

Iterative Data Programming for Expanding Text Classification Corpora

Neil Mallinar,^{†*} Abhishek Shah, Tin Kam Ho, Rajendra Ugrani, Ayush Gupta

IBM Watson, New York, NY, [†]Pryon Research, New York, NY

Abstract

Real-world text classification tasks often require many labeled training examples that are expensive to obtain. Recent advancements in machine teaching, specifically the data programming paradigm, facilitate the creation of training data sets quickly via a general framework for building weak models, also known as labeling functions, and denoising them through ensemble learning techniques. We present a fast, simple data programming method for augmenting text data sets by generating neighborhood-based weak models with minimal supervision. Furthermore, our method employs an iterative procedure to identify sparsely distributed examples from large volumes of unlabeled data. The iterative data programming techniques improve newer weak models as more labeled data is confirmed with human-in-loop. We show empirical results on sentence classification tasks, including those from a task of improving intent recognition in conversational agents.

Introduction

A chat bot that automates handling of common customer service requests often uses a text classifier to identify the intent in incoming requests. Many machine learning techniques are suitable for building these classifiers. The biggest bottleneck in applying such classifiers is the availability of labeled training data.

Labeling vast amounts of text data is expensive, involving manual annotation by domain experts with complex background knowledge. It is even more challenging when the classes are highly skewed, which renders a sequential process to look for each class largely ineffective, as examples of the minority class are buried in much larger amounts of data of other classes. Deciding between multiple classes for each example, on the other hand, is also stressful as the annotator has to recall the definitions of all classes all the time.

Much work has been done to make the labeling effort easier by expanding the manually provided labels in some automated way. Commonly used data expansion methods include active learning (Karlos et al. 2017), self-training (Guzmán-Cabrera et al. 2009), and co-training (Du, Ling,

and Zhou 2010). Active learning utilizes a query-selection strategy and an oracle (e.g., a domain expert) to select an important subset of the unlabeled examples to send to the oracle for labeling. The learner is retrained with the newly obtained labels and the process repeats. Without using an oracle, self-training and co-training leverage the more confident decisions of classifiers trained with available examples to provide additional examples to retrain the classifiers. In self-training, an automatic learner is first trained on any available labeled data and evaluated on the unlabeled data. Any unlabeled examples that the learner can classify with confidence passing a threshold are labeled by the decision and used in retraining, and the process repeats. In co-training, multiple classifiers are trained on different views of the labeled data, and the decision with highest confidence on the unlabeled data from the consensus is added to the labeled set.

Semi-supervised learning methods seek to utilize available labels to maximum effect. Examples are techniques like Expectation-Maximization (Nigam et al. 2011) to assist class discovery, and training generative models through data programming (Ratner et al. 2016; Nashaat et al. 2018) with large amounts of unlabeled text data. In this paper, we aim to explore the latter, data programming, in an iterative procedure for the purpose of expanding labeled data efficiently.

We design an iterative expansion process that does not utilize the downstream classifier (for which the training data is built) in any way, so as to stay generally applicable to many choices for the classifier, and to avoid computationally expensive training of that classifier over many iterations of expansion. We experiment with the process using public and private text corpora and confirm that the approach is effective and superior to simpler iterative expansion techniques.

Data Programming, Weak Supervision, and Machine Teaching

Previous work has shown that interactive procedures aid creation and bootstrapping of new class-labeled data (Williams et al. 2015). They show that domain experts, who are not necessarily knowledgeable about machine learning systems, are able to effectively build classes through a combination of model definition, labeling, active learning, and more. Domain experts are often expensive and their time is valuable.

^{*}Work done while at IBM Watson

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Data: Labeled training data TR ; unlabeled corpus U ; selection model S ; labeler L ; termination criteria $t = 30$; and batch size $b = 10$

Result: Augmented labeled training data, TR'

$done = \mathbf{False}$;

while *not done* **do**

$R = \text{sort}_S(U)$; $found = \mathbf{False}$;

for $i = 0$; $i < \min(t, |R|)$; $i = i + b$ **do**

$P = \text{label}_L(R[i : i + b - 1])$;

if $|P| > 0$;

then

$N = R[i : i + b - 1] - P$;

$TR = TR + P + N$;

$U = U - P - N$; $found = \mathbf{True}$;

break ;

end

end

if not found then

$done = \mathbf{True}$;

end

end

▷ retrieve rankings of elements of U

▷ label small batches from R

▷ if any positive examples found

▷ get negative examples

▷ expand training data

▷ remove newly labeled data from U

▷ break out of inner for-loop

▷ terminate if no positive examples found

Algorithm 1: In the expansion procedure, a selection model extracts small, rank-ordered batches of data from an unlabeled corpus that are relevant to the current labeled data. These batches are sent, in order, to the labeler for annotation. Once positive labels are found, the selection model is updated with the new data and the procedure is repeated until a termination condition is met. Here, termination is when no more positive labels can be found within a maximum number of batches.

Guided learning approaches aim to lessen the amount of data to be labeled, by using search to select relevant examples for labeling, and are shown to be more effective than random labeling and active learning techniques in settings where data exhibits high skew, as ours is given the iterative nature of the classifier building process. (Attenberg and Provost 2010).

In this work, we employ weak supervision techniques to select batches of examples for labeling. Such techniques have recently been reintroduced through the data programming framework (Ratner et al. 2016; 2018), by the notion of “labeling functions” that generate weak labels. These functions can encode heuristics or can simply be noisy processes for labeling data. Often such heuristics can be given by domain experts, using significantly less of their time. Once many labeling functions are obtained, the next step is to denoise their decisions to obtain a singular label for each example. The work on Snorkel (Ratner et al. 2018) uses a generative model to denoise the ensemble of labeling functions. After training the generative model, probabilistic labels are obtained per example and are used to train a downstream learner which generalizes the learned heuristics and apply them to any unseen examples.

These ideas have been used to expand intent training data for conversational agents (Mallinar et al. 2019) and to quickly label large industrial data sets (Nashaat et al. 2018). We build on these works by using an iterative procedure to automatically construct weak models and show that the Snorkel generative model can select more relevant sentences for labeling than search-based methods alone.

Our work is also connected to the recent machine teaching literature, following definitions and processes provided by (Simard, Amershi, and others 2017). Machine teaching

is defined to be a procedure that can make a “teacher” more efficient at constructing machine learning models. A teacher is defined as somebody who can transfer knowledge (a concept) to a machine learning system (referred to as “student” or “learner”). Modes of knowledge transfer include selection, where a teacher can find useful examples that exhibit characteristics of a concept, and labeling, where a teacher can provide annotations to directly relate concepts to examples. Earlier definitions of machine teaching are given by (Zhu 2015), where the task is for a teacher that knows the optimal decision boundary between classes in feature space to provide as few examples as possible to the student learner such that the learner can build the same decision boundary. Recently, work has been done in iterative machine teaching (Liu et al. 2017), showing that employing a continuous loop allows the teacher to react to the student as it learns. Summarizing over all of these definitions, the overarching goal of machine teaching is to avoid expensive searches in data space by making use of teacher models that already know the decision boundary, allowing the teacher to show the student as few examples as possible.

The closest existing literature to machine teaching is in the active learning sub-field. Active learning operates in a closed-loop where a learning algorithm is trained on any labeled data available. That learner then evaluates a large unlabeled corpus and a query-selection strategy is employed to find a small set of examples that will be most useful, if labeled, to add to the training data. An oracle, usually a human annotator or crowd-sourced worker, is asked to assign true labels to the selected set of examples.

We distinguish our work from active learning frameworks as in (Settles 2009) by that we never access or query the

downstream learner during our procedure. Often such learners require a decent amount of data and tuning on validation sets to be effective models. In our setting, the amount of labeled data provided at the start is very limited (< 20 examples per class) to the point that even a validation set is impossible to extract. Effectively, we are asking the question: can a “teacher” find a “curriculum” for a concept that is versatile enough for an unknown learner to understand?

Proposed Methods

We present a procedure for expanding a small set of labeled training data with examples from a larger unlabeled corpus. Our procedure repeats a basic step that consists of leveraging the labeled examples to rank and select the unlabeled examples, and then consulting an oracle for labeling the selected ones. In each step the selection is limited in size to minimize consumption of the oracle resources. This procedure continues until a termination criteria is reached. The resultant labeled data set is taken to train a downstream text classifier for final evaluation and eventual use in the application.

Iterative Expansion Procedure


We define the main expansion process, provided in Algorithm 1, in the context of binary label data (positive or negative for a specific class). For multi-class data, we simply apply the same expansion process independently for each class and merge the resultant positive labeled sets to get our final expanded training set to avoid extreme data skew that might be caused by intermediate steps prior to the completion of expansion of all the classes.

The algorithm takes as input: an initial training set, TR , composed of a small number of examples per class; an unlabeled corpus, U , to search within for relevant examples; a selection model, S , used to find and evaluate relevancy of unlabeled data with respect to labeled data; and a labeler (oracle), L , that can provide true labels quickly. The labeler can be expensive to evaluate in practice (often a domain expert), so we include a batch size setting that restrict the per-iteration labeling effort, and a termination criteria to determine when to stop expanding the data.


At the beginning of each iteration we retrieve a rank-ordered set of examples from the unlabeled corpus that are relevant to the current set of labeled training data via the selection model. We then iterate, in-order, over these relevant examples in batches of size b which are sent to the labeler for annotation. When a batch is found with positive examples we augment the training data set with the newly found positive and negative examples, where the negative set is the complement of the positive set within the batch.

This is repeated until a termination criteria, parameterized by t , is met. We define t to be the maximum number of consecutive examples we are willing to send to the labeler, in each iteration, without receiving any positive labels back. We consider the data to be fully expanded when we reach this point. In this work, we use $t = 30$ and $b = 10$ to simulate a user looking through three pages of relevant examples with ten examples per page before quitting.

In the case of multi-class data, we convert the data to binary labels per class and use Algorithm 1 to expand each



(a) Reddit: graphics_tablet



(b) Reddit: home_security

Figure 1: Plotting number of examples added to training data per iteration on Y-axis vs. number of iterations of expansion on X-axis for two classes from the Reddit dataset.

of these binarized data sets independently. While we track negative expansions relative to each set, we simply take the union of the positive expansions to be the final augmented multi-class training data.

Ranking and Selection Methods

We consider two methods for ranking and selecting the unlabeled examples. The first method uses search-based methods only, while the latter employs data programming.

Iterative Expansion by Search (I-MLT) Search methods have been used to extract data when the class distribution of the unlabeled data is highly skewed (Attenberg and Provost 2010). In expansion by search, we use a search engine as a selection model to find examples relevant to the input examples. Here we use the “More Like This” (MLT) feature of the search engine *Elasticsearch*¹ to expand batches of examples and rank them based on the returned score. As such, this selection model only uses positively labeled data for expansion and does not make use of negative labels.

With expansion by search, the unlabeled corpus is ingested into the search engine once and continuously used throughout the process. This allows to perform selection efficiently, compared to re-training a statistical classifier at every iteration.

Iterative Data Programming with Expansion by Search (I-DP) As introduced earlier, data programming proposes to write simple “labeling functions” that are fast to evaluate and can coarsely assign labels to, potentially large, subsets of unlabeled data. We refer to such labeling functions as

¹<https://www.elastic.co/guide/en/elasticsearch/reference/6.4/search-more-like-this.html>

weak models hereafter, and their label assignment as weak labels (not to be confused with the labels obtained from the labeler L in Algorithm 1, which may also be referred to as strict labels). Such weak models generally do not have access to many ground truth labels and make their assignments based on heuristics encoded by human domain experts, crowdsourced workers, or external knowledge bases. We note that while this work uses lexical matching to build neighborhoods, the way we construct weak models can be generalized to any reasonable way to define a subset in a metric space. We leave the exploration of different ways to construct weak models to future work.

Given a collection of n weak models that collectively assign weak labels to m examples, a label matrix, $\Lambda \in \mathbb{R}^{m \times n}$, is constructed out of all of the assignments. Where example i is not labeled by weak model j , we set $\Lambda_{i,j} = 0$ to represent abstention from labeling the example by the given model. As such, Λ can have varied sparsity depending on how large the coverage of the constituent weak models is. As described in (Ratner et al. 2018), when Λ is very sparse or very dense a simple majority vote can be used to denoise the weak labels assigned between the various weak models.

In the medium density setting, a probabilistic generative model $p_\theta(\Lambda, Y)$ is trained over Λ and the true labels, Y , of the examples. It is shown that in the medium density setting, such a model can better denoise the assignments between weak models than a simple majority vote. When Y is not available, it is treated as a latent variable in the model and optimization is done over the negative log marginal likelihood, $-\log \sum_Y p_\theta(\Lambda, Y)$. This is done via a procedure that alternates between Gibbs sampling steps and stochastic gradient descent steps.

Furthermore, dependencies between weak models are learned through a structure learning process whereby $p_\theta(\Lambda, Y)$ is treated as a factor graph. Additional factors are learned and added to the model to represent common dependencies between weak models (Bach et al. 2017).

We construct weak models with two independent MLT expansions on the newly labeled positive and negative examples obtained at each iteration. Each search expansion brings a set of ranked examples relevant to the input data. Here, the ranking scores from the search engine are discarded and each set of results is weakly labeled as positive or negative, depending on whether the expansion was on positive or negative labeled data, respectively. Thus, at the k^{th} iteration in Algorithm 1 we are able to construct $2k$ weak models (half of which are positive expansions, and the other half negative). This automatic construction of weak models eliminates the need for domain experts to encode relevant heuristics about the classes, or the need for many crowdsourced workers. Instead, a single labeler can quickly label a small number of examples over a few iterations and generate a larger label matrix, Λ .

We further truncate the expanded ranked examples list in each weak model to the top 50 to ensure we retrieve high precision subsets from the unlabeled data and to control the level of noise introduced in the weak labeling process. After training the generative model, we finally retrieve a list of computed probabilistic labels over examples in our un-

labeled corpus (where we default to a score of 0 for those examples that are not covered by at least one weak model). These probabilistic labels are used to rank examples for selection in future iterations of Algorithm 1.

We employ this in our procedure to create a hybrid expansion method where rankings obtained both by search expansions and data programming expansions at each iteration are merged into a final ranked examples lists via round-robin merging. This is done because training the probabilistic model incurs time and computational overhead while using search does not, so we introduce a re-train frequency parameter, f , to control how often re-training of the probabilistic model occurs.

The most recent ranked examples obtained from the generative model are merged with the latest search expansion rankings until the next time the probabilistic model is re-trained (e.g. in the f subsequent iterations of the algorithm after the generative model is trained). This merging scheme ensures that at each iteration there is an expansion method that reflects the latest qualities of the labeled data (search), as well as an expansion method that produces stronger, less noisy rankings at the cost of being slightly outdated relative to the current set of labeled data (data programming).

The effect of using $f > 1$ is two-fold, in that an iteration may not provide enough newly labeled data to be worth re-training the generative model on. As such, we continue to expand via search until we have retrieved enough examples to affect the rankings list in a more substantial way. In this work, we set $f = 3$, meaning the model is only re-trained when there are at least 6 new weak models to evaluate.

Experiments and Results

We use three sentence classification sets for experiments. From each dataset, we keep only the classes with over 110 examples, and merge the rest into an “others” class. Our experience indicates that most users can afford to hand-write or hand-select a very small set of examples per class. So we hand-select such examples as the initial training set in our experiments. The remaining training examples are stripped of labels and used as the unlabeled corpus ingested into the search engine.

Reddit: Reddit self-post classification dataset² has 1,013 sub-reddits from which we choose titles as examples from 81 sub-reddits related to the technology domain. We preprocess the titles (and subreddit names as labels) by removing newlines, tabs, special characters, and extraneous spaces and any duplicates. The final Reddit dataset consists of 810 labeled titles to start with, 63,990 unlabeled titles and 16,200 titles as test set.

Customer Service Inquiries: We construct two data sets of utterances, one in English and one in Japanese, from privately sourced customer service dialogues. They cover topics from general customer service requests to domain-specific questions in several industries like banking, insurance, and utilities. This data primarily consists of short questions and demands mapped to intents (categories) defined by

²The dataset can be downloaded from <https://www.kaggle.com/mswarbrickjones/reddit-selfposts/>

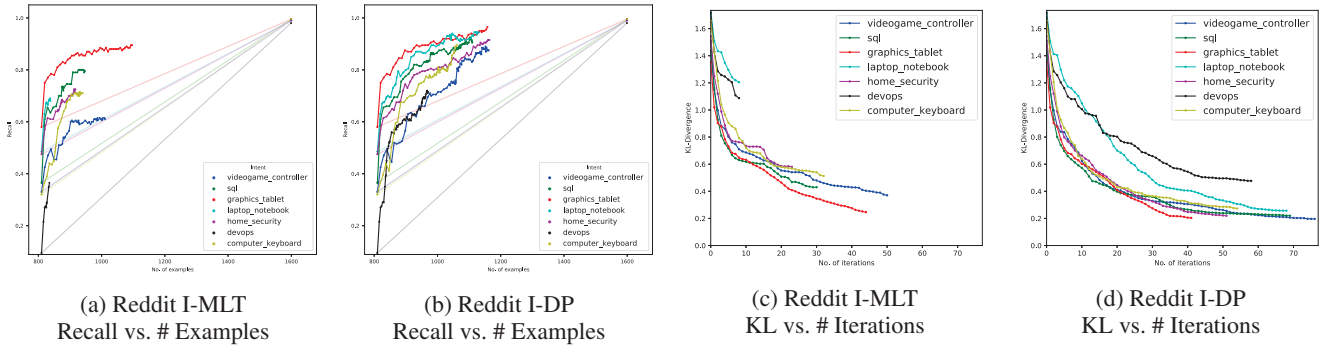


Figure 2: Comparison plots for 7 classes sampled from the Reddit data set. (a,b): Expansion paths plotting recall on test set vs. number of examples added. The points at the far right of each curve mark the maximum recall if all relevant training samples are found and included. (c,d): KL-Divergence vs. number of iterations of expansion. We measure KL-Divergence between the vocabulary distribution at each iteration to the optimal distribution per-class.

subject-matter experts. The English data set, CS-En, is split into an unlabeled corpus of 15,236 examples, initially labeled set of size 1,998, and test set of size 15,177 spanning 97 categories. The Japanese data set, CS-Ja, is split into an unlabeled corpus of 1,709 examples, initially labeled set of size 733, and test set of size 1,691 covering 38 categories that are a subset of those in English.

Experimental Procedure

To evaluate the effectiveness of our expansion procedure under different selection models, we start with the initial labeled set and expand each class independently as described in Algorithm 1. After the expansion procedure finishes for each class, we combine all of the fully expanded classes together to create the final labeled training data set.

Finally, we train a text classifier³ on this labeled set and evaluate accuracy on the held-out test set. This accuracy is called the terminal accuracy for the given method used in expansion. We compare the terminal accuracy to the optimal accuracy, which is obtained by training the same classifier on using the ground truth for the entire corpus from which the examples are selected, and evaluating on the test set. For further analysis, we measure the KL-Divergence between the vocabulary distribution (at each iteration) of the class being expanded to the optimal vocabulary distribution of that class (as computed from all examples of that class from the corpus).

Results and Discussion

We present our main results for I-MLT and I-DP in Table 1, along with comparisons to a baseline system that uses a random ranker as selection “model”. The accuracy results there reflect the quality of a classifier trained on labeled examples obtained by each method, and are evaluated on the held out examples. Another assessment of the quality of the obtained training set is given by the estimated KL divergence of the word distributions in the training set and the word distribution in the optimal set. It is clear in all cases that the I-DP

Methods	Reddit	CS-En	CS-Ja
Initial	0.3035	0.7434	0.5523
Optimal	0.6219	0.8725	0.7776
Random	0.3070	0.7541	0.6375
I-MLT	0.4081	0.7806	0.6606
I-DP	0.4660	0.8081	0.6972

Table 1: Accuracy results for Random, I-MLT, and I-DP on Reddit, English and Japanese Customer Service test sets. We compare the initial and optimal train sets to the terminal train sets, using a downstream classifier for evaluation.

Method	Reddit	CS-En	CS-Ja
Initial	810	1,998	733
Optimal	64,715	17,234	2,442
Random	842	2,030	974
I-MLT	6,934	6,975	1,517
I-DP	16,006	10,125	1,769

Table 2: Constructed training data size comparisons on our benchmark data for each expansion method. We also present the counts for the initial labeled data set and optimal set.

approach is able to construct a training set that performs best on the held-out test sets.

Figure 1 shows how many examples are able to be labeled at each iteration of I-MLT and I-DP for two classes sampled from the Reddit data. We see that I-DP is able to consistently find more relevant data for labeling and for more iterations before terminating. With these results, we built a tool to assist human annotators to find examples in chat logs to build up intent training data. Early users are finding it very helpful.

Figure 2 shows the expansion paths, per-class, of seven classes sampled from the Reddit data. In each plot, the terminal point to the far right of the opaque lines depict the optimal recall of each class, being 1.0 as all training data

³<https://www.ibm.com/cloud/watson-assistant/>

is included. Each plotted point in the main expansion lines represents an iteration step of the expansion procedure.

One way to assess the ability of these methods is to evaluate the recall of the system at different steps. Also, the assessment on how much recall the end system gives gives a better indication of how much useful data can be pulled from the unlabeled corpus when the data is skewed. Hence, we compare the expansion paths of I-MLT in Figure 2a to that of I-DP in Figure 2b to see how the different selection models affect the rate of recall increase. It is evident that with the I-DP procedure, we have a more rapid ascent than I-MLT, and are able to continue to find useful examples for many more rounds before hitting the termination condition. This can be seen more clearly for the class “videogame_controller”, where I-MLT seems to converge earlier and reach a recall asymptote whereas I-DP, at the same number of examples sent for labeling, continues to expand.

We further analyze the effectiveness of each expansion method in Figures 2c - 2d by plotting the KL-Divergence between the vocabulary distribution at each iteration to the optimal vocabulary distribution. We see that the KL-Divergence curves approach 0 at similar rates as their relative recall curves. We see the KL-Divergence seems to taper off for a few iterations and then begins a steeper descent again, indicating that our procedures are able to explore and push into new areas of the data space and not prematurely terminate. We notice such exploration more so in Figure 2d, which utilizes the I-DP expansion process.

While the I-DP algorithm expands the training data more than I-MLT, we see in Figure 2 that the per-class recall curves, though they taper off as more examples are added, still do not reach convergence yet. This indicates that there is room for future research on steering the expansion. For instance, the method assumes that the full scope of a class can be reached through chains of neighbors. When this is not true with heavily fragmented classes, we need alternative ways to scatter some seeds to every fragment, by randomization or in-depth analysis of the data characteristics. The generative model also needs to be made more efficient if we are to use it in every iteration.

Conclusion

We present iterative applications, of a search-based selection strategy and a data programming strategy employing weak learning, for expanding text classification data that is independent of knowledge of downstream classifiers. Observations from experiments confirm that the iterative data programming approach can sustain a longer expansion path, allowing the downstream learner to reach better accuracy than the simpler approach. Iterative data programming also leads to better quality training sets sooner, and terminate at obtaining non-exhaustive, labeled training sets with which the learner can attain accuracies very close to those allowed by an optimal training set where all examples must be labeled.

References

Attenberg, J., and Provost, F. 2010. Why label when you can search? alternatives to active learning for applying human re-

sources to build classification models under extreme class imbalance. *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*.

Bach, S.; He, B.; Ratner, A.; and Ré, C. 2017. Learning the structure of generative models without labeled data. *Proceedings of the 34th International Conference on Machine Learning*.

Du, J.; Ling, C. X.; and Zhou, Z.-H. 2010. When does co-training work in real data. *IEEE Transactions on Knowledge and Data Engineering*.

Guzmán-Cabrera, R.; Montes-y Gómez, M.; Rosso, P.; and Vilaseñor-Pineda, L. 2009. Using the web as corpus for self-training text categorization. *Information Retrieval Journal* 400–415.

Karlos, S.; Fazakisb, N.; Kotsiantisc, S.; and Sgarbas, K. 2017. An empirical study of active learning for text classification. *International Conference on Knowledge Based and Intelligent Information and Engineering Systems*.

Liu, W.; Dai, B.; Humayun, A.; Tay, C.; Yu, C.; Smith, L. B.; Rehg, J. M.; and Song, L. 2017. Iterative machine teaching. *Proceedings of the 34th International Conference on Machine Learning*.

Mallinar, N.; Shah, A.; Ugrani, R.; Gupta, A.; Gurusankar, M.; Ho, T. K.; Liao, Q. V.; Zhang, Y.; Bellamy, R. K.; Yates, R.; Desmarais, C.; and McGregor, B. 2019. Bootstrapping conversational agents with weak supervision. *Innovative Applications of Artificial Intelligence Conferences*.

Nashaat, M.; Ghosh, A.; Miller, J.; Quader, S.; Marston, C.; and Puget, J.-F. 2018. Hybridization of active learning and data programming for labeling large industrial datasets. *2018 IEEE International Conference on Big Data*.

Nigam, B.; Ahirwal, P.; Salve, S.; and Vamney, S. 2011. Document classification using expectation maximization with semi supervised learning. *International Journal on Soft Computing*.

Ratner, A.; De Sa, C.; Wu, S.; Selsam, D.; and Ré, C. 2016. Data programming: Creating large training sets, quickly. *30th Conference on Neural Information Processing Systems*.

Ratner, A.; Bach, S.; Ehrenberg, H.; Fries, J.; Wu, S.; and Ré, C. 2018. Snorkel: Rapid training data creation with weak supervision. *The 44th International Conference on Very Large Data Bases*.

Settles, B. 2009. Active learning literature survey. Computer Sciences Technical Report 1648, University of Wisconsin–Madison.

Simard, P. Y.; Amershi, S.; et al. 2017. Machine teaching: A new paradigm for building machine learning systems. *arXiv preprint arXiv:1707.06742*.

Williams, J. D.; Niraula, N. B.; Dasigi, P.; Lakshmiratan, A.; Suarez, C. G. J.; Reddy, M.; and Zweig, G. 2015. Rapidly scaling dialog systems with interactive learning. In *Natural Language Dialog Systems and Intelligent Assistants*. Springer. 1–13.

Zhu, X. 2015. Machine teaching: An inverse problem to machine learning and an approach toward optimal education. *AAAI*.