionally, we also investigate the effect of training the classifier with a subset of the dataset within a defined uncertainty budget on pseudo labels. We accomplish this by formalizing the subset selection problem as a submodular maximization objective with a knapsack constraint on the entropy of pseudo labels. We conduct experiments on multiple datasets and demonstrate the efficacy of our methods on several tasks vis-a-vis the current state-of-the-art methods. Our implementation is available at <https://github.com/kyrs/subpws-gan>

Introduction

The success of many deep learning methods is often credited to the availability of large amounts of labeled data (Deng et al. 2009; Cao et al. 2018). However, labeling necessitates manual annotation, which is a time-consuming and labor-intensive process. Distant supervision or weak supervision methods address this issue by using an external knowledge base (Hoffmann et al. 2011), pre-trained models (Bach et al. 2019), and heuristics approaches (Awasthi et al. 2020) to

generate inexpensive but potentially noisy class labels for downstream tasks. These methods use an alternative source of information to assign labels to the training data and thus avoid the need for extensive manual annotation. However, the pseudo labels generated by such methods are often noisy in nature and can have a detrimental impact on downstream performance.

Programmatic weak supervision (Zhang et al. 2022a; Ratner et al. 2016, 2020) addresses this issue by combining the prediction of multiple such noisy sources of labels often known as label functions (LFs) systematically. In PWS, a label model is trained to combine the prediction of these label functions to estimate the unobserved ground truth label for a given training sample. The generated pseudo label is then used for downstream training. For a given set of label functions, the quality of the pseudo label is often determined by the performance of the label model. Various approaches have been suggested to effectively train the label model (Zhang et al. 2022a) without knowing the underlying ground truth class labels. While approaches like (Ratner et al. 2016, 2020, 2019; Fu et al. 2020) use only the prediction of label functions to estimate the pseudo label, encoder-based label models also consider data features to estimate pseudo labels (Cachay et al. 2021; Boecking et al. 2023).

Recently, Boecking et al. (2023) proposed WSGAN, a novel training strategy to efficiently generate pseudo labels for a given image dataset. Their approach combines InfoGAN with an encoder-based label model (Cachay et al. 2021). The key idea is to train both InfoGAN (Chen et al. 2016) and label model simultaneously and leverage the disentangled latent factors learned by InfoGAN to align the label model’s predictions. This training process offers several advantages. The synchronized training of the label model and InfoGAN allows the efficient utilization of the disentangled latent factors of InfoGAN for the training of the label model. Simultaneously, the pseudo labels generated by the label model facilitate the generation of class-conditioned images, which can benefit downstream tasks as an additional form of data augmentation.

Nevertheless, the implementation of such a training process presents several challenges. The disentangled latent representation of InfoGAN is learned by maximizing the mutual information between the discrete conditional input and the generated images. However, the representations

(latent factors) learned from InfoGAN may not be class-specific. In such a scenario, aligning the prediction of the label model with such latent factors can hurt the accuracy of the label model. Furthermore, the predictions of the label models are often noisy in nature, and using the pseudo labels directly for GAN training can make the overall training process unstable, impacting the quality of images generated by InfoGAN.

In this study, we address these challenges by training a noise-aware classifier-guided conditional GAN (Odena, Olah, and Shlens 2017) for the training of label model and synthetic image generation. In contrast to conventional conditional GANs that rely on authentic ground truth labels during training, our approach operates within a weak supervision framework, where the pseudo labels generated by the label model are used to train the classifier. To handle the noise associated with the pseudo labels, the classifier is trained using a noise-aware symmetric cross entropy loss (Wang et al. 2019) in an adaptive fashion (Morero et al. 2020). Within the current setup, as training proceeds, a new set of the refined dataset and associated pseudo labels generated by the label model are used to train the classifier.

Further, a recent line of research by Lang, Vijayaraghavan, and Sontag (2022) shows that, in weak supervision, a subset of the most representative training samples yields superior performance compared to using the entire dataset for downstream training of a classifier. Motivated by this, we propose to utilize a subset of highly representative and diverse examples that exhibit minimal uncertainty with the pseudo labels. We accomplish this by incorporating a data subset selection process for weak supervision by formulating it as a submodular maximization problem (Bilmes 2022; Krause and Golovin 2014) with a knapsack constraint defined over the overall entropy of the pseudo labels of the selected samples. (Krause and Golovin 2014; Xiong, Mehta, and Singh 2019). Our contributions are summarized below:

(1) We propose a novel technique to fuse a classifier-guided conditional GAN with an encoder-based label model. Within the framework of programmatic weak supervision (PWS), this helps in efficient training of the label model and generation of class conditional images, (2) We present a novel approach for subset selection to be used in tandem with PWS by identifying the most diverse and representative samples via a submodular maximization technique, aiding in the reduction of the uncertainty associated with the training dataset, and (3) We investigate the impact on the overall performance of the label model as well as on the quality of images using the proposed subset selection scheme and compare our method with the current state-of-the-art.

Related Work

Programmatic Weak Supervision

Programmatic Weak Supervision (Ratner et al. 2016) is an efficient technique that addresses the issue of lack of ground truth labels for a downstream task. Within this framework, a subject matter expert uses different sources of noisy labeling schemes to annotate the unlabeled dataset. These sources often known as label functions (LFs) provide

partial information about the true labels and exhibit superior performance than a random model. Generally, these label functions are often approximated using an external knowledge base (Hoffmann et al. 2011), pre-trained models (Bach et al. 2019), and other similar techniques (Zhang et al. 2022a). A common practice in PWS is to define a threshold associated with each label function so that an LF can abstain from making an uncertain prediction about any example. One of the main tasks in PWS is to train a label model that combines the prediction from different LFs to estimate the underlying ground truth labels. This is done by estimating the accuracy and dependency between LFs and then using it to weight their prediction to estimate the pseudo labels. Different methods have used different approximation techniques to make informed predictions about the pseudo labels. Ratner et al. (2016) modeled the label model using a factor graph, Dawid and Skene (1979) used expectation maximization to estimate the pseudo labels, Ratner et al. (2019) used a Markov network and matrix completion techniques to recover the associated parameters, FlyingSquid (Fu et al. 2020) used Ising model to predict pseudo labels. While most of the methods estimate the pseudo labels by only considering the prediction made by the label functions, Cachay et al. (2021) proposed to use a neural network, which implicitly captures these dependencies and use data-dependent features to estimate the accuracy of the label functions. WSGAN (Boecking et al. 2023) and WSVAE (Tonolini et al. 2023) further built on top of the idea proposed by Cachay et al. (2021), and use generative models such as InfoGAN and VAE to train the label model. While WSGAN proposed the given fusion for image data, WSVAE primarily focused on improving the label model’s performance. To improve the performance of the downstream model, recent works (Zhang et al. 2022b; Yu, Ding, and Bach 2022) have proposed to use noise-aware loss in PWS. However, these losses utilize the weights of label model to improve the downstream prediction task and, hence, are not suitable for the training of the label model. Label models are generally data agnostic and only depend on the prediction made by the label functions. However, the label function used to generate noisy labels can differ for different data modalities. For instance, Varma and Ré (2018) proposed to generate pseudo labels using heuristic rules. Similarly, Chen et al. (2019) define rules on top of the prediction made by a surrogate model. Boecking et al. (2020) used the information from other modalities of data like text to estimate the labels.

Conditional GAN and Noisy Labels

Synthetic data generated using class-conditional generative models, such as Generative Adversarial Networks (GANs) (Goodfellow et al. 2014) is another approach to efficiently utilize the available dataset. However, in a weak supervision setting, such a model needs to be trained with noisy class labels. In a noisy environment, several methods have been proposed to efficiently train conditional GAN (Han et al. 2020; Kaneko and Harada 2021). However, most of these methods use a predefined set of noisy

class labels to train the GAN. Recently, Morerio et al. (2020) have shown that in a multi-domain setting, the performance of a classifier can further be improved by utilizing a conditional GAN. They proposed a unique training paradigm where the GAN and the classifier are jointly trained in an iterative fashion, in which the classifier provides refined labels for the GAN, and the GAN provides better samples to train the classifier. However, the training of such a model is limited only to a multi-domain setting, where accurate class labels for at least one domain are available.

Subset Selection for Weak Supervision

The performance of a classifier in a weakly supervised setting is influenced by the data used for its training (Lang, Vijayaraghavan, and Sontag 2022; Maheshwari et al. 2020; Mirzasoleiman, Cao, and Leskovec 2020). There exists a trade-off between the quantity of data used and the accuracy of the corresponding pseudo labels (Lang, Vijayaraghavan, and Sontag 2022). Generally, it is customary to utilize the entire dataset for training purposes. However, this approach may introduce unnecessary noise and potentially hinder the overall performance of the classifier. A possible solution is to formulate the given problem as a subset selection problem and select the most representative samples with the least entropy associated with the pseudo labels for the training of the classifier (Lang, Vijayaraghavan, and Sontag 2022). However, these methods do not provide any guarantee for the representativeness and diversity of the samples used for training and are not computationally suited for an adaptive training setting, which requires selecting multiple pseudo-labeled subsets.

Selecting an optimal subset of representative samples from the training data is an NP-hard problem. However, if the function used to measure the utility of a given set is shown to be submodular, then the set generated by maximizing the given utility using a greedy based selection scheme with constraints like cardinality (Nemhauser, Wolsey, and Fisher 1978) and knapsack (Xiong, Mehta, and Singh 2019; Leskovec et al. 2007; Krause and Golovin 2014), is guaranteed to be very close to the optimal set, thereby approximately solves the subset selection problem.

Recent works (Bilmes 2022; Xiong, Mehta, and Singh 2019; Joseph et al. 2019; Simsar et al. 2023; Wei, Iyer, and Bilmes 2015; Maheshwari et al. 2021; Sinha et al. 2020), have used submodular functions to select a subset of the entire data for efficient training of deep learning models.

Proposed Method

Problem Setting and Overview

Let there be n training data samples $\tilde{\mathcal{D}} = \{x_1, x_2, \dots, x_n\}$ drawn i.i.d from a distribution \mathcal{P}_x . The objective of our formulation is two folds: firstly, we want to infer the class labels for these samples i.e., $y \in \{0 \dots \mathcal{C}\}$ and secondly, we want to sample synthetic images whose marginal is represented as \mathcal{P}_x . In our setting, instead of observing true labels, we have access to the predictions of \mathcal{K} label functions on given training data. Each label function $\lambda_k(x_i)$ (where $k \in \{1 \dots \mathcal{K}\}$) generates a noisy prediction of the true class label for a

sample x_i . These label functions can either predict among a given set of possible class labels i.e, $\lambda_k(x_i) \in \{0 \dots \mathcal{C}\}$ or can abstain from making a prediction because of low confidence (Ratner et al. 2016). Let us refer to the data with a prediction from at least one label function as $\tilde{\mathcal{D}}_t \subseteq \tilde{\mathcal{D}}$ (non-abstained dataset). For each data sample we have an associated prediction from \mathcal{K} label functions represented as $\Lambda(x_i) = (\lambda_1(x_i), \lambda_2(x_i) \dots \lambda_{\mathcal{K}}(x_i))$. A label model (\mathcal{L}) utilizes the predictions made by these LFs to predict a final pseudo label (\hat{y}_i) for a given input sample x_i . To train our noise-aware classifier, we generate a subset of original training data $\tilde{\mathcal{D}}_0 \subseteq \tilde{\mathcal{D}}_t$ and their associated pseudo labels ($\hat{\mathcal{Y}} = \{\hat{y}_1 \dots \hat{y}_{|\tilde{\mathcal{D}}_0|}\}$) and study the impact of training the noise-aware classifier with the given subset on the overall performance of the conditional GAN and the label model.




Figure 1: Overview of the proposed fusion between conditional GAN and Programmatic Weak Supervision. In the current setup, a subset of the non-abstained samples, selected by submodular maximization under a knapsack constraint (subset selection (SS)), is used to train a noise-aware classifier ($\mathcal{P}(y|x)$). This classifier shares its network weights with the discriminator (\mathcal{D}) and label model (\mathcal{L}). The noise-aware classifier is then employed to train the generator (\mathcal{G}) and refine the label model.

Noise-Aware Classifier Training

As stated previously, we have used a classifier to facilitate the training of the conditional GAN and the alignment of the label model. Given the absence of ground truth labels, the classifier is trained using the pseudo labels generated by the label model.

The cross-entropy loss is a common choice of the loss function used to train a classifier. However, recent studies (Wang et al. 2019; Feng et al. 2021) have shown that a classifier trained only with cross-entropy loss under a noisy setting tends to exhibit prediction bias and is susceptible to overfitting. In the context of weak supervision, such a classifier can not only impact the performance of the conditional GAN but can also lead to further deterioration of the quality of pseudo labels generated by the label model.

In our study, we address this issue by training our classifier using a symmetric cross-entropy loss (l_{sce}) (Wang et al.

2019) between the pseudo label ($\hat{y} \sim \hat{\mathcal{Y}}$) generated by the label model, and the class conditional distribution modeled by the classifier ($\mathcal{P}(y|x)$). The classifier is trained on a subset of the training dataset ($\tilde{\mathcal{D}}_0$) selected by our proposed submodular maximization scheme. Symmetric cross-entropy loss is a weighted combination of standard cross-entropy loss (\mathbf{l}_{ce}) and reverse cross-entropy loss (\mathbf{l}_{rce}). While the cross-entropy loss provides a better convergence to the classifier, reverse cross-entropy loss provides robustness to different types of noise such as uniform and class conditional noise (Wang et al. 2019; Ghosh, Kumar, and Sastry 2017), thereby ensuring that the classifier does not overfit to noisy class labels. Mathematically, the symmetric cross-entropy loss (\mathbf{l}_{sce}) is described in Eq. 1.

$$\begin{aligned} \mathbf{l}_{ce}(x, \hat{y}) &= - \sum_{t=0}^c \hat{y}_t \log(\mathcal{P}(t|x)) \\ \mathbf{l}_{rce}(x, \hat{y}) &= - \sum_{t=0}^c \mathcal{P}(t|x) \log(\hat{y}_t) \\ \mathbf{l}_{sce}(x, \hat{y}) &= \alpha \mathbf{l}_{ce}(x, \hat{y}) + \beta \mathbf{l}_{rce}(x, \hat{y}) \\ \mathbf{l}_{sce}^{\tilde{\mathcal{D}}_0} &= \frac{1}{|\tilde{\mathcal{D}}_0|} \sum_{x \sim \tilde{\mathcal{D}}_0, \hat{y} \sim \hat{\mathcal{Y}}} \mathbf{l}_{sce}(x, \hat{y}) \end{aligned} \quad (1)$$

The classifier is trained in an adaptive manner (Moreno et al. 2020), where a new set of data ($\tilde{\mathcal{D}}_0$) is introduced for training after a specific number of epochs. During the training, the given examples and their associated pseudo labels remain unchanged until the next iteration. This ensures that as the accuracy and performance of the label model improve, a more refined dataset and its corresponding pseudo labels are utilized for training. Simultaneously, it prevents the model from diverging due to frequent alterations in data samples and their associated pseudo labels.

Conditional GAN

In our study, we have employed a classifier-guided generative adversarial network (GAN) to generate class conditional images (Odena, Olah, and Shlens 2017). The generator (\mathcal{G}) is fed a noise sample $z \sim N(0, I)$ and a one-hot representation of class label $c \sim \mathcal{P}(\mathcal{C})$ to sample an image ($x_g = \mathcal{G}(z, c)$) from the parameterized distribution (\mathcal{P}_θ). The discriminator (\mathcal{D}) tries to distinguish between the images sampled from real data distribution (\mathcal{P}_x) and the images sampled from parameterized distribution (\mathcal{P}_θ). For conditional image generation, the generator further maximizes the likelihood of the image (x_g) belonging to class c with respect to our noise-aware classifier. The objective function that is optimized for Conditional GAN training is as follows:

$$\begin{aligned} \mathbf{l}(\mathcal{G}, \mathcal{D}) &= \mathbb{E}_{x \sim \mathcal{P}_x} [\log(\mathcal{D}(x))] + \mathbb{E}_{x \sim \mathcal{P}_\theta} [\log(1 - \mathcal{D}(x))] \\ \min_{\mathcal{G}} \max_{\mathcal{D}} \mathbf{l}_{gan} &= \mathbf{l}(\mathcal{G}, \mathcal{D}) + \mathbb{E}_{x \sim \mathcal{G}(z, c), c \sim \mathcal{P}(\mathcal{C})} [\mathbf{l}_{ce}(x, c)] \end{aligned} \quad (2)$$

Construction of Label Models

In programmatic weak supervision, the label model accumulates the predictions made by different label functions to make an informed final estimation of the pseudo label associated with the given training sample. Generally, this is done by defining an accuracy potential function (ϕ_k) to capture the association between the prediction made by a given label function and any class label y ($\phi_k(\lambda_k, y)$). Similar to previous works (Boecking et al. 2023; Ratner et al. 2020), we have defined the accuracy potential as an indicator function, $\phi_k(\lambda_k, y) = \mathbf{1}(\lambda_k = y)$ which captures the notion whether a given label function (λ_k) has predicted class y for a given input sample. With each label function, there is an associated accuracy parameter ($\mathcal{A}_k : \mathcal{R}^N \rightarrow (0, 1)$), which determines the accuracy of its prediction. We have used sample-dependent accuracy parameters, where the accuracy of every such association is different for different input samples (Cachay et al. 2021; Boecking et al. 2023). The label model ($\mathcal{L}(\Lambda(x))_j$), defined below, captures the probability associated with class j for any input sample x .

$$\mathcal{L}(\Lambda(x))_j = \frac{\exp(\sum_{i=0}^{\mathcal{K}} \mathcal{A}(x)_i \phi_i(\lambda_i(x), j))}{\sum_{s \in \mathcal{Y}} \exp(\sum_{i=0}^{\mathcal{K}} \mathcal{A}(x)_i \phi_i(\lambda_i(x), s))} \quad \text{where } j \in \{0 \dots \mathcal{C}\} \quad (3)$$

The final pseudo label, associated with the input sample x , is defined as the class with the maximum probability.

$$\hat{y} = \arg \max_j \mathcal{L}(\Lambda(x))_j \quad (4)$$

To enhance the precision of the generated pseudo labels, the label model is trained using the probabilities generated by the classifier for the samples from the non-abstained dataset ($\tilde{\mathcal{D}}_t$). In this approach, we employ a linear layer followed by a softmax ($\hat{\mathcal{F}}$) for the alignment of probabilities between the output of the label model and the classifier ($\mathcal{P}(t|x)$) (Boecking et al. 2023), giving rise to the following loss function.

$$\mathbf{l}_{align}(x) = - \sum_{t=0}^c \mathcal{P}(t|x) \log(\hat{\mathcal{F}}(\mathcal{L}(\Lambda(x))_t)) \quad (5)$$

Since our classifier is trained iteratively by utilizing the pseudo labels generated through the label model (\hat{y} in Eq. 1), we ensure that the pseudo labels are not entirely random, especially during initial epochs. To accomplish this, we employ a new loss \mathbf{l}_{decay} that assigns equal weights to all the accuracy parameters and gradually decay this loss as the label model becomes more accurate (Boecking et al. 2023). Under this loss, the accuracy parameter for all the label functions is approximated to a value of 0.5, and the overall loss is decayed by a rate μ as the training proceeds. Mathematically, it is equivalent to generating an accuracy parameter of 0.5 for all the label functions, which is approximated by a vector $0.5 * \vec{\mathbf{1}}$ where $\vec{\mathbf{1}}$ is a vector of ones with dimension equal to the number of label functions, giving rise to the following loss function :

$$\mathbf{l}_{decay}(x) = \mu ||\mathcal{A}(x) - 0.5 * \vec{\mathbf{1}}||^2 \quad (6)$$

This process ensures that during the initial epochs, the label model approximates a majority vote, and as training progresses, it becomes more data-specific. Finally, the label model is trained using l_{lm} , by combining alignment and decay loss $l_{lm} = l_{align} + l_{decay}$.

Subset Selection for Classifier Training

To facilitate better training of our noise-aware classifier, we select a subset of data (\tilde{D}_0) from the training examples that have received at least one vote from the label functions (\tilde{D}_t). We formulate this as a submodular optimization problem as in (Bilmes 2022; Mirzasoleiman et al. 2013).

Preliminary on Submodularity

Definition 1 A function $\mathcal{F} : 2^{\mathcal{V}} \rightarrow R$ is said to be submodular for a finite set \mathcal{V} , if for every $A \subseteq B \subseteq \mathcal{V}$ and $\forall v \in \mathcal{V} \setminus B$ $\mathcal{F}(A \cup v) - \mathcal{F}(A) \geq \mathcal{F}(B \cup v) - \mathcal{F}(B)$

Further, a specific class of submodular functions is called monotone if $\forall (A, B \subseteq \mathcal{V})$ s.t. $A \subseteq B$, $\mathcal{F}(A) \leq \mathcal{F}(B)$.

A monotonous submodular function can be used to determine the utility of a given set of data for a downstream task, and an efficient subset of the data can be selected by maximizing the given function. For a cardinality constraint, the process begins with an empty set (\mathcal{S}_0), and the selection proceeds iteratively in a greedy manner, selecting the data point that maximizes the submodular function until the given budget on the size is exhausted. Mathematically, if v denotes the datapoint to be selected, then $v = \arg \max_{v \in \mathcal{V} \setminus \mathcal{S}_{n-1}} \mathcal{F}(v \cup \mathcal{S}_{n-1}) - \mathcal{F}(\mathcal{S}_{n-1})$ where $(\mathcal{S}_n = \mathcal{S}_{n-1} \cup v)$. Although this heuristic method doesn't guarantee the generation of an optimal subset due to the NP-hard nature of the problem, it can still yield a solution that is assured to be in proximity to the optimal set (Nemhauser, Wolsey, and Fisher 1978) and hence approximately solves the submodular maximization problem. Knapsack constraint (Krause and Golovin 2014) is another special case of submodular maximization, where there are different costs associated with adding every new element.

A Submodular Framework for Subset Selection In our work, we have used a generalized graphcut-based monotonous submodular function (Iyer et al. 2021; Bilmes 2022) that uses a greedy approach to select a set of representative and diverse examples to train the classifier. Further, the selection is done under a knapsack constraint where the addition of any new data sample has a cost equal to the entropy of the pseudo label associated with it, and the selection of samples is done under a budget on the overall entropy of the subset. Eq. 7 illustrates the formulation of the generalized graphcut algorithm. The function ($\mathcal{F}(S)$) represents the utility of the set (S), and the variable $s_{i,j}$ represents a kernel function that quantifies the similarity between the i^{th} and j^{th} samples in the dataset. We have used cosine similarity over the Inception features (DeVries, Drozdal, and Taylor 2020) of the images to define our kernel function. The term $(\sum_{i \in \tilde{D}_t} \sum_{j \in S} s_{ij})$ represents a cumulative score for the similarity between the samples in set (S) and the data present in the non-abstained dataset (\tilde{D}_t), capturing

the representativeness of the selected samples in comparison to (\tilde{D}_t). Additionally, the term $(\sum_{l,m \in S} s_{l,m})$ represents the similarity between the selected samples in set (S), minimizing which, leads to the selection of diverse data points. The hyperparameter γ controls the tradeoff between representativeness and diversity. Further, choosing $\gamma \geq 2$ ensures that the given function exhibits monotonicity.

$$\mathcal{F}(S) = \gamma \sum_{i \in \tilde{D}_t} \sum_{j \in S} s_{ij} - \sum_{l,m \in S} s_{l,m} \quad (7)$$

We have defined a knapsack constraint on the overall entropy of the pseudo labels generated by label model for the selected samples. Under the given constraint, every data sample in \tilde{D}_t has an associated selection cost equal to the entropy of the pseudo label, and the overall entropy of the selected samples cannot exceed a given budget (\mathcal{B}) (Eq. 8). This enforces the selection of diverse and representative samples while ensuring that uncertainty associated with the pseudo labels for the selected samples is within a limit. The complete objective of our formulation is stated as follows:

$$\begin{aligned} & \max_{S \subseteq \tilde{D}_t} \mathcal{F}(S) \\ & \text{s.t. Cost}(S) \leq \mathcal{B} \end{aligned} \quad (8)$$

The budget \mathcal{B} on the overall entropy of the pseudo labels for the selected samples is defined to be a fraction (η where $0 < \eta < 1$) of the total entropy of the pseudo labels ($Cost(\tilde{D}_t)$) for the non-abstained data set. Eq. 9 illustrates the overall constraint on the given submodular maximization problem, where \mathcal{H} is the entropy defined for a given sample (x_l), and $\mathcal{L}(\Lambda(x_l))$ is the probability distribution of the label model associated with it.

$$\begin{aligned} Cost(S) &= \sum_{x_l \in S} \mathcal{H}(x_l; \mathcal{L}(\Lambda(x_l))) \\ Cost(\tilde{D}_t) &= \sum_{x_l \in \tilde{D}_t} \mathcal{H}(x_l; \mathcal{L}(\Lambda(x_l))) \\ \mathcal{B} &= \eta * Cost(\tilde{D}_t) \end{aligned} \quad (9)$$

For the submodular maximization problem under a knapsack constraint, a naive greedy algorithm can perform arbitrarily bad as it is indifferent to the cost associated with each sample. To account for the selection cost associated with every sample, the normal greedy-based selection scheme can further be modified to create a cost-effective greedy algorithm (CEG) (Eq. 10) to select the optimal set, where $c(\cdot)$ is the cost function defined for every input sample, v_k is the data point that is selected at k^{th} iteration, $\mathcal{F}(S_{k-1})$ is the utility of the set (S_{k-1}) as per the the submodular function defined in Eq. 7 and S_k is the new set created by including the data point v_k in the set S_{k-1} . In our formulation, the function $c(\cdot)$ is set to be equal to the entropy associated with the pseudo label generated by label model ($c(v) = \mathcal{H}(v; \mathcal{L}(\Lambda(v)))$). The cost-effective greedy can be summarized as below :

$$\begin{aligned} v_k &= \arg \max_{v \in \tilde{D}_t / S_{k-1}} \frac{\mathcal{F}(S_{k-1} \cup v) - \mathcal{F}(S_{k-1})}{c(v)} \\ & \text{where } S_k = S_{k-1} \cup v_k \end{aligned} \quad (10)$$

The given algorithm is run iteratively until the total cost of the selected set is less than the predefined budget (\mathcal{B}).

Even though cost-effective greedy considers the selection cost for each sample while selecting a subset, it can still perform arbitrarily badly in some extreme scenarios (Leskovec et al. 2007). Fortunately, for the given submodular maximization objective, the cost-effective greedy algorithm can be adapted to provide a constant factor of $\frac{1}{2}(1 - \frac{1}{e})$ to the optimal solution (Prop. 1 in supplementary material) (Leskovec et al. 2007; Xiong, Mehta, and Singh 2019), hence generating a subset close to the optimal set. In fact, at least one among the final set ($\mathcal{S}_{uniform}$) selected by using uniform selection cost ($(c(v) = 1)$) and the final set (\mathcal{S}_{cost}) selected using desired cost function ($(c(v) = \mathcal{H}(v; \mathcal{L}(\Lambda(v))))$) will have a constant approximation bound with respect to the optimal solution, and the set with the maximum utility score can be used for the training of the classifier ($\tilde{\mathcal{D}}_0 = \arg \max_{\mathcal{S}} (\mathcal{F}(\mathcal{S}_{uniform}), \mathcal{F}(\mathcal{S}_{cost}))$).

Experiments

We assess the effectiveness of our method by experimenting with it on different datasets and label functions. Specifically, we use the label functions provided by the authors of WSGAN (Boecking et al. 2023). The primary experiments are done on five datasets, namely Animals with Attributes 2 (AWA2) (Xian et al. 2018), DomainNet (Peng et al. 2019), CIFAR10 (Krizhevsky, Hinton et al. 2009), MNIST (LeCun et al. 1998), FashionMNIST (Xiao, Rasul, and Vollgraf 2017) and German Traffic Sign Benchmark (GTSRB) (Stalnkamp et al. 2012). These experiments are conducted on the following four types of label functions:

Domain transfer: Experiments related to DomainNet use a domain transfer framework, wherein the label functions are trained using images from distinct source domains, like paintings, art and are subsequently used to generate weak labels for a new target domain like real-world images (Mazzetto et al. 2021). **Attribute heuristics:** Experiments related to AWA2 use attribute classifiers that are trained on a fixed set of animal images. The weak labels for the unseen dataset are generated by training a decision tree on top of the predictions made by these classifiers. **Self-supervised learning:** Experiments related to CIFAR10-B, MNIST, GTSRB, and FashionMNIST use label functions generated by finetuning a shallow MLP network over a small validation dataset. A SimCLR (Chen et al. 2020) model pretrained on ImageNet was utilized to generate features for this experiment. **Synthetic:** Experiments related to CIFAR10-A are conducted using label functions generated by synthetically corrupting the true class labels of data under given propensity and accuracy rate. Further details about the label functions can be found in WSGAN (Boecking et al. 2023).

Baselines and Evaluation Metrics

We analyze the performance of two versions of our model. Firstly, we conduct experiments by considering all the samples where the label function has provided at least one vote for the classifier’s training i.e., $\tilde{\mathcal{D}}_0 = \tilde{\mathcal{D}}_t$. The associated results related to these experiments are reported as

Ours (comp). Secondly, we conduct experiments where the dataset for classifier training is selected using the proposed submodular scheme. The results for these experiments are reported as *Ours (sub)*. We used Apricot library (Schreiber, Bilmes, and Noble 2020) to perform submodular optimization for subset selection using the lazy greedy method. All the major experiments are conducted using a DCGAN (Radford, Metz, and Chintala 2015) architecture. In the current design, the discriminator network shares weights with the classifier and accuracy parameter-based model. Further, we have provided a network ablation over styleGAN-ADA (Karras et al. 2020) architecture under a similar configuration for CIFAR10-B and high-resolution images like LSUN (Yu et al. 2015) in the supplementary material, which also includes the hyperparameters of subset selection and other implementation details.

We utilized the FID (Heusel et al. 2017) score to assess and compare the image quality produced by our approach. Similar to WSGAN, we conducted five independent iterations of our proposed technique and recorded the average FID score. Additionally, as PWS operates in a transductive setting, we follow the prior work and have reported the mean accuracy of the pseudo labels on the training data to compare the performance of the label model with other baselines. Metrics related to the F1 score, mean precision and other related scores are provided in the supplementary material. We have also evaluated the performance of the proposed method on the downstream data augmentation task.

To show the improvement in the accuracy of the classifier, we have reported the Adjusted Rank Index score (ARI) for the training epochs. We have conducted a comparative analysis with different label models, including Majority Voting (MV), MeTaL (Ratner et al. 2019), FlyingSquid (FS) (Fu et al. 2020), Snorkel (Ratner et al. 2020), hyper-label model(HLM) (Wu et al. 2023) and DawidSkene (DS) (Dawid and Skene 1979). To facilitate this comparison, we employed the label model codebase from Wrench (Zhang et al. 2021) and utilized the official codebase provided by WSGAN and hyper-label model. If any conflicts emerged during the implementation process due to version disparities, we referred to the numerical values reported in the paper. When information was unavailable, we omitted to report those particular numbers.

Results on Classifier Performance

Figure-2 shows the Adjusted Rank Index between the prediction made by our classifier and the underlying ground truth label. We have further compared it with the prediction made by WSGAN. The results indicate that the noise-aware classifier is more effective than WSGAN in learning the underlying ground truth labels.

Results on Label Model Performance

Table-1 shows the comparison of the mean posterior accuracy of our proposed method with other baselines. Compared to WSGAN, our proposed method gives a 3.0% improvement in AWA2 dataset and around 1.8% improvement in CIFAR10-A. Among the subset selection and the non-

Dataset	MV	MeTaL	FS	Snorkel	DS	HLM	WSGAN	Ours (comp)	Ours (sub)
AWA2	0.627	0.622	0.621	0.617	-	0.651	0.670	0.701	0.698
CIFAR10-A	0.831	0.804	0.800	0.804	0.850	0.842	0.872	0.890	0.890
CIFAR10-B	0.719	0.708	0.708	0.709	0.677	0.726	0.729	0.740	0.740
DomainNet	0.595	0.486	0.635	0.507	0.658	0.646	0.643	0.650	0.653
FashionMNIST	0.735	0.730	0.734	0.725	0.717	0.735	0.744	0.754	0.754
GTSRB	0.816	0.619	0.815	0.671	0.814	-	0.825	0.829	0.828
MNIST	0.778	0.758	0.773	0.759	0.729	0.781	0.816	0.818	0.819

Table 1: Comparison between the average posterior accuracy of the label models for samples with at least one vote from the label function. The best performer across methods is denoted as bold.




Figure 2: Comparison of Adjusted Rank Index(ARI) of proposed methods and WSGAN.

Dataset	WSGAN	Ours (comp)	Ours (sub)
AWA2	36.00	25.92	25.23
CIFAR10-A	22.71	18.00	17.50
CIFAR10-B	24.41	18.72	18.72
DomainNet	44.35	37.42	36.82
GTSRB	73.96	62.55	59.09
FashionMNIST	15.94	11.78	11.81
MNIST	5.35	4.00	4.00

Table 2: Comparison of image quality (mean FID) score of the proposed method with WSGAN.

Dataset	W-S	W-L	Os-L	Os-L	Oc-C	Os-C
AWA2	0.99	0.99	2.58	1.59	1.59	1.79
CIFAR10-A	1.38	0.46	0.50	0.76	1.74	2.16
CIFAR10-B	0.66	0.36	0.42	0.82	1.22	1.56
DomainNet	0.84	0.84	1.88	3.14	3.77	1.88
FashionMNIST	0.36	2.72	2.74	2.92	0.74	0.64
GTSRB	0.05	0.32	1.06	0.60	0.81	1.27
MNIST	0.30	0.08	0.34	0.22	0.52	0.38

Table 3: Percentage increase in the test accuracy using synthetic data augmentation. The WSGAN model with synthetic pseudo labels (W-S) and label function-based pseudo labels (W-L) is compared with our complete data setting (Oc) and subset selection setting (Os) using conditional latent vectors (L) and classifier-based (C) pseudo labels.

abstained data, we found a marginal difference in the performance of the label model.

Results on Image Generation

The results in Table-2 highlight the average FID scores achieved by our methods compared to those of WSGAN. On average, our method achieves 7.08 point FID improvement compared to WSGAN. Additionally, the subset selection scheme for the classifier further improves the quality of generated images. This distinction is particularly notable for

the GTSRB dataset, with a 3.46-point FID enhancement.

Results on Data Augmentation

To examine the improvement in test accuracy for the data augmentation task, we trained a ResNet-18 (He et al. 2016) model using around 1000 synthetic images generated by the GAN, in addition to the original training dataset. For pseudo labels, we have considered two settings. Firstly, we use the conditional discrete latent vector (L) used for image generation. Secondly, we use the output of the classifier to generate the associated class label for the images (C). The results of the experiment are reported in Table-3. Comparatively, we found that our proposed method generates better results than the baseline. Among the subset selection scheme and the non-abstained data, we found the performance of the subset selection scheme to be better for DomainNet, FashionMNIST, CIFAR10-B, CIFAR10-A for conditional discrete latent vector (L), and AWA2, CIFAR10-A, GTSRB, and CIFAR10-B for classifier based pseudo labels (C).

Conclusion

In this work, we introduced a new framework to fuse the training of conditional GAN and label models using a noise-aware classifier. We have further investigated the impact of training the classifier using a subset of representative and diverse samples with limited uncertainty associated with the pseudo labels. Our empirical results indicate an improvement in the quality of images and the accuracy of pseudo labels compared to existing state-of-the-art techniques. We further demonstrate that employing a subset selection strategy helps in improving the image quality of GAN, and a small training set can generate comparable performance of label model. However, it is worth noting that the subset selection approach has limitations tied to the size of the non-abstained sample set, which can result in computational challenges for very large datasets. In future work, we aim to extend our approach to bigger datasets.

Acknowledgments

This work was supported (in part for setting up the GPU compute) by the Indian Institute of Science through a start-up grant. Prathosh is supported by the Infosys Foundation Young Investigator Award.

References

- Awasthi, A.; Ghosh, S.; Goyal, R.; and Sarawagi, S. 2020. Learning from rules generalizing labeled exemplars. *arXiv preprint arXiv:2004.06025*.
- Bach, S. H.; Rodriguez, D.; Liu, Y.; Luo, C.; Shao, H.; Xia, C.; Sen, S.; Ratner, A.; Hancock, B.; Alborzi, H.; et al. 2019. Snorkel drybell: A case study in deploying weak supervision at industrial scale. In *Proceedings of the 2019 International Conference on Management of Data*, 362–375.
- Bilmes, J. 2022. Submodularity in machine learning and artificial intelligence. *arXiv preprint arXiv:2202.00132*.
- Boecking, B.; Neiswanger, W.; Xing, E.; and Dubrawski, A. 2020. Interactive weak supervision: Learning useful heuristics for data labeling. *arXiv preprint arXiv:2012.06046*.
- Boecking, B.; Roberts, N.; Neiswanger, W.; Ermon, S.; Sala, F.; and Dubrawski, A. 2023. Generative Modeling Helps Weak Supervision (and Vice Versa). In *ICLR*.
- Cachay; Rühling, S.; Boecking, B.; and Dubrawski, A. 2021. End-to-end weak supervision. *NeurIPS*, 34: 1845–1857.
- Cao; Jiajiong, Y.; Li, Z.; and Zhang. 2018. Celeb-500k: A large training dataset for face recognition. In *2018 25th IEEE International Conference on Image Processing (ICIP)*, 2406–2410. IEEE.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *ICML*, 1597–1607. PMLR.
- Chen, V.; Wu, S.; Ratner, A. J.; Weng, J.; and Ré, C. 2019. Slice-based learning: A programming model for residual learning in critical data slices. *NeurIPS*, 32.
- Chen, X.; Duan, Y.; Houthoofd, R.; Schulman, J.; Sutskever, I.; and Abbeel, P. 2016. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. *NeurIPS*, 29.
- Dawid, A. P.; and Skene, A. M. 1979. Maximum likelihood estimation of observer error-rates using the EM algorithm. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 28(1): 20–28.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. IEEE.
- DeVries, T.; Drozdal, M.; and Taylor, G. W. 2020. Instance selection for gans. *NeurIPS*, 33: 13285–13296.
- Feng, L.; Shu, S.; Lin, Z.; Lv, F.; Li, L.; and An, B. 2021. Can cross entropy loss be robust to label noise? In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, 2206–2212.
- Fu, D.; Chen, M.; Sala, F.; Hooper, S.; Fatahalian, K.; and Ré, C. 2020. Fast and three-rious: Speeding up weak supervision with triplet methods. In *ICML*, 3280–3291. PMLR.
- Ghosh, A.; Kumar, H.; and Sastry, P. S. 2017. Robust loss functions under label noise for deep neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *NeurIPS*, 27.
- Han, L.; Gao, R.; Kim, M.; Tao, X.; Liu, B.; and Metaxas, D. 2020. Robust conditional GAN from uncertainty-aware pairwise comparisons. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 10909–10916.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *NeurIPS*, volume 30. Curran Associates, Inc.
- Hoffmann, R.; Zhang, C.; Ling, X.; Zettlemoyer, L.; and Weld, D. S. 2011. Knowledge-based weak supervision for information extraction of overlapping relations. In *Proceedings of the 49th annual meeting of the association for computational linguistics: human language technologies*, 541–550.
- Iyer, R.; Khargoankar, N.; Bilmes, J.; and Asanani, H. 2021. Submodular combinatorial information measures with applications in machine learning. In *Algorithmic Learning Theory*, 722–754. PMLR.
- Joseph, K.; Singh, K.; Balasubramanian, V. N.; et al. 2019. Submodular batch selection for training deep neural networks. *arXiv preprint arXiv:1906.08771*.
- Kaneko, T.; and Harada, T. 2021. Blur, noise, and compression robust generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13579–13589.
- Karras, T.; Aittala, M.; Hellsten, J.; Laine, S.; Lehtinen, J.; and Aila, T. 2020. Training generative adversarial networks with limited data. *NeurIPS*, 33: 12104–12114.
- Krause, A.; and Golovin, D. 2014. Submodular function maximization. *Tractability*, 3(71-104): 3.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Lang, H.; Vijayaraghavan, A.; and Sontag, D. 2022. Training subset selection for weak supervision. *NeurIPS*.
- LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11): 2278–2324.
- Leskovec, J.; Krause, A.; Guestrin, C.; Faloutsos, C.; VanBriesen, J.; and Glance, N. 2007. Cost-effective outbreak detection in networks. In *Proceedings of the 13th ACM*

- SIGKDD international conference on Knowledge discovery and data mining*, 420–429.
- Maheshwari, A.; Chatterjee, O.; Killamsetty, K.; Ramakrishnan, G.; and Iyer, R. 2020. Semi-supervised data programming with subset selection. *arXiv preprint arXiv:2008.09887*.
- Maheshwari, A.; Killamsetty, K.; Ramakrishnan, G.; Iyer, R.; Danilevsky, M.; and Popa, L. 2021. Learning to robustly aggregate labeling functions for semi-supervised data programming. *arXiv preprint arXiv:2109.11410*.
- Mazzetto, A.; Cousins, C.; Sam, D.; Bach, S. H.; and Upfal, E. 2021. Adversarial multi class learning under weak supervision with performance guarantees. In *International Conference on Machine Learning*, 7534–7543. PMLR.
- Mirzasoleiman, B.; Cao, K.; and Leskovec, J. 2020. Core-sets for robust training of deep neural networks against noisy labels. *NeurIPS*, 33: 11465–11477.
- Mirzasoleiman, B.; Karbasi, A.; Sarkar, R.; and Krause, A. 2013. Distributed submodular maximization: Identifying representative elements in massive data. *NeurIPS*, 26.
- Morerio, P.; Volpi, R.; Ragonesi, R.; and Murino, V. 2020. Generative pseudo-label refinement for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 3130–3139.
- Nemhauser, G. L.; Wolsey, L. A.; and Fisher, M. L. 1978. An analysis of approximations for maximizing submodular set functions—I. *Mathematical programming*, 14: 265–294.
- Odena, A.; Olah, C.; and Shlens, J. 2017. Conditional image synthesis with auxiliary classifier gans. In *ICML*, 2642–2651. PMLR.
- Peng, X.; Bai, Q.; Xia, X.; Huang, Z.; Saenko, K.; and Wang, B. 2019. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 1406–1415.
- Radford, A.; Metz, L.; and Chintala, S. 2015. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*.
- Ratner, A.; Bach, S. H.; Ehrenberg, H.; Fries, J.; Wu, S.; and Ré, C. 2020. Snorkel: Rapid training data creation with weak supervision. *The VLDB Journal*, 29(2-3): 709–730.
- Ratner, A.; Hancock, B.; Dunnmon, J.; Sala, F.; Pandey, S.; and Ré, C. 2019. Training complex models with multi-task weak supervision. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 4763–4771.
- Ratner, A. J.; De Sa, C. M.; Wu, S.; Selsam, D.; and Ré, C. 2016. Data programming: Creating large training sets, quickly. *NeurIPS*, 29.
- Schreiber, J.; Bilmes, J.; and Noble, W. S. 2020. apricot: Submodular selection for data summarization in Python. *The Journal of Machine Learning Research*, 21(1): 6474–6479.
- Simsar, E.; Kocasari, U.; Er, E. G.; and Yanardag, P. 2023. Fantastic Style Channels and Where to Find Them: A Submodular Framework for Discovering Diverse Directions in GANs. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 4731–4740.
- Sinha, S.; Zhang, H.; Goyal, A.; Bengio, Y.; Larochelle, H.; and Odena, A. 2020. Small-gan: Speeding up gan training using core-sets. In *ICML*, 9005–9015. PMLR.
- Stallkamp, J.; Schlipsing, M.; Salmen, J.; and Igel, C. 2012. Man vs. computer: Benchmarking machine learning algorithms for traffic sign recognition. *Neural networks*, 32: 323–332.
- Tonolini, F.; Aletras, N.; Jiao, Y.; and Kazai, G. 2023. Robust Weak Supervision with Variational Auto-Encoders.
- Varma, P.; and Ré, C. 2018. Snuba: Automating weak supervision to label training data. In *Proceedings of the VLDB Endowment. International Conference on Very Large Data Bases*, volume 12, 223. NIH Public Access.
- Wang, Y.; Ma, X.; Chen, Z.; Luo, Y.; Yi, J.; and Bailey, J. 2019. Symmetric cross entropy for robust learning with noisy labels. In *Proceedings of the IEEE/CVF international conference on computer vision*, 322–330.
- Wei, K.; Iyer, R.; and Bilmes, J. 2015. Submodularity in data subset selection and active learning. In *ICML*, 1954–1963. PMLR.
- Wu, R.; Chen, S.-E.; Zhang, J.; and Chu, X. 2023. Learning Hyper Label Model for Programmatic Weak Supervision. In *ICLR*.
- Xian, Y.; Lampert, C. H.; Schiele, B.; and Akata, Z. 2018. Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly. *IEEE transactions on pattern analysis and machine intelligence*, 41(9): 2251–2265.
- Xiao, H.; Rasul, K.; and Vollgraf, R. 2017. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*.
- Xiong, Y.; Mehta, R.; and Singh, V. 2019. Resource constrained neural network architecture search: Will a submodularity assumption help? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1901–1910.
- Yu, F.; Seff, A.; Zhang, Y.; Song, S.; Funkhouser, T.; and Xiao, J. 2015. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*.
- Yu, P.; Ding, T.; and Bach, S. H. 2022. Learning from Multiple Noisy Partial Labelers. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, 11072–11095. PMLR.
- Zhang, J.; Hsieh, C.-Y.; Yu, Y.; Zhang, C.; and Ratner, A. 2022a. A survey on programmatic weak supervision. *arXiv preprint arXiv:2202.05433*.
- Zhang, J.; Wang, H.; Hsieh, C.-Y.; and Ratner, A. J. 2022b. Understanding Programmatic Weak Supervision via Source-aware Influence Function. In *NeurIPS*, volume 35, 2862–2875.
- Zhang, J.; Yu, Y.; Li, Y.; Wang, Y.; Yang, Y.; Yang, M.; and Ratner, A. 2021. Wrench: A comprehensive benchmark for weak supervision. *arXiv preprint arXiv:2109.11377*.