
Data Programming: Creating Large Training Sets, Quickly

Alexander Ratner, Christopher De Sa, Sen Wu, Daniel Selsam, Christopher Ré
Stanford University
{ajratner, cdesa, senwu, dselsam, chrisrmre}@stanford.edu

Abstract

Large labeled training sets are the critical building blocks of supervised learning methods and are key enablers of deep learning techniques. For some applications, creating labeled training sets is the most time-consuming and expensive part of applying machine learning. We therefore propose a paradigm for the programmatic creation of training sets called *data programming* in which users provide a set of *labeling functions*, which are programs that heuristically label subsets of the data, but that are noisy and may conflict. By viewing these labeling functions as implicitly describing a generative model for this noise, we show that we can recover the parameters of this model to “denoise” the generated training set, and establish theoretically that we can recover the parameters of these generative models in a handful of settings. We then show how to modify a discriminative loss function to make it noise-aware, and demonstrate our method over a range of discriminative models including logistic regression and LSTMs. Experimentally, on the 2014 TAC-KBP Slot Filling challenge, we show that data programming would have led to a new winning score, and also show that applying data programming to an LSTM model leads to a TAC-KBP score almost 6 F1 points over a state-of-the-art LSTM baseline (and into second place in the competition). Additionally, in initial user studies we observed that data programming may be an easier way for non-experts to create machine learning models when training data is limited or unavailable.

1 Introduction

Many of the major machine learning breakthroughs of the last decade have been catalyzed by the release of a new labeled training dataset.¹ Supervised learning approaches that use such datasets have increasingly become key building blocks of applications throughout science and industry. This trend has also been fueled by the recent empirical success of automated feature generation approaches, notably deep learning methods such as long short term memory (LSTM) networks [14], which ameliorate the burden of feature engineering given large enough labeled training sets. For many real-world applications, however, large hand-labeled training sets do not exist, and are prohibitively expensive to create due to requirements that labelers be experts in the application domain. Furthermore, applications’ needs often change, necessitating new or modified training sets.

To help reduce the cost of training set creation, we propose *data programming*, a paradigm for the programmatic creation of training datasets. Data programming extends the idea of *distant supervision*, in which an external knowledge base is mapped onto an input dataset to generate training examples [22]. In data programming, users provide a set of heuristic *labeling functions*, which are user-defined programs that each provide a label for some subset of the data, and collectively generate a large but potentially overlapping set of training labels. These labeling functions can be more general than distant supervision mappings—they can use external knowledge bases (as in distant

¹<http://www.spacemachine.net/views/2016/3/datasets-over-algorithms>

supervision), model an individual annotator’s labels (as in crowdsourcing), or leverage a combination of domain-specific patterns and dictionaries—and thus may have widely varying error rates and may conflict on certain data points. To address this, we model the labeling functions as a generative process, which lets us automatically denoise the resulting training set by learning the accuracies of the labeling functions along with their correlation structure. In turn, we use this model of the training set to optimize a stochastic version of the loss function of the discriminative model that we desire to train. We show that, given certain conditions on the labeling functions, our method achieves the same asymptotic scaling as supervised learning methods, but that our scaling depends on the amount of *unlabeled* data, and uses only a fixed number of labeling functions.

Data programming is in part motivated by the challenges that users faced when applying prior programmatic supervision approaches, and is intended to be a new software engineering paradigm for the creation and management of training sets. For example, consider the scenario when two labeling functions of differing quality and scope overlap and possibly conflict on certain training examples; in prior approaches the user would have to decide which one to use, or how to somehow integrate the signal from both. In data programming, we accomplish this automatically by learning a model of the training set that includes both labeling functions. Additionally, users are often aware of, or able to induce, dependencies between their labeling functions. In data programming, users can provide a dependency graph to indicate, for example, that two labeling functions are similar, or that one “fixes” or “reinforces” another. We describe cases in which we can learn the strength of these dependencies, and for which our generalization is again asymptotically identical to the supervised case.

One further motivation for our method is driven by the observation that users often struggle with selecting *features* for their models, which is a traditional development bottleneck given fixed-size training sets. However, initial feedback from users suggests that writing labeling functions in the framework of data programming may be easier [12]. While the impact of a feature on end performance is dependent on the training set and on statistical characteristics of the model, a labeling function has a simple and intuitive optimality criterion: that it labels data correctly. Motivated by this, we explore whether we can flip the traditional machine learning development process on its head, having users instead focus on generating training sets large enough to support automatically-generated features.

Summary of Contributions and Outline Our first contribution is the *data programming* framework, in which users can implicitly describe a rich generative model for a training set in a more flexible and general way than in previous approaches. In Section 3, we first explore a simple model in which labeling functions are conditionally independent. We show here that under certain conditions, the sample complexity is nearly the same as in the labeled case. In Section 4, we extend our results to more sophisticated data programming models, generalizing related results in crowdsourcing [17]. In Section 5, we validate our approach experimentally on large real-world text relation extraction tasks in genomics, pharmacogenomics and news domains, where we show an average 2.34 point F1 score improvement over a baseline distant supervision approach—including what would have been a new competition-winning score for the 2014 TAC-KBP Slot Filling competition. Using LSTM-generated features, we additionally would have placed second in this competition, achieving a 5.98 point F1 score gain over a state-of-the-art LSTM baseline [32]. Additionally, we describe promising feedback from a usability study with a group of bioinformatics users.

2 Related Work

Our work builds on many previous approaches in machine learning. *Distant supervision* is one preceding paradigm for programmatically creating training sets. The canonical example is relation extraction from text, wherein a knowledge base of known relations is heuristically mapped to label a set of mentions in an input corpus as ground truth examples [8, 22]. Basic extensions group these mapped examples by the particular textual pattern w that they occur with, and cast the problem as a *multiple instance learning* one [15, 25]. Other extensions actually model the accuracy of this pattern w using a discriminative feature-based model [26], or generative models such as hierarchical topic models [1, 27, 31]. Like our approach, these latter methods model a generative process of training set creation, however in a proscribed way that is not based on user input as in our approach. There is also a wealth of examples where additional heuristic patterns used to label training data are collected from unlabeled data [7] or directly from users [21, 29], in a similar manner to our approach, but without any framework to deal with the fact that said labels are explicitly noisy.

Crowdsourcing is widely used for various machine learning tasks [13, 18]. Of particular relevance to our problem setting is the theoretical question of how to model the accuracy of various experts without ground truth available, classically raised in the context of crowdsourcing [10]. More recent results provide formal guarantees even in the absence of labeled data using various approaches [4, 9, 16, 17, 24, 33]. Our model can capture the basic model of the crowdsourcing setting, and can be considered equivalent in the independent case (Sec. 3). However, in addition to generalizing beyond getting inputs solely from human annotators, we also model user-supplied dependencies between the “labelers” in our model, which is not natural within the context of crowdsourcing. Additionally, while crowdsourcing results focus on the regime of a large number of labelers each labeling a small subset of the data, we consider a small set of labeling functions each labeling a large portion of the dataset.

Co-training is a classic procedure for effectively utilizing both a small amount of labeled data and a large amount of unlabeled data by selecting two conditionally independent *views* of the data [5]. In addition to not needing a set of labeled data, and allowing for more than two views (labeling functions in our case), our approach allows explicit modeling of dependencies between views, for example allowing observed issues with dependencies between views to be explicitly modeled [19].

Boosting is a well known procedure for combining the output of many “weak” classifiers to create a strong classifier in a supervised setting [28]. Recently, boosting-like methods have been proposed which leverage unlabeled data in addition to labeled data, which is also used to set constraints on the accuracies of the individual classifiers being ensembled [3]. This is similar in spirit to our approach, except that labeled data is not explicitly necessary in ours, and richer dependency structures between our “heuristic” classifiers (labeling functions) are supported.

The general case of *learning with noisy labels* is treated both in classical [20] and more recent contexts [23]. It has also been studied specifically in the context of *label-noise robust* logistic regression [6]. We consider the more general scenario where multiple noisy labeling functions can conflict and have dependencies.

3 The Data Programming Paradigm

In many applications, we would like to use machine learning, but we face the following challenges: (i) *hand-labeled* training data is not available, and is prohibitively expensive to obtain in sufficient quantities as it requires expensive domain expert labelers; (ii) *related external knowledge bases* are either unavailable or insufficiently specific, precluding a traditional distant supervision or co-training approach; (iii) *application specifications* are in flux, changing the model we ultimately wish to learn.

In such a setting, we would like a simple, scalable and adaptable approach for supervising a model applicable to our problem. More specifically, we would ideally like our approach to achieve ϵ expected loss with high probability, given $O(1)$ inputs of some sort from a domain-expert user, rather than the traditional $\tilde{O}(\epsilon^{-2})$ hand-labeled training examples required by most supervised methods (where \tilde{O} notation hides logarithmic factors). To this end, we propose *data programming*, a paradigm for the programmatic creation of training sets, which enables domain-experts to more rapidly train machine learning systems and has the potential for this type of scaling of expected loss. In data programming, rather than manually labeling each example, users instead describe the *processes* by which these points could be labeled by providing a set of heuristic rules called *labeling functions*.

In the remainder of this paper, we focus on a binary classification task in which we have a distribution π over object and class pairs $(x, y) \in \mathcal{X} \times \{-1, 1\}$, and we are concerned with minimizing the logistic loss under a linear model given some *features*,

$$l(w) = \mathbf{E}_{(x,y) \sim \pi} \left[\log(1 + \exp(-w^T f(x)y)) \right],$$

where without loss of generality, we assume that $\|f(x)\| \leq 1$. Then, a labeling function $\lambda_i : \mathcal{X} \mapsto \{-1, 0, 1\}$ is a user-defined function that encodes some domain heuristic, which provides a (non-zero) label for some subset of the objects. As part of a *data programming specification*, a user provides some m labeling functions, which we denote in vectorized form as $\lambda : \mathcal{X} \mapsto \{-1, 0, 1\}^m$.

Example 3.1. To gain intuition about labeling functions, we describe a simple text relation extraction example. In Figure 1, we consider the task of classifying co-occurring gene and disease mentions as either expressing a causal relation or not. For example, given the sentence “Gene A causes disease B”, the object $x = (A, B)$ has true class $y = 1$. To construct a training set, the user writes three labeling

```


def lambda_1(x):
    return 1 if (x.gene, x.pheno) in KNOWN_RELATIONS_1 else 0

def lambda_2(x):
    return -1 if re.match(r'.*not_cause.*', x.text_between) else 0

def lambda_3(x):
    return 1 if re.match(r'.*associated.*', x.text_between)
        and (x.gene, x.pheno) in KNOWN_RELATIONS_2 else 0

```

(a) An example set of three labeling functions written by a user.



(b) The generative model of a training set defined by the user input (unary factors omitted).

Figure 1: An example of extracting mentions of gene-disease relations from the scientific literature.

functions (Figure 1a). In λ_1 , an external structured knowledge base is used to label a few objects with relatively high accuracy, and is equivalent to a traditional distant supervision rule (see Sec. 2). λ_2 uses a purely heuristic approach to label a much larger number of examples with lower accuracy. Finally, λ_3 is a “hybrid” labeling function, which leverages a knowledge base and a heuristic.

A labeling function need not have perfect accuracy or recall; rather, it represents a pattern that the user wishes to impart to their model and that is easier to encode as a labeling function than as a set of hand-labeled examples. As illustrated in Ex. 3.1, labeling functions can be based on external knowledge bases, libraries or ontologies, can express heuristic patterns, or some hybrid of these types; we see evidence for the existence of such diversity in our experiments (Section 5). The use of labeling functions is also strictly more general than manual annotations, as a manual annotation can always be directly encoded by a labeling function. Importantly, labeling functions can overlap, conflict, and even have dependencies which users can provide as part of the data programming specification (see Section 4); our approach provides a simple framework for these inputs.

Independent Labeling Functions We first describe a model in which the labeling functions label independently, given the true label class. Under this model, each labeling function λ_i has some probability β_i of labeling an object and then some probability α_i of labeling the object correctly; for simplicity we also assume here that each class has probability 0.5. This model has distribution

$$\mu_{\alpha, \beta}(\Lambda, Y) = \frac{1}{2} \prod_{i=1}^m (\beta_i \alpha_i \mathbf{1}_{\{\Lambda_i=Y\}} + \beta_i (1 - \alpha_i) \mathbf{1}_{\{\Lambda_i=-Y\}} + (1 - \beta_i) \mathbf{1}_{\{\Lambda_i=0\}}), \quad (1)$$

where $\Lambda \in \{-1, 0, 1\}^m$ contains the labels output by the labeling functions, and $Y \in \{-1, 1\}$ is the predicted class. If we allow the parameters $\alpha \in \mathbb{R}^m$ and $\beta \in \mathbb{R}^m$ to vary, (1) specifies a family of generative models. In order to expose the scaling of the expected loss as the size of the unlabeled dataset changes, we will assume here that $0.3 \leq \beta_i \leq 0.5$ and $0.8 \leq \alpha_i \leq 0.9$. We note that while these arbitrary constraints can be changed, they are roughly consistent with our applied experience, where users tend to write high-accuracy and high-coverage labeling functions.

Our first goal will be to learn which parameters (α, β) are most consistent with our observations—our unlabeled training set—using maximum likelihood estimation. To do this for a particular training set $S \subset X$, we will solve the problem

$$(\hat{\alpha}, \hat{\beta}) = \arg \max_{\alpha, \beta} \sum_{x \in S} \log \mathbf{P}_{(\Lambda, Y) \sim \mu_{\alpha, \beta}} (\Lambda = \lambda(x)) = \arg \max_{\alpha, \beta} \sum_{x \in S} \log \left(\sum_{y' \in \{-1, 1\}} \mu_{\alpha, \beta}(\lambda(x), y') \right) \quad (2)$$

In other words, we are maximizing the probability that the observed labels produced on our training examples occur under the generative model in (1). In our experiments, we use stochastic gradient descent to solve this problem; since this is a standard technique, we defer its analysis to the appendix.

Noise-Aware Empirical Loss Given that our parameter learning phase has successfully found some $\hat{\alpha}$ and $\hat{\beta}$ that accurately describe the training set, we can now proceed to estimate the parameter w which minimizes the expected risk of a linear model over our feature mapping f , given $\hat{\alpha}, \hat{\beta}$. To do so, we define the *noise-aware empirical risk* $L_{\hat{\alpha}, \hat{\beta}}$ with regularization parameter ρ , and compute the *noise-aware empirical risk minimizer*

$$\hat{w} = \arg \min_w L_{\hat{\alpha}, \hat{\beta}}(w; S) = \arg \min_w \frac{1}{|S|} \sum_{x \in S} \mathbf{E}_{(\Lambda, Y) \sim \mu_{\hat{\alpha}, \hat{\beta}}} \left[\log \left(1 + e^{-w^T f(x) Y} \right) \middle| \Lambda = \lambda(x) \right] + \rho \|w\|^2 \quad (3)$$

This is a logistic regression problem, so it can be solved using stochastic gradient descent as well.

We can in fact prove that stochastic gradient descent running on (2) and (3) is guaranteed to produce accurate estimates, under conditions which we describe now. First, the problem distribution π needs to be accurately modeled by some distribution μ in the family that we are trying to learn. That is, for some α^* and β^* ,

$$\forall \Lambda \in \{-1, 0, 1\}^m, Y \in \{-1, 1\}, \mathbf{P}_{(x,y) \sim \pi^*}(\lambda(x) = \Lambda, y = Y) = \mu_{\alpha^*, \beta^*}(\Lambda, Y). \quad (4)$$

Second, given an example $(x, y) \sim \pi^*$, the class label y must be independent of the features $f(x)$ given the labels $\lambda(x)$. That is,

$$(x, y) \sim \pi^* \Rightarrow y \perp f(x) \mid \lambda(x). \quad (5)$$

This assumption encodes the idea that the labeling functions, while they may be arbitrarily dependent on the features, provide sufficient information to accurately identify the class. Third, we assume that the algorithm used to solve (3) has bounded generalization risk such that for some parameter χ ,

$$\mathbf{E}_{\hat{w}} \left[\mathbf{E}_S \left[L_{\hat{\alpha}, \hat{\beta}}(\hat{w}; S) \right] - \min_w \mathbf{E}_S \left[L_{\hat{\alpha}, \hat{\beta}}(w; S) \right] \right] \leq \chi. \quad (6)$$

Under these conditions, we make the following statement about the accuracy of our estimates, which is a simplified version of a theorem that is detailed in the appendix.

Theorem 1. *Suppose that we run data programming, solving the problems in (2) and (3) using stochastic gradient descent to produce $(\hat{\alpha}, \hat{\beta})$ and \hat{w} . Suppose further that our setup satisfies the conditions (4), (5), and (6), and suppose that $m \geq 2000$. Then for any $\epsilon > 0$, if the number of labeling functions m and the size of the input dataset S are large enough that*

$$|S| \geq \frac{356}{\epsilon^2} \log \left(\frac{m}{3\epsilon} \right)$$

then our expected parameter error and generalization risk can be bounded by

$$\mathbf{E} \left[\|\hat{\alpha} - \alpha^*\|^2 \right] \leq m\epsilon^2 \quad \mathbf{E} \left[\|\hat{\beta} - \beta^*\|^2 \right] \leq m\epsilon^2 \quad \mathbf{E} \left[l(\hat{w}) - \min_w l(w) \right] \leq \chi + \frac{\epsilon}{27\rho}.$$

We select $m \geq 2000$ to simplify the statement of the theorem and give the reader a feel for how ϵ scales with respect to $|S|$. The full theorem with scaling in each parameter (and for arbitrary m) is presented in the appendix. This result establishes that to achieve both expected loss and parameter estimate error ϵ , it suffices to have only $m = O(1)$ labeling functions and $|S| = \tilde{O}(\epsilon^{-2})$ training examples, which is the same asymptotic scaling exhibited by methods that use labeled data. This means that data programming achieves the same learning rate as methods that use labeled data, while requiring asymptotically less work from its users, who need to specify $O(1)$ labeling functions rather than manually label $\tilde{O}(\epsilon^{-2})$ examples. In contrast, in the crowdsourcing setting [17], the number of workers m tends to infinity while here it is constant while the dataset grows. These results provide some explanation of why our experimental results suggest that a small number of rules with a large unlabeled training set can be effective at even complex natural language processing tasks.

4 Handling Dependencies

In our experience with data programming, we have found that users often write labeling functions that have clear dependencies among them. As more labeling functions are added as the system is developed, an implicit dependency structure arises naturally amongst the labeling functions: modeling these dependencies can in some cases improve accuracy. We describe a method by which the user can specify this dependency knowledge as a *dependency graph*, and show how the system can use it to produce better parameter estimates.

Label Function Dependency Graph To support the injection of dependency information into the model, we augment the data programming specification with a *label function dependency graph*, $G \subset \mathcal{D} \times \{1, \dots, m\} \times \{1, \dots, m\}$, which is a directed graph over the labeling functions, each of the edges of which is associated with a *dependency type* from a class of dependencies \mathcal{D} appropriate to the domain. From our experience with practitioners, we identified four commonly-occurring types of dependencies as illustrative examples: *similar*, *fixing*, *reinforcing*, and *exclusive* (see Figure 2).

For example, suppose that we have two functions λ_1 and λ_2 , and λ_2 typically labels only when (i) λ_1 also labels, (ii) λ_1 and λ_2 disagree in their labeling, and (iii) λ_2 is actually correct. We call this a *fixing dependency*, since λ_2 *fixes* mistakes made by λ_1 . If λ_1 and λ_2 were to typically agree rather than disagree, this would be a *reinforcing dependency*, since λ_2 reinforces a subset of the labels of λ_1 .




Figure 2: Examples of labeling function dependency predicates.

Modeling Dependencies The presence of dependency information means that we can no longer model our labels using the simple Bayesian network in (1). Instead, we model our distribution as a factor graph. This standard technique lets us describe the family of generative distributions in terms of a known *factor function* $h : \{-1, 0, 1\}^m \times \{-1, 1\} \mapsto \{-1, 0, 1\}^M$ (in which each entry h_i represents a factor), and an unknown parameter $\theta \in \mathbb{R}^M$ as

$$\mu_\theta(\Lambda, Y) = Z_\theta^{-1} \exp(\theta^T h(\Lambda, Y)),$$

where Z_θ is the *partition function* which ensures that μ is a distribution. Next, we will describe how we define h using information from the dependency graph.

To construct h , we will start with some base factors, which we inherit from (1), and then augment them with additional factors representing dependencies. For all $i \in \{1, \dots, m\}$, we let

$$h_0(\Lambda, Y) = Y, \quad h_i(\Lambda, Y) = \Lambda_i Y, \quad h_{m+i}(\Lambda, Y) = \Lambda_i, \quad h_{2m+i}(\Lambda, Y) = \Lambda_i^2 Y, \quad h_{3m+i}(\Lambda, Y) = \Lambda_i^2.$$

These factors alone are sufficient to describe any distribution for which the labels are mutually independent, given the class: this includes the independent family in (1).

We now proceed by adding additional factors to h , which model the dependencies encoded in G . For each dependency edge (d, i, j) , we add one or more factors to h as follows. For a near-duplicate dependency on (i, j) , we add a single factor $h_i(\Lambda, Y) = \mathbf{1}\{\Lambda_i = \Lambda_j\}$, which increases our prior probability that the labels will agree. For a fixing dependency, we add two factors, $h_i(\Lambda, Y) = -\mathbf{1}\{\Lambda_i = 0 \wedge \Lambda_j \neq 0\}$ and $h_{i+1}(\Lambda, Y) = \mathbf{1}\{\Lambda_i = -Y \wedge \Lambda_j = Y\}$, which encode the idea that λ_j labels only when λ_i does, and that λ_j fixes errors made by λ_i . The factors for a reinforcing dependency are the same, except that $h_{i+1}(\Lambda, Y) = \mathbf{1}\{\Lambda_i = Y \wedge \Lambda_j = Y\}$. Finally, for an exclusive dependency, we have a single factor $h_i(\Lambda, Y) = -\mathbf{1}\{\Lambda_i \neq 0 \wedge \Lambda_j \neq 0\}$.

Learning with Dependencies We can again solve a maximum likelihood problem like (2) to learn the parameter $\hat{\theta}$. Using the results, we can continue on to find the noise-aware empirical loss minimizer by solving the problem in (3). In order to solve these problems in the dependent case, we typically invoke stochastic gradient descent, using Gibbs sampling to sample from the distributions used in the gradient update. Under conditions similar to those in Section 3, we can again provide a bound on the accuracy of these results. We define these conditions now. First, there must be some set $\Theta \subset \mathbb{R}^M$ that we know our parameter lies in. This is analogous to the assumptions on α_i and β_i we made in Section 3, and we can state the following analogue of (4):

$$\exists \theta^* \in \Theta \text{ s.t. } \forall (\Lambda, Y) \in \{-1, 0, 1\}^m \times \{-1, 1\}, \mathbf{P}_{(x,y) \sim \pi^*}(\lambda(x) = \Lambda, y = Y) = \mu_{\theta^*}(\Lambda, Y). \quad (7)$$

Second, for any $\theta \in \Theta$, it must be possible to accurately learn θ from full (i.e. labeled) samples of μ_θ . More specifically, there exists an unbiased estimator $\hat{\theta}(T)$ that is a function of some dataset T of independent samples from μ_θ such that, for some $c > 0$ and for all $\theta \in \Theta$,

$$\text{Cov}(\hat{\theta}(T)) \leq (2c|T|)^{-1}I. \quad (8)$$

Third, for any two feasible models θ_1 and $\theta_2 \in \Theta$,

$$\mathbf{E}_{(\Lambda_1, Y_1) \sim \mu_{\theta_1}} [\text{Var}_{(\Lambda_2, Y_2) \sim \mu_{\theta_2}} (Y_2 | \Lambda_1 = \Lambda_2)] \leq cM^{-1}. \quad (9)$$

That is, we'll usually be reasonably sure in our guess for the value of Y , even if we guess using distribution μ_{θ_2} while the the labeling functions were actually sampled from (the possibly totally different) μ_{θ_1} . We can now prove the following result about the accuracy of our estimates.

Features	Method	KBP (News)			Genomics			Pharmacogenomics		
		Prec.	Rec.	F1	Prec.	Rec.	F1	Prec.	Rec.	F1
Hand-tuned	ITR	51.15	26.72	35.10	83.76	41.67	55.65	68.16	49.32	57.23
	DP	50.52	29.21	37.02	83.90	43.43	57.24	68.36	54.80	60.83
LSTM	ITR	37.68	28.81	32.66	69.07	50.76	58.52	32.35	43.84	37.23
	DP	47.47	27.88	35.78	75.48	48.48	58.99	37.63	47.95	42.17

Table 1: Precision/Recall/F1 scores using data programming (DP), as compared to distant supervision ITR approach, with both hand-tuned and LSTM-generated features.

Theorem 2. Suppose that we run stochastic gradient descent to produce $\hat{\theta}$ and \hat{w} , and that our setup satisfies the conditions (5)-(9). Then for any $\epsilon > 0$, if the input dataset S is large enough that

$$|S| \geq \frac{2}{c^2 \epsilon^2} \log \left(\frac{2 \|\theta_0 - \theta^*\|^2}{\epsilon} \right),$$

then our expected parameter error and generalization risk can be bounded by

$$\mathbf{E} \left[\|\hat{\theta} - \theta^*\|^2 \right] \leq M \epsilon^2 \quad \mathbf{E} \left[l(\hat{w}) - \min_w l(w) \right] \leq \chi + \frac{c\epsilon}{2\rho}.$$

As in the independent case, this shows that we need only $|S| = \tilde{O}(\epsilon^{-2})$ unlabeled training examples to achieve error $O(\epsilon)$, which is the same asymptotic scaling as supervised learning methods. This suggests that while we pay a computational penalty for richer dependency structures, we are no less statistically efficient. In the appendix, we provide more details, including an explicit description of the algorithm and the step size used to achieve this result.

5 Experiments

We seek to experimentally validate three claims about our approach. Our first claim is that data programming can be an effective paradigm for building high quality machine learning systems, which we test across three real-world relation extraction applications. Our second claim is that data programming can be used successfully in conjunction with automatic feature generation methods, such as LSTM models. Finally, our third claim is that data programming is an intuitive and productive framework for domain-expert users, and we report on our initial user studies.

Relation Mention Extraction Tasks In the *relation mention extraction* task, our objects are relation mention *candidates* $x = (e_1, e_2)$, which are pairs of entity mentions e_1, e_2 in unstructured text, and our goal is to learn a model that classifies each candidate as either a true textual assertion of the relation $R(e_1, e_2)$ or not. We examine a news application from the 2014 TAC-KBP Slot Filling challenge², where we extract relations between real-world entities from articles [2]; a clinical genomics application, where we extract causal relations between genetic mutations and phenotypes from the scientific literature³; and a pharmacogenomics application where we extract interactions between genes, also from the scientific literature [21]; further details are included in the Appendix.

For each application, we or our collaborators originally built a system where a training set was programmatically generated by ordering the labeling functions as a sequence of if-then-return statements, and for each candidate, taking the first label emitted by this script as the training label. We refer to this as the *if-then-return (ITR)* approach, and note that it often required significant domain expert development time to tune (weeks or more). For this set of experiments, we then used the same labeling function sets within the framework of data programming. For all experiments, we evaluated on a blind hand-labeled evaluation set. In Table 1, we see that we achieve consistent improvements: on average by 2.34 points in F1 score, including what would have been a winning score on the 2014 TAC-KBP challenge [30].

We observed these performance gains across applications with very different labeling function sets. We describe the labeling function summary statistics—*coverage* is the percentage of objects that had at least one label, *overlap* is the percentage of objects with more than one label, and *conflict* is

²<http://www.nist.gov/tac/2014/KBP/>

³<https://github.com/HazyResearch/dd-genomics>

the percentage of objects with conflicting labels—and see in Table 2 that even in scenarios where m is small, and conflict and overlap is relatively less common, we still realize performance gains. Additionally, on a disease mention extraction task (see Usability Study), which was written from scratch within the data programming paradigm, allowing developers to supply dependencies of the basic types outlined in Sec. 4 led to a 2.3 point F1 score boost.

Application	# of LFs	Coverage	$ S_{\lambda \neq 0} $	Overlap	Conflict	F1 Score Improvement	
						HT	LSTM
KBP (News)	40	29.39	2.03M	1.38	0.15	1.92	3.12
Genomics	146	53.61	256K	26.71	2.05	1.59	0.47
Pharmacogenomics	7	7.70	129K	0.35	0.32	3.60	4.94
Diseases	12	53.32	418K	31.81	0.98	N/A	N/A

Table 2: Labeling function (LF) summary statistics, sizes of generated training sets $S_{\lambda \neq 0}$ (only counting non-zero labels), and relative F1 score improvement over baseline IRT methods for hand-tuned (HT) and LSTM-generated (LSTM) feature sets.

Automatically-generated Features We additionally compare both hand-tuned and automatically-generated features, where the latter are learned via an LSTM recurrent neural network (RNN) [14]. Conventional wisdom states that deep learning methods such as RNNs are prone to overfitting to the biases of the imperfect rules used for programmatic supervision. In our experiments, however, we find that using data programming to denoise the labels can mitigate this issue, and we report a 9.79 point boost to precision and a 3.12 point F1 score improvement on the benchmark 2014 TAC-KBP (News) task, over the baseline *if-then-return* approach. Additionally for comparison, our approach is a 5.98 point F1 score improvement over a state-of-the-art LSTM approach [32].

Usability Study One of our hopes is that a user without expertise in ML will be more productive iterating on labeling functions than on features. To test this, we arranged a hackathon involving a handful of bioinformatics researchers, using our open-source information extraction framework Snorkel⁴ (formerly DDLite). Their goal was to build a disease tagging system which is a common and important challenge in the bioinformatics domain [11]. The hackathon participants did not have access to a labeled training set nor did they perform any feature engineering. The entire effort was restricted to iterative labeling function development and the setup of candidates to be classified. In under eight hours, they had created a training set that led to a model which scored within 10 points of F1 of the supervised baseline; the gap was mainly due to recall issue in the candidate extraction phase. This suggests data programming may be a promising way to build high quality extractors, quickly.

6 Conclusion and Future Work

We introduced data programming, a new approach to generating large labeled training sets. We demonstrated that our approach can be used with automatic feature generation techniques to achieve high quality results. We also provided anecdotal evidence that our methods may be easier for domain experts to use. We hope to explore the limits of our approach on other machine learning tasks that have been held back by the lack of high-quality supervised datasets, including those in other domains such as imaging and structured prediction.

Acknowledgements Thanks to Theodoros Rekatsinas, Manas Joglekar, Henry Ehrenberg, Jason Fries, Percy Liang, the DeepDive and DDLite users and many others for their helpful conversations. The authors acknowledge the support of: DARPA FA8750-12-2-0335; NSF IIS-1247701; NSFCCF-1111943; DOE 108845; NSF CCF-1337375; DARPA FA8750-13-2-0039; NSF IIS-1353606; ONR N000141210041 and N000141310129; NIH U54EB020405; DARPA’s SIMPLEX program; Oracle; NVIDIA; Huawei; SAP Labs; Sloan Research Fellowship; Moore Foundation; American Family Insurance; Google; and Toshiba. The views and conclusions expressed in this material are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA, AFRL, NSF, ONR, NIH, or the U.S. Government.

⁴snorkel.stanford.edu

References

- [1] E. Alfonseca, K. Filippova, J.-Y. Delort, and G. Garrido. Pattern learning for relation extraction with a hierarchical topic model. In *Proceedings of the ACL*.
- [2] G. Angeli, S. Gupta, M. Jose, C. D. Manning, C. Ré, J. Tibshirani, J. Y. Wu, S. Wu, and C. Zhang. Stanford’s 2014 slot filling systems. *TAC KBP*, 695, 2014.
- [3] A. Balsubramani and Y. Freund. Scalable semi-supervised aggregation of classifiers. In *Advances in Neural Information Processing Systems*, pages 1351–1359, 2015.
- [4] D. Berend and A. Kontorovich. Consistency of weighted majority votes. In *NIPS 2014*.
- [5] A. Blum and T. Mitchell. Combining labeled and unlabeled data with co-training. In *Proceedings of the eleventh annual conference on Computational learning theory*, pages 92–100. ACM, 1998.
- [6] J. Bootkrajang and A. Kabán. Label-noise robust logistic regression and its applications. In *Machine Learning and Knowledge Discovery in Databases*, pages 143–158. Springer, 2012.
- [7] R. Bunescu and R. Mooney. Learning to extract relations from the web using minimal supervision. In *Annual meeting-association for Computational Linguistics*, volume 45, page 576, 2007.
- [8] M. Craven, J. Kumlien, et al. Constructing biological knowledge bases by extracting information from text sources. In *ISMB*, volume 1999, pages 77–86, 1999.
- [9] N. Dalvi, A. Dasgupta, R. Kumar, and V. Rastogi. Aggregating crowdsourced binary ratings. In *Proceedings of the 22Nd International Conference on World Wide Web, WWW ’13*, pages 285–294, 2013.
- [10] A. P. Dawid and A. M. Skene. Maximum likelihood estimation of observer error-rates using the em algorithm. *Applied statistics*, pages 20–28, 1979.
- [11] R. I. Doğan and Z. Lu. An improved corpus of disease mentions in pubmed citations. In *Proceedings of the 2012 workshop on biomedical natural language processing*.
- [12] H. R. Ehrenberg, J. Shin, A. J. Ratner, J. A. Fries, and C. Ré. Data programming with ddlite: putting humans in a different part of the loop. In *HILDA@ SIGMOD*, page 13, 2016.
- [13] H. Gao, G. Barbier, R. Goolsby, and D. Zeng. Harnessing the crowdsourcing power of social media for disaster relief. Technical report, DTIC Document, 2011.
- [14] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [15] R. Hoffmann, C. Zhang, X. Ling, L. Zettlemoyer, and D. S. Weld. Knowledge-based weak supervision for information extraction of overlapping relations. In *Proceedings of the ACL*.
- [16] M. Joglekar, H. Garcia-Molina, and A. Parameswaran. Comprehensive and reliable crowd assessment algorithms. In *Data Engineering (ICDE), 2015 IEEE 31st International Conference on*.
- [17] D. R. Karger, S. Oh, and D. Shah. Iterative learning for reliable crowdsourcing systems. In *Advances in neural information processing systems*, pages 1953–1961, 2011.
- [18] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L.-J. Li, D. A. Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *arXiv preprint arXiv:1602.07332*, 2016.
- [19] M.-A. Krogel and T. Scheffer. Multi-relational learning, text mining, and semi-supervised learning for functional genomics. *Machine Learning*, 57(1-2):61–81, 2004.
- [20] G. Lugosi. Learning with an unreliable teacher. *Pattern Recognition*, 25(1):79 – 87, 1992.
- [21] E. K. Mallory, C. Zhang, C. Ré, and R. B. Altman. Large-scale extraction of gene interactions from full-text literature using deepdive. *Bioinformatics*, 2015.
- [22] M. Mintz, S. Bills, R. Snow, and D. Jurafsky. Distant supervision for relation extraction without labeled data. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL*, 2009.
- [23] N. Natarajan, I. S. Dhillon, P. K. Ravikumar, and A. Tewari. Learning with noisy labels. In *Advances in Neural Information Processing Systems* 26.
- [24] F. Parisi, F. Strino, B. Nadler, and Y. Kluger. Ranking and combining multiple predictors without labeled data. *Proceedings of the National Academy of Sciences*, 111(4):1253–1258, 2014.
- [25] S. Riedel, L. Yao, and A. McCallum. Modeling relations and their mentions without labeled text. In *Machine Learning and Knowledge Discovery in Databases*, pages 148–163. Springer, 2010.
- [26] B. Roth and D. Klakow. Feature-based models for improving the quality of noisy training data for relation extraction. In *Proceedings of the 22nd ACM Conference on Knowledge management*.
- [27] B. Roth and D. Klakow. Combining generative and discriminative model scores for distant supervision. In *EMNLP*, pages 24–29, 2013.
- [28] R. E. Schapire and Y. Freund. *Boosting: Foundations and algorithms*. MIT press, 2012.
- [29] J. Shin, S. Wu, F. Wang, C. De Sa, C. Zhang, and C. Ré. Incremental knowledge base construction using deepdive. *Proceedings of the VLDB Endowment*, 8(11):1310–1321, 2015.
- [30] M. Surdeanu and H. Ji. Overview of the english slot filling track at the tac2014 knowledge base population evaluation. In *Proc. Text Analysis Conference (TAC2014)*, 2014.
- [31] S. Takamatsu, I. Sato, and H. Nakagawa. Reducing wrong labels in distant supervision for relation extraction. In *Proceedings of the ACL*.
- [32] P. Verga, D. Belanger, E. Strubell, B. Roth, and A. McCallum. Multilingual relation extraction using compositional universal schema. *arXiv preprint arXiv:1511.06396*, 2015.
- [33] Y. Zhang, X. Chen, D. Zhou, and M. I. Jordan. Spectral methods meet em: A provably optimal algorithm for crowdsourcing. In *Advances in Neural Information Processing Systems* 27, pages 1260–1268. 2014.

A General Theoretical Results

In this section, we will state the full form of the theoretical results we alluded to in the body of the paper. First, we restate, in long form, our setup and assumptions.

We assume that, for some function $h : \{-1, 0, 1\}^m \times \{-1, 1\} \mapsto \{-1, 0, 1\}^M$ of *sufficient statistics*, we are concerned with learning distributions, over the set $\Omega = \{-1, 0, 1\}^m \times \{-1, 1\}$, of the form

$$\pi_\theta(\Lambda, Y) = \frac{1}{Z_\theta} \exp(\theta^T h(\Lambda, Y)), \quad (10)$$

where $\theta \in \mathbb{R}^M$ is a parameter, and Z_θ is the partition function that makes this a distribution. We assume that we are given, i.e. can derive from the data programming specification, some set Θ of *feasible parameters*. This set must have the following two properties.

First, for any $\theta \in \Theta$, learning the parameter θ from (full) samples from π_θ is possible, at least in some sense. More specifically, there exists an unbiased estimator $\hat{\theta}$ that is a function of some number D samples from π_θ (and is unbiased for all $\theta \in \Theta$) such that, for all $\theta \in \Theta$ and for some $c > 0$,

$$\text{Cov}(\hat{\theta}) \leq \frac{I}{2cD}. \quad (11)$$

Second, for any $\theta_1, \theta_2 \in \Theta$,

$$\mathbf{E}_{(\lambda_2, y_2) \sim \pi_{\theta_2}} \left[\mathbf{Var}_{(\lambda_1, y_1) \sim \pi_{\theta_1}} (y_1 | \lambda_1 = \lambda_2) \right] \leq \frac{c}{M}. \quad (12)$$

That is, we'll always be reasonably certain in our guess for the value of y , even if we are totally wrong about the true parameter θ .

On the other hand, we are also concerned with a distribution π^* which ranges over the set $\mathcal{X} \times \{-1, 1\}$, and represents the distribution of training and test examples we are using to learn. These objects are associated with a labeling function $\lambda : \mathcal{X} \mapsto \{-1, 0, 1\}^m$ and a feature function $f : \mathcal{X} \mapsto \mathbb{R}^n$. We make three assumptions about this distribution. First, we assume that, given $(x, y) \sim \pi^*$, the class label y is independent of the features $f(x)$ given the labels $\lambda(x)$. That is,

$$(x, y) \sim \pi^* \Rightarrow y \perp f(x) \mid \lambda(x). \quad (13)$$

Second, we assume that we can describe the relationship between $\lambda(x)$ and y in terms of our family in (10) above. That is, for some parameter $\theta^* \in \Theta$,

$$\mathbf{P}_{(x, y) \sim \pi^*} (\lambda(x) = \Lambda, y = Y) = \pi_{\theta^*}(\Lambda, Y). \quad (14)$$

Third, we assume that the features themselves are bounded; for all $x \in \mathcal{X}$,

$$\|f(x)\| \leq 1. \quad (15)$$

Our goal is twofold. First, we want to recover some estimate $\hat{\theta}$ of the true parameter θ^* . Second, we want to produce a parameter \hat{w} that minimizes the regularized logistic loss

$$l(w) = \mathbf{E}_{(x, y) \sim \pi^*} \left[\log(1 + \exp(-w^T f(x)y)) \right] + \rho \|w\|^2.$$

We actually accomplish this by minimizing a noise-aware loss function, given our recovered parameter $\hat{\theta}$,

$$l_{\hat{\theta}}(w) = \mathbf{E}_{(\bar{x}, \bar{y}) \sim \pi^*} \left[\mathbf{E}_{(\Lambda, Y) \sim \pi_{\hat{\theta}}} \left[\log(1 + \exp(-w^T f(\bar{x})Y)) \mid \Lambda = \lambda(\bar{x}) \right] \right] + \rho \|w\|^2.$$

In fact we can't even minimize this; rather, we will be minimizing the empirical noise-aware loss function, which is only this in expectation. Since the analysis of logistic regression is not itself interesting, we assume that we are able to run some algorithm that produces an estimate \hat{w} which satisfies, for some $\chi > 0$,

$$\mathbf{E} \left[l_{\hat{\theta}}(\hat{w}) - \min_w l_{\hat{\theta}}(w) \mid \hat{\theta} \right] \leq \chi. \quad (16)$$

The algorithm chosen can be anything, but in practice, we use stochastic gradient descent.

We learn $\hat{\theta}$ and \hat{w} by running the following algorithm.

Under these assumptions, we are able to prove the following theorem about the behavior of Algorithm 1.

Algorithm 1 Data Programming

Require: Step size η , dataset $S \subset \mathcal{X}$, and initial parameter $\theta_0 \in \Theta$.

$\theta \rightarrow \theta_0$

for all $x \in S$ **do**

Independently sample (Λ, Y) from π_θ , and $(\bar{\Lambda}, \bar{Y})$ from π_θ conditionally given $\Lambda = \lambda(x)$.

$\theta \leftarrow \theta + \eta(h(\Lambda, Y) - h(\bar{\Lambda}, \bar{Y}))$.

$\theta = P_\Theta(\theta)$ \triangleright Here, P_Θ denotes orthogonal projection onto Θ .

end for

Compute \hat{w} using the algorithm described in (15)

return (θ, \hat{w}) .

Theorem A.1. Suppose that we run Algorithm 1 on a data programming specification that satisfies conditions (11), (12), (13), (14), (15), and (16). Suppose further that, for some parameter $\epsilon > 0$, we use step size

$$\eta = \frac{c\epsilon^2}{4}$$

and our dataset is of a size that satisfies

$$|S| = \frac{2}{c^2\epsilon^2} \log \left(\frac{2\|\theta_0 - \theta^*\|^2}{\epsilon} \right).$$

Then, we can bound the expected parameter error with

$$\mathbf{E} \left[\|\hat{\theta} - \theta^*\|^2 \right] \leq \epsilon^2 M$$

and the expected risk with

$$\mathbf{E} \left[l(\hat{w}) - \min_w l(w) \right] \leq \chi + \frac{c\epsilon}{2\rho}.$$

This theorem's conclusions and assumptions can readily be seen to be identical to those of Theorem 2 in the main body of the paper, except that they apply to the slightly more general case of arbitrary h , rather than h of the explicit form described in the body. Therefore, in order to prove Theorem 2, it suffices to prove Theorem A.1, which we will do in Section C.

B Theoretical Results for Independent Model

For the independent model, we can obtain a more specific version of Theorem A.1. In the independent model, the variables are, as before, $\Lambda \in \{-1, 0, 1\}^m$ and $Y \in \{-1, 1\}$. The sufficient statistics are $\Lambda_i Y$ and Λ_i^2 .

To produce results that make intuitive sense, we also define the alternate parameterization

$$\mathbf{P}_\pi(\Lambda_i | Y) = \begin{cases} \beta_i^{\frac{1+\gamma_i}{2}} & \Lambda_i = Y \\ (1 - \beta_i) & \Lambda_i = 0 \\ \beta_i^{\frac{1-\gamma_i}{2}} & \Lambda_i = -Y \end{cases}.$$

In comparison to the parameters used in the body of the paper, we have

$$\alpha_i = \frac{1 + \gamma_i}{2}.$$

Now, we are concerned with models that are feasible. For a model to be feasible (i.e. for $\theta \in \Theta$), we require that it satisfy, for some constants $\gamma_{\min} > 0$, $\gamma_{\max} > 0$, and β_{\min} ,

$$\gamma_{\min} \leq \gamma_i \leq \gamma_{\max} \quad \beta_{\min} \leq \beta_i \leq \frac{1}{2}.$$

For $0 \leq \beta \leq 1$ and $-1 \leq \gamma \leq 1$.

For this model, we can prove the following corollary to Theorem A.1

Corollary B.1. Suppose that we run Algorithm 1 on an independent data programming specification that satisfies conditions (13), (14), (15), and (16). Furthermore, assume that the number of labeling functions we use satisfies

$$m \geq \frac{9.34 \operatorname{artanh}(\gamma_{\max})}{(\gamma\beta)_{\min} \gamma_{\min}^2} \log \left(\frac{24m}{\beta_{\min}} \right).$$

Suppose further that, for some parameter $\epsilon > 0$, we use step size

$$\eta = \frac{\beta_{\min} \epsilon^2}{16}$$

and our dataset is of a size that satisfies

$$|S| = \frac{32}{\beta_{\min}^2 \epsilon^2} \log \left(\frac{2 \|\theta_0 - \theta^*\|^2}{\epsilon} \right).$$

Then, we can bound the expected parameter error with

$$\mathbf{E} \left[\|\hat{\theta} - \theta^*\|^2 \right] \leq \epsilon^2 M$$

and the expected risk with

$$\mathbf{E} \left[l(\hat{w}) - \min_w l(w) \right] \leq \chi + \frac{\beta_{\min} \epsilon}{8\rho}.$$

We can see that if, as stated in the body of the paper, $\beta_i \geq 0.3$ and $0.8 \leq \alpha_i \leq 0.9$ (which is equivalent to $0.6 \leq \gamma_i \leq 0.8$), then

$$2000 \geq 1896.13 = \frac{9.34 \operatorname{artanh}(0.8)}{0.3 \cdot 0.6^3} \log \left(\frac{24 \cdot 2000}{0.3} \right).$$

This means that, as stated in the paper, $m = 2000$ is sufficient for this corollary to hold with

$$|S| = \frac{32}{0.3^2 \cdot \epsilon^2} \log \left(\frac{2m(\operatorname{artanh}(0.8) - \operatorname{artanh}(0.6))^2}{\epsilon} \right) = \frac{356}{\epsilon^2} \log \left(\frac{m}{3\epsilon} \right).$$

Thus, proving Corollary B.1 is sufficient to prove Theorem 1 from the body of the paper. We will prove Corollary B.1 in Section E

C Proof of Theorem A.1

First, we state some lemmas that will be useful in the proof to come.

Lemma D.1. Given a family of maximum-entropy distributions

$$\pi_{\theta}(x) = \frac{1}{Z_{\theta}} \exp(\theta^T h(x)),$$

for some function of sufficient statistics $h : \Omega \mapsto \mathbb{R}^M$, if we let $J : \mathbb{R}^M \mapsto \mathbb{R}$ be the maximum log-likelihood objective for some event $A \subseteq \Omega$,

$$J(\theta) = \log \mathbf{P}_{x \sim \pi_{\theta}}(x \in A),$$

then its gradient is

$$\nabla J(\theta) = \mathbf{E}_{x \sim \pi_{\theta}} [h(x)|x \in A] - \mathbf{E}_{x \sim \pi_{\theta}} [h(x)]$$

and its Hessian is

$$\nabla^2 J(\theta) = \mathbf{Cov}_{x \sim \pi_{\theta}}(h(x)|x \in A) - \mathbf{Cov}_{x \sim \pi_{\theta}}(h(x)).$$

Lemma D.2. Suppose that we are looking at a distribution from a data programming label model. That is, our maximum-entropy distribution can now be written in terms of two variables, the labeling function values $\lambda \in \{-1, 0, 1\}$ and the class $y \in \{-1, 1\}$, as

$$\pi_{\theta}(\lambda, y) = \frac{1}{Z_{\theta}} \exp(\theta^T h(\lambda, y)),$$

where we assume without loss of generality that for some M , $h(\lambda, y) \in \mathbb{R}^M$ and $\|h(\lambda, y)\|_\infty \leq 1$. If we let $J : \mathbb{R}^M \mapsto \mathbb{R}$ be the maximum expected log-likelihood objective, under another distribution π^* , for the event associated with the observed labeling function values λ ,

$$J(\theta) = \mathbf{E}_{(\lambda^*, y^*) \sim \pi^*} \left[\log \mathbf{P}_{(\lambda, y) \sim \pi_\theta} (\lambda = \lambda^*) \right],$$

then its Hessian can be bounded with

$$\nabla^2 J(\theta) \leq M \mathbf{E}_{(\lambda^*, y^*) \sim \pi^*} \left[\mathbf{Var}_{(\lambda, y) \sim \pi_\theta} (y | \lambda = \lambda^*) \right] - \mathcal{I}(\theta),$$

where $\mathcal{I}(\theta)$ is the Fisher information.

Lemma D.3. Suppose that we are looking at a data programming distribution, as described in the text of Lemma D.2. Suppose further that we are concerned with some feasible set of parameters $\Theta \subset \mathbb{R}^M$, such that the any model with parameters in this space satisfies the following two conditions.

First, for any $\theta \in \Theta$, learning the parameter θ from (full) samples from π_θ is possible, at least in some sense. More specifically, there exists an unbiased estimator $\hat{\theta}$ that is a function of some number D samples from π_θ (and is unbiased for all $\theta \in \Theta$) such that, for all $\theta \in \Theta$ and for some $c > 0$,

$$\mathbf{Cov}(\hat{\theta}) \leq \frac{\mathcal{I}}{2cD}.$$

Second, for any $\theta, \theta^* \in \Theta$,

$$\mathbf{E}_{(\lambda^*, y^*) \sim \pi^*} \left[\mathbf{Var}_{(\lambda, y) \sim \pi_{\theta^*}} (y | \lambda = \lambda^*) \right] \leq \frac{c}{M}.$$

That is, we'll always be reasonably certain in our guess for the value of y , even if we are totally wrong about the true parameter θ^* .

Under these conditions, the function J is strongly concave on Θ with parameter of strong convexity c .

Lemma D.4. Suppose that we are looking at a data programming maximum likelihood estimation problem, as described in the text of Lemma D.2. Suppose further that the objective function J is strongly concave with parameter $c > 0$.

If we run stochastic gradient descent on objective J , using unbiased samples from a true distribution π_{θ^*} , where $\theta^* \in \Theta$, then if we use step size

$$\eta = \frac{c\epsilon^2}{4}$$

and run (using a fresh sample at each iteration) for T steps, where

$$T = \frac{2}{c^2\epsilon^2} \log \left(\frac{2 \|\theta_0 - \theta^*\|^2}{\epsilon} \right)$$

then we can bound the expected parameter estimation error with

$$\mathbf{E} \left[\|\hat{\theta} - \theta^*\|^2 \right] \leq \epsilon^2 M.$$

Lemma D.5. Assume in our model that, without loss of generality, $\|f(x)\| \leq 1$ for all x , and that in our true model π^* , the class y is independent of the features $f(x)$ given the labels $\lambda(x)$.

Suppose that we now want to solve the expected loss minimization problem wherein we minimize the objective

$$l(w) = \mathbf{E}_{(x, y) \sim \pi^*} \left[\log(1 + \exp(-w^T f(x)y)) \right] + \rho \|w\|^2.$$

We actually accomplish this by minimizing our noise-aware loss function, given our chosen parameter $\hat{\theta}$,

$$l_{\hat{\theta}}(w) = \mathbf{E}_{(\bar{x}, \bar{y}) \sim \pi^*} \left[\mathbf{E}_{(\Lambda, Y) \sim \pi_{\hat{\theta}}} \left[\log(1 + \exp(-w^T f(\bar{x})Y)) \mid \Lambda = \lambda(\bar{x}) \right] \right] + \rho \|w\|^2.$$

In fact we can't even minimize this; rather, we will be minimizing the empirical noise-aware loss function, which is only this in expectation. Suppose that doing so produces an estimate \hat{w} which satisfies, for some $\chi > 0$,

$$\mathbf{E} \left[l_{\hat{\theta}}(\hat{w}) - \min_w l_{\hat{\theta}}(w) \mid \hat{\theta} \right] \leq \chi.$$

(Here, the expectation is taken with respect to only the random variable \hat{w} .) Then, we can bound the expected risk with

$$\mathbf{E} \left[l(\hat{w}) - \min_w l(w) \right] \leq \chi + \frac{c\epsilon}{2\rho}.$$

Now, we restate and prove our main theorem.

Theorem A.1. Suppose that we run Algorithm 1 on a data programming specification that satisfies conditions (11), (12), (13), (14), (15), and (16). Suppose further that, for some parameter $\epsilon > 0$, we use step size

$$\eta = \frac{c\epsilon^2}{4}$$

and our dataset is of a size that satisfies

$$|S| = \frac{2}{c^2\epsilon^2} \log\left(\frac{2\|\theta_0 - \theta^*\|^2}{\epsilon}\right).$$

Then, we can bound the expected parameter error with

$$\mathbf{E} \left[\|\hat{\theta} - \theta^*\|^2 \right] \leq \epsilon^2 M$$

and the expected risk with

$$\mathbf{E} \left[l(\hat{w}) - \min_w l(w) \right] \leq \chi + \frac{c\epsilon}{2\rho}.$$

Proof. The bounds on the expected parameter estimation error follow directly from Lemma D.4, and the remainder of the theorem follows directly from Lemma D.5. \square

D Proofs of Lemmas

Lemma D.1. Given a family of maximum-entropy distributions

$$\pi_\theta(x) = \frac{1}{Z_\theta} \exp(\theta^T h(x)),$$

for some function of sufficient statistics $h : \Omega \mapsto \mathbb{R}^M$, if we let $J : \mathbb{R}^M \mapsto \mathbb{R}$ be the maximum log-likelihood objective for some event $A \subseteq \Omega$,

$$J(\theta) = \log \mathbf{P}_{x \sim \pi_\theta} (x \in A),$$

then its gradient is

$$\nabla J(\theta) = \mathbf{E}_{x \sim \pi_\theta} [h(x)|x \in A] - \mathbf{E}_{x \sim \pi_\theta} [h(x)]$$

and its Hessian is

$$\nabla^2 J(\theta) = \mathbf{Cov}_{x \sim \pi_\theta} (h(x)|x \in A) - \mathbf{Cov}_{x \sim \pi_\theta} (h(x)).$$

Proof. For the gradient,

$$\begin{aligned} \nabla J(\theta) &= \nabla \log \mathbf{P}_{\pi_\theta} (A) \\ &= \nabla \log \left(\frac{\sum_{x \in A} \exp(\theta^T h(x))}{\sum_{x \in \Omega} \exp(\theta^T h(x))} \right) \\ &= \nabla \log \left(\sum_{x \in A} \exp(\theta^T h(x)) \right) - \nabla \log \left(\sum_{x \in \Omega} \exp(\theta^T h(x)) \right) \\ &= \frac{\sum_{x \in A} h(x) \exp(\theta^T h(x))}{\sum_{x \in A} \exp(\theta^T h(x))} - \frac{\sum_{x \in \Omega} h(x) \exp(\theta^T h(x))}{\sum_{x \in \Omega} \exp(\theta^T h(x))} \\ &= \mathbf{E}_{x \sim \pi_\theta} [h(x)|x \in A] - \mathbf{E}_{x \sim \pi_\theta} [h(x)]. \end{aligned}$$

And for the Hessian,

$$\begin{aligned}
\nabla^2 J(\theta) &= \nabla \frac{\sum_{x \in A} h(x) \exp(\theta^T h(x))}{\sum_{x \in A} \exp(\theta^T h(x))} - \nabla \frac{\sum_{x \in \Omega} h(x) \exp(\theta^T h(x))}{\sum_{x \in \Omega} \exp(\theta^T h(x))} \\
&= \frac{\sum_{x \in A} h(x) h(x)^T \exp(\theta^T h(x))}{\sum_{x \in A} \exp(\theta^T h(x))} - \frac{\left(\sum_{x \in A} h(x) \exp(\theta^T h(x)) \right) \left(\sum_{x \in A} h(x) \exp(\theta^T h(x)) \right)^T}{\left(\sum_{x \in A} \exp(\theta^T h(x)) \right)^2} \\
&\quad - \left(\frac{\sum_{x \in \Omega} h(x) h(x)^T \exp(\theta^T h(x))}{\sum_{x \in \Omega} \exp(\theta^T h(x))} - \frac{\left(\sum_{x \in \Omega} h(x) \exp(\theta^T h(x)) \right) \left(\sum_{x \in \Omega} h(x) \exp(\theta^T h(x)) \right)^T}{\left(\sum_{x \in \Omega} \exp(\theta^T h(x)) \right)^2} \right) \\
&= \mathbf{E}_{x \sim \pi_\theta} \left[h(x) h(x)^T | x \in A \right] - \mathbf{E}_{x \sim \pi_\theta} [h(x) | x \in A] \mathbf{E}_{x \sim \pi_\theta} [h(x) | x \in A]^T \\
&\quad - \left(\mathbf{E}_{x \sim \pi_\theta} [h(x) h(x)^T] - \mathbf{E}_{x \sim \pi_\theta} [h(x)] \mathbf{E}_{x \sim \pi_\theta} [h(x)]^T \right) \\
&= \mathbf{Cov}_{x \sim \pi_\theta} (h(x) | x \in A) - \mathbf{Cov}_{x \sim \pi_\theta} (h(x)).
\end{aligned}$$

□

Lemma D.2. Suppose that we are looking at a distribution from a data programming label model. That is, our maximum-entropy distribution can now be written in terms of two variables, the labeling function values $\lambda \in \{-1, 0, 1\}$ and the class $y \in \{-1, 1\}$, as

$$\pi_\theta(\lambda, y) = \frac{1}{Z_\theta} \exp(\theta^T h(\lambda, y)),$$

where we assume without loss of generality that for some M , $h(\lambda, y) \in \mathbb{R}^M$ and $\|h(\lambda, y)\|_\infty \leq 1$. If we let $J : \mathbb{R}^M \mapsto \mathbb{R}$ be the maximum expected log-likelihood objective, under another distribution π^* , for the event associated with the observed labeling function values λ ,

$$J(\theta) = \mathbf{E}_{(\lambda^*, y^*) \sim \pi^*} \left[\log \mathbf{P}_{(\lambda, y) \sim \pi_\theta} (\lambda = \lambda^*) \right],$$

then its Hessian can be bounded with

$$\nabla^2 J(\theta) \leq M \mathbf{E}_{(\lambda^*, y^*) \sim \pi^*} \left[\mathbf{Var}_{(\lambda, y) \sim \pi_\theta} (y | \lambda = \lambda^*) \right] - \mathcal{I}(\theta),$$

where $\mathcal{I}(\theta)$ is the Fisher information.

Proof. From the result of Lemma D.1, we have that

$$\nabla^2 J(\theta) = \mathbf{E}_{(\lambda^*, y^*) \sim \pi^*} \left[\mathbf{Cov}_{(\lambda, y) \sim \pi_\theta} (h(\lambda, y) | \lambda = \lambda^*) \right] - \mathbf{Cov}_{(\lambda, y) \sim \pi_\theta} (h(\lambda, y)). \quad (17)$$

We start by defining $h_0(\lambda)$ and $h_1(\lambda)$ such that

$$h(\lambda, y) = h(\lambda, 1) \frac{1+y}{2} + h(\lambda, -1) \frac{1-y}{2} = \frac{h(\lambda, 1) + h(\lambda, -1)}{2} + y \frac{h(\lambda, 1) - h(\lambda, -1)}{2} = h_0(\lambda) + y h_1(\lambda).$$

This allows us to reduce (17) to

$$\nabla^2 J(\theta) = \mathbf{E}_{(\lambda^*, y^*) \sim \pi^*} \left[h_1(\lambda^*) h_1(\lambda^*)^T \mathbf{Var}_{(\lambda, y) \sim \pi_\theta} (y | \lambda = \lambda^*) \right] - \mathbf{Cov}_{(\lambda, y) \sim \pi_\theta} (h(\lambda, y)).$$

On the other hand, the Fisher information of this model at θ is

$$\begin{aligned}
\mathcal{I}(\theta) &= \mathbf{E} \left[(\nabla_\theta \log \pi_\theta(x))^2 \right] \\
&= \mathbf{E} \left[\left(\nabla_\theta \log \left(\frac{\exp(\theta^T h(x))}{\sum_{z \in \Omega} \exp(\theta^T h(z))} \right) \right)^2 \right] \\
&= \mathbf{E} \left[\left(\nabla_\theta \log (\exp(\theta^T h(x))) - \nabla_\theta \log \left(\sum_{z \in \Omega} \exp(\theta^T h(z)) \right) \right)^2 \right] \\
&= \mathbf{E} \left[\left(h(x) - \frac{\sum_{z \in \Omega} h(z) \exp(\theta^T h(z))}{\sum_{z \in \Omega} \exp(\theta^T h(z))} \right)^2 \right] \\
&= \mathbf{E} \left[(h(x) - \mathbf{E} [h(z)])^2 \right] \\
&= \mathbf{Cov} (h(x)).
\end{aligned}$$

Therefore, we can write the second derivative of J as

$$\nabla^2 J(\theta) = \mathbf{E}_{(\lambda^*, y^*) \sim \pi^*} \left[h_1(\lambda^*) h_1(\lambda^*)^T \mathbf{Var}_{(\lambda, y) \sim \pi_\theta} (y | \lambda = \lambda^*) \right] - I(\theta).$$

If we apply the fact that

$$h_1(\lambda^*) h_1(\lambda^*)^T \leq I \|h_1(\lambda^*)\|^2 \leq MI \|h_1(\lambda^*)\|_\infty^2 \leq MI,$$

then we can reduce this to

$$\nabla^2 J(\theta) \leq MI \mathbf{E}_{(\lambda^*, y^*) \sim \pi^*} \left[\mathbf{Var}_{(\lambda, y) \sim \pi_\theta} (y | \lambda = \lambda^*) \right] - I(\theta).$$

This is the desired result. \square

Lemma D.3. *Suppose that we are looking at a data programming distribution, as described in the text of Lemma D.2. Suppose further that we are concerned with some feasible set of parameters $\Theta \subset \mathbb{R}^M$, such that the any model with parameters in this space satisfies the following two conditions.*

First, for any $\theta \in \Theta$, learning the parameter θ from (full) samples from π_θ is possible, at least in some sense. More specifically, there exists an unbiased estimator $\hat{\theta}$ that is a function of some number D samples from π_θ (and is unbiased for all $\theta \in \Theta$) such that, for all $\theta \in \Theta$ and for some $c > 0$,

$$\mathbf{Cov}(\hat{\theta}) \leq \frac{I}{2cD}.$$

Second, for any $\theta, \theta^ \in \Theta$,*

$$\mathbf{E}_{(\lambda^*, y^*) \sim \pi^*} \left[\mathbf{Var}_{(\lambda, y) \sim \pi_\theta} (y | \lambda = \lambda^*) \right] \leq \frac{c}{M}.$$

That is, we'll always be reasonably certain in our guess for the value of y , even if we are totally wrong about the true parameter θ^ .*

Under these conditions, the function J is strongly concave on Θ with parameter of strong convexity c .

Proof. From the Cramér-Rao bound, we know in general that the variance of any unbiased estimator is bounded by the reciprocal of the Fisher information

$$\mathbf{Cov}(\hat{\theta}) \geq (I(\theta))^{-1}.$$

Since for the estimator described in the lemma statement, we have D independent samples from the distribution, it follows that the Fisher information of this experiment is D times the Fisher information of a single sample. Combining this with the bound in the lemma statement on the covariance, we get

$$\frac{I}{2cD} \geq \mathbf{Cov}(\hat{\theta}) \geq (DI(\theta))^{-1}.$$

It follows that

$$I(\theta) \geq 2cI.$$

On the other hand, also from the lemma statement, we can conclude that

$$MI \mathbf{E}_{(\lambda^*, y^*) \sim \pi^*} \left[\mathbf{Var}_{(\lambda, y) \sim \pi_\theta} (y | \lambda = \lambda^*) \right] \leq cI.$$

Therefore, for all $\theta \in \Theta$,

$$\nabla^2 J(\theta) \leq MI \mathbf{E}_{(\lambda^*, y^*) \sim \pi^*} \left[\mathbf{Var}_{(\lambda, y) \sim \pi_\theta} (y | \lambda = \lambda^*) \right] - I(\theta) \leq -cI.$$

This implies that J is strongly concave over Θ , with constant c , as desired. \square

Lemma D.4. *Suppose that we are looking at a data programming maximum likelihood estimation problem, as described in the text of Lemma D.2. Suppose further that the objective function J is strongly concave with parameter $c > 0$.*

If we run stochastic gradient descent on objective J , using unbiased samples from a true distribution π_{θ^} , where $\theta^* \in \Theta$, then if we use step size*

$$\eta = \frac{c\epsilon^2}{4}$$

and run (using a fresh sample at each iteration) for T steps, where

$$T = \frac{2}{c^2 \epsilon^2} \log \left(\frac{2 \|\theta_0 - \theta^*\|^2}{\epsilon} \right)$$

then we can bound the expected parameter estimation error with

$$\mathbf{E} \left[\|\hat{\theta} - \theta^*\|^2 \right] \leq \epsilon^2 M.$$

Proof. First, we note that, in the proof to follow, we can ignore the projection onto the feasible set Θ , since this projection always takes us closer to the optimum θ^* .

If we track the expected distance to the optimum θ^* , then at the next timestep,

$$\|\theta_{t+1} - \theta^*\|^2 = \|\theta_t - \theta^*\|^2 + 2\gamma(\theta_t - \theta^*)^\top \nabla \tilde{J}(\theta_t) + \gamma^2 \|\nabla \tilde{J}(\theta_t)\|^2.$$

Since we can write our stochastic samples in the form

$$\nabla \tilde{J}_t(\theta_t) = h(\lambda_t, y_t) - h(\bar{\lambda}_t, \bar{y}_t),$$

for some samples $\lambda_t, y_t, \bar{\lambda}_t$, and \bar{y}_t , we can conclude that

$$\|\nabla \tilde{J}_t(\theta_t)\|^2 \leq M \|\nabla \tilde{J}_t(\theta_t)\|_\infty^2 \leq 4M.$$

Therefore, taking the expected value conditioned on the filtration,

$$\mathbf{E} \left[\|\theta_{t+1} - \theta^*\|^2 \middle| \mathcal{F}_t \right] = \|\theta_t - \theta^*\|^2 + 2\gamma(\theta_t - \theta^*)^\top \nabla J(\theta_t) + 4\gamma^2 M.$$

Since J is strongly concave,

$$(\theta_t - \theta^*)^\top \nabla J(\theta_t) \leq -c \|\theta_t - \theta^*\|^2;$$

and so,

$$\mathbf{E} \left[\|\theta_{t+1} - \theta^*\|^2 \middle| \mathcal{F}_t \right] \leq (1 - 2\gamma c) \|\theta_t - \theta^*\|^2 + 4\gamma^2 M.$$

If we take the full expectation and subtract the fixed point from both sides,

$$\mathbf{E} \left[\|\theta_{t+1} - \theta^*\|^2 \right] - \frac{2\gamma M}{c} \leq (1 - 2\gamma c) \mathbf{E} \left[\|\theta_t - \theta^*\|^2 \right] + 4\gamma^2 M - \frac{2\gamma M}{c} = (1 - 2\gamma c) \left(\mathbf{E} \left[\|\theta_t - \theta^*\|^2 \right] - \frac{2\gamma M}{c} \right).$$

Therefore,

$$\mathbf{E} \left[\|\theta_t - \theta^*\|^2 \right] - \frac{2\gamma M}{c} \leq (1 - 2\gamma c)^t \left(\|\theta_0 - \theta^*\|^2 - \frac{2\gamma M}{c} \right),$$

and so

$$\mathbf{E} \left[\|\theta_t - \theta^*\|^2 \right] \leq \exp(-2\gamma c t) \|\theta_0 - \theta^*\|^2 + \frac{2\gamma M}{c}.$$

In order to ensure that

$$\mathbf{E} \left[\|\theta_t - \theta^*\|^2 \right] \leq \epsilon^2,$$

it therefore suffices to pick

$$\gamma = \frac{c\epsilon^2}{4M}$$

and

$$t = \frac{2M}{c^2 \epsilon^2} \log \left(\frac{2 \|\theta_0 - \theta^*\|^2}{\epsilon} \right).$$

Substituting $\epsilon^2 \rightarrow \epsilon^2 M$ produces the desired result. \square

Lemma D.5. Assume in our model that, without loss of generality, $\|f(x)\| \leq 1$ for all x , and that in our true model π^* , the class y is independent of the features $f(x)$ given the labels $\lambda(x)$.

Suppose that we now want to solve the expected loss minimization problem wherein we minimize the objective

$$l(w) = \mathbf{E}_{(x,y) \sim \pi^*} \left[\log(1 + \exp(-w^T f(x)y)) \right] + \rho \|w\|^2.$$

We actually accomplish this by minimizing our noise-aware loss function, given our chosen parameter $\hat{\theta}$,

$$l_{\hat{\theta}}(w) = \mathbf{E}_{(\bar{x}, \bar{y}) \sim \pi^*} \left[\mathbf{E}_{(\Lambda, Y) \sim \pi_{\hat{\theta}}} \left[\log(1 + \exp(-w^T f(\bar{x})Y)) \middle| \Lambda = \lambda(\bar{x}) \right] \right] + \rho \|w\|^2.$$

In fact we can't even minimize this; rather, we will be minimizing the empirical noise-aware loss function, which is only this in expectation. Suppose that doing so produces an estimate \hat{w} which satisfies, for some $\chi > 0$,

$$\mathbf{E} \left[l_{\hat{\theta}}(\hat{w}) - \min_w l_{\hat{\theta}}(w) \middle| \hat{\theta} \right] \leq \chi.$$

(Here, the expectation is taken with respect to only the random variable \hat{w} .) Then, we can bound the expected risk with

$$\mathbf{E} \left[l(\hat{w}) - \min_w l(w) \right] \leq \chi + \frac{c\epsilon}{2\rho}.$$

Proof. (To simplify the symbols in this proof, we freely use θ when we mean $\hat{\theta}$.)

The loss function we want to minimize is, in expectation,

$$l(w) = \mathbf{E}_{(x, y) \sim \pi^*} \left[\log(1 + \exp(-w^T f(x)y)) \right] + \rho \|w\|^2.$$

By the law of total expectation,

$$l(w) = \mathbf{E}_{(\bar{x}, \bar{y}) \sim \pi^*} \left[\mathbf{E}_{(x, y) \sim \pi^*} \left[\log(1 + \exp(-w^T f(\bar{x})y)) \middle| x = \bar{x} \right] \right] + \rho \|w\|^2,$$

and by our conditional independence assumption,

$$l(w) = \mathbf{E}_{(\bar{x}, \bar{y}) \sim \pi^*} \left[\mathbf{E}_{(x, y) \sim \pi^*} \left[\log(1 + \exp(-w^T f(\bar{x})y)) \middle| \lambda(x) = \lambda(\bar{x}) \right] \right] + \rho \|w\|^2.$$

Since we know from our assumptions that, for the optimum parameter θ^* ,

$$\mathbf{P}_{(x, y) \sim \pi^*} (\lambda(x) = \Lambda, y = Y) = \mathbf{P}_{(\Lambda, Y) \sim \pi_{\theta^*}} (\Lambda = \Lambda, y = Y),$$

we can rewrite this as

$$l(w) = \mathbf{E}_{(\bar{x}, \bar{y}) \sim \pi^*} \left[\mathbf{E}_{(\Lambda, Y) \sim \pi_{\theta^*}} \left[\log(1 + \exp(-w^T f(\bar{x})Y)) \middle| \Lambda = \lambda(\bar{x}) \right] \right] + \rho \|w\|^2.$$

On the other hand, if we are minimizing the model we got from the previous step, we will be actually minimizing

$$l_{\theta}(w) = \mathbf{E}_{(\bar{x}, \bar{y}) \sim \pi^*} \left[\mathbf{E}_{(\Lambda, Y) \sim \pi_{\theta}} \left[\log(1 + \exp(-w^T f(\bar{x})Y)) \middle| \Lambda = \lambda(\bar{x}) \right] \right] + \rho \|w\|^2.$$

We can reduce this further by noticing that

$$\begin{aligned} & \mathbf{E}_{(\Lambda, Y) \sim \pi_{\theta}} \left[\log(1 + \exp(-w^T f(\bar{x})Y)) \middle| \Lambda = \lambda(\bar{x}) \right] \\ &= \mathbf{E}_{(\Lambda, Y) \sim \pi_{\theta}} \left[\log(1 + \exp(-w^T f(\bar{x}))) \frac{1+Y}{2} + \log(1 + \exp(w^T f(\bar{x}))) \frac{1-Y}{2} \middle| \Lambda = \lambda(\bar{x}) \right] \\ &= \frac{\log(1 + \exp(-w^T f(\bar{x}))) + \log(1 + \exp(w^T f(\bar{x})))}{2} \\ &\quad + \frac{\log(1 + \exp(-w^T f(\bar{x}))) - \log(1 + \exp(w^T f(\bar{x})))}{2} \mathbf{E}_{(\Lambda, Y) \sim \pi_{\theta}} [Y | \Lambda = \lambda(\bar{x})] \\ &= \frac{\log(1 + \exp(-w^T f(\bar{x}))) + \log(1 + \exp(w^T f(\bar{x})))}{2} \\ &\quad - \frac{w^T f(\bar{x})}{2} \mathbf{E}_{(\Lambda, Y) \sim \pi_{\theta}} [Y | \Lambda = \lambda(\bar{x})]. \end{aligned}$$

It follows that the difference between the loss functions will be

$$|l(w) - l_{\theta}(w)| = \left| \mathbf{E}_{(\bar{x}, \bar{y}) \sim \pi^*} \left[\frac{w^T f(\bar{x})}{2} (\mathbf{E}_{(\Lambda, Y) \sim \pi_{\theta}} [Y | \Lambda = \lambda(\bar{x})] - \mathbf{E}_{(\Lambda, Y) \sim \pi_{\theta^*}} [Y | \Lambda = \lambda(\bar{x})]) \right] \right|.$$

Now, we can compute that

$$\begin{aligned}
\nabla_{\theta} \mathbf{E}_{(\Lambda, Y) \sim \pi_{\theta}} [Y | \Lambda = \lambda] &= \nabla_{\theta} \frac{\exp(\theta^T h(\lambda, 1)) - \exp(\theta^T h(\lambda, -1))}{\exp(\theta^T h(\lambda, 1)) + \exp(\theta^T h(\lambda, -1))} \\
&= \nabla_{\theta} \frac{\exp(\theta^T h_1(\lambda)) - \exp(-\theta^T h_1(\lambda))}{\exp(\theta^T h_1(\lambda)) + \exp(\theta^T h_1(\lambda))} \\
&= \nabla_{\theta} \tanh(\theta^T h_1(\lambda)) \\
&= h_1(\lambda) (1 - \tanh^2(\theta^T h_1(\lambda))) \\
&= h_1(\lambda) \mathbf{Var}_{(\Lambda, Y) \sim \pi_{\theta}} (Y | \Lambda = \lambda).
\end{aligned}$$

It follows by the mean value theorem that for some ψ , a linear combination of θ and θ^* ,

$$|l(w) - l_{\theta}(w)| = \left| \mathbf{E}_{(\bar{x}, \bar{y}) \sim \pi^*} \left[\frac{w^T f(\bar{x})}{2} (\theta - \theta^*)^T h_1(\lambda) \mathbf{Var}_{(\Lambda, Y) \sim \pi_{\psi}} (Y | \Lambda = \lambda) \right] \right|.$$

Since Θ is convex, clearly $\psi \in \Theta$. From our assumption on the bound of the variance, we can conclude that

$$\mathbf{E}_{(\bar{x}, \bar{y}) \sim \pi^*} [\mathbf{Var}_{(\Lambda, Y) \sim \pi_{\psi}} (Y | \Lambda = \lambda)] \leq \frac{c}{M}.$$

By the Cauchy-Schwarz inequality,

$$|l(w) - l_{\theta}(w)| \leq \frac{1}{2} \left| \mathbf{E}_{(\bar{x}, \bar{y}) \sim \pi^*} [\|w\| \|f(\bar{x})\| \|\theta - \theta^*\| \|h_1(\lambda)\| \mathbf{Var}_{(\Lambda, Y) \sim \pi_{\psi}} (Y | \Lambda = \lambda)] \right|.$$

Since (by assumption) $\|f(x)\| \leq 1$ and $\|h_1(\lambda)\| \leq \sqrt{M}$,

$$\begin{aligned}
|l(w) - l_{\theta}(w)| &\leq \frac{\|w\| \|\theta - \theta^*\| \sqrt{M}}{2} \left| \mathbf{E}_{(\bar{x}, \bar{y}) \sim \pi^*} [\mathbf{Var}_{(\Lambda, Y) \sim \pi_{\psi}} (Y | \Lambda = \lambda)] \right| \\
&\leq \frac{\|w\| \|\theta - \theta^*\| \sqrt{M}}{2} \cdot \frac{c}{M} \\
&= \frac{c \|w\| \|\theta - \theta^*\|}{2 \sqrt{M}}.
\end{aligned}$$

Now, for any w that could conceivably be a solution, it must be the case that

$$\|w\| \leq \frac{1}{2\rho},$$

since otherwise the regularization term would be too large. Therefore, for any possible solution w ,

$$|l(w) - l_{\theta}(w)| \leq \frac{c \|\theta - \theta^*\|}{4\rho \sqrt{M}}.$$

Now, we apply the assumption that we are able to solve the empirical problem, producing an estimate \hat{w} that satisfies

$$\mathbf{E} [l_{\theta}(\hat{w}) - l_{\theta}(w_{\theta}^*)] \leq \chi,$$

where w_{θ}^* is the true solution to

$$w_{\theta}^* = \arg \min_w l_{\theta}(w).$$

Therefore,

$$\begin{aligned}
\mathbf{E} [l(\hat{w}) - l(w^*)] &= \mathbf{E} [l_{\theta}(\hat{w}) - l_{\theta}(w_{\theta}^*) + l_{\theta}(w_{\theta}^*) - l_{\theta}(\hat{w}) + l(\hat{w}) - l(w^*)] \\
&\leq \chi + \mathbf{E} [l_{\theta}(w^*) - l_{\theta}(\hat{w}) + l(\hat{w}) - l(w^*)] \\
&\leq \chi + \mathbf{E} [|l_{\theta}(w^*) - l(w^*)| + |l_{\theta}(\hat{w}) - l(\hat{w})|] \\
&\leq \chi + \mathbf{E} \left[\frac{c \|\theta - \theta^*\|}{2\rho \sqrt{M}} \right] \\
&= \chi + \frac{c}{2\rho \sqrt{M}} \mathbf{E} [\|\theta - \theta^*\|] \\
&\leq \chi + \frac{c}{2\rho \sqrt{M}} \sqrt{\mathbf{E} [\|\theta - \theta^*\|^2]}.
\end{aligned}$$

We can now bound this using the result of Lemma D.4, which results in

$$\begin{aligned}\mathbf{E} [l(\hat{w}) - l(w^*)] &\leq \chi + \frac{c}{2\rho\sqrt{M}} \sqrt{M\epsilon^2} \\ &= \chi + \frac{c\epsilon}{2\rho}.\end{aligned}$$

This is the desired result. \square

E Proofs of Results for the Independent Model

To restate, in the independent model, the variables are, as before, $\Lambda \in \{-1, 0, 1\}^m$ and $Y \in \{-1, 1\}$. The sufficient statistics are $\Lambda_i Y$ and Λ_i^2 . That is, for expanded parameter $\theta = (\psi, \phi)$,

$$\pi_\theta(\Lambda, Y) = \frac{1}{Z} \exp(\psi^T \Lambda Y + \phi^T \Lambda^2).$$

This can be combined with the simple assumption that $\mathbf{P}(Y) = \frac{1}{2}$ to complete a whole distribution. Using this, we can prove the following simple result about the moments of the sufficient statistics.

Lemma E.1. *The expected values and covariances of the sufficient statistics are, for all $i \neq j$,*

$$\begin{aligned}\mathbf{E} [\Lambda_i Y] &= \beta_i \gamma_i \\ \mathbf{E} [\Lambda_i^2] &= \beta_i \\ \mathbf{Var} (\Lambda_i Y) &= \beta_i - \beta_i^2 \gamma_i^2 \\ \mathbf{Var} (\Lambda_i^2) &= \beta_i - \beta_i^2 \\ \mathbf{Cov} (\Lambda_i Y, \Lambda_j Y) &= 0 \\ \mathbf{Cov} (\Lambda_i^2, \Lambda_j^2) &= 0 \\ \mathbf{Cov} (\Lambda_i Y, \Lambda_j^2) &= 0.\end{aligned}$$

We also prove the following basic lemma that relates ψ_i to γ_i .

Lemma E.2. *It holds that*

$$\gamma_i = \tanh(\psi_i).$$

We also make the following claim about feasible models.

Lemma E.3. *For any feasible model, it will be the case that, for any other feasible parameter vector $\hat{\psi}$,*

$$\mathbf{P} \left(\hat{\psi}^T \Lambda Y \leq \frac{m}{2} \gamma_{\min} (\gamma \beta)_{\min} \right) \leq \exp \left(- \frac{m(\gamma \beta)_{\min} \gamma_{\min}^2}{9.34 \operatorname{artanh}(\gamma_{\max})} \right).$$

We can also prove the following simple result about the conditional covariances

Lemma E.4. *The covariances of the sufficient statistics, conditioned on Λ , are for all $i \neq j$,*

$$\begin{aligned}\mathbf{Cov} (\Lambda_i Y, \Lambda_j Y | \Lambda) &= \Lambda_i \Lambda_j \operatorname{sech}^2(\psi^T \Lambda) \\ \mathbf{Cov} (\Lambda_i^2, \Lambda_j^2 | \Lambda) &= 0.\end{aligned}$$

We can combine these two results to bound the expected variance of these conditional statistics.

Lemma E.5. *If θ and θ^* are two feasible models, then for any u ,*

$$\mathbf{E}_{\theta^*} [\mathbf{Var}_\theta (Y | \Lambda)] \leq 3 \exp \left(- \frac{m \beta_{\min}^2 \gamma_{\min}^3}{8 \operatorname{artanh}(\gamma_{\max})} \right).$$

We can now proceed to restate and prove the main corollary of Theorem A.1 that applies in the independent case.

Corollary B.1. *Suppose that we run Algorithm 1 on an independent data programming specification that satisfies conditions (13), (14), (15), and (16). Furthermore, assume that the number of labeling functions we use satisfies*

$$m \geq \frac{9.34 \operatorname{artanh}(\gamma_{\max})}{(\gamma\beta)_{\min} \gamma_{\min}^2} \log \left(\frac{24m}{\beta_{\min}} \right).$$

Suppose further that, for some parameter $\epsilon > 0$, we use step size

$$\eta = \frac{\beta_{\min} \epsilon^2}{16}$$

and our dataset is of a size that satisfies

$$|S| = \frac{32}{\beta_{\min}^2 \epsilon^2} \log \left(\frac{2 \|\theta_0 - \theta^*\|^2}{\epsilon} \right).$$

Then, we can bound the expected parameter error with

$$\mathbf{E} \left[\|\hat{\theta} - \theta^*\|^2 \right] \leq \epsilon^2 M$$

and the expected risk with

$$\mathbf{E} \left[l(\hat{w}) - \min_w l(w) \right] \leq \chi + \frac{\beta_{\min} \epsilon}{8\rho}.$$

Proof. In order to apply Theorem A.1, we have to verify all its conditions hold in the independent case.

First, we notice that (11) is used only to bound the covariance of the sufficient statistics. From Lemma E.1, we know that these can be bounded by $\beta_i - \beta_i^2 \gamma_i^2 \geq \frac{\beta_{\min}}{2}$. It follows that we can choose

$$c = \frac{\beta_{\min}}{4},$$

and we can consider (11) satisfied, for the purposes of applying the theorem.

Second, to verify (12), we can use Lemma E.5. For this to work, we need

$$3 \exp \left(-\frac{m(\gamma\beta)_{\min} \gamma_{\min}^2}{9.34 \operatorname{artanh}(\gamma_{\max})} \right) \leq \frac{c}{M} = \frac{\beta_{\min}}{8m}.$$

This happens whenever the number of labeling functions satisfies

$$m \geq \frac{9.34 \operatorname{artanh}(\gamma_{\max})}{(\gamma\beta)_{\min} \gamma_{\min}^2} \log \left(\frac{24m}{\beta_{\min}} \right).$$

The remaining assumptions, (13), (14), (15), and (16), are satisfied directly by the assumptions of this corollary. So, we can apply Theorem A.1, which produces the desired result. \square

F Proofs of Independent Model Lemmas

Lemma E.1. *The expected values and covariances of the sufficient statistics are, for all $i \neq j$,*

$$\begin{aligned} \mathbf{E} [\Lambda_i Y] &= \beta_i \gamma_i \\ \mathbf{E} [\Lambda_i^2] &= \beta_i \\ \mathbf{Var} (\Lambda_i Y) &= \beta_i - \beta_i^2 \gamma_i^2 \\ \mathbf{Var} (\Lambda_i^2) &= \beta_i - \beta_i^2 \\ \mathbf{Cov} (\Lambda_i Y, \Lambda_j Y) &= 0 \\ \mathbf{Cov} (\Lambda_i^2, \Lambda_j^2) &= 0 \\ \mathbf{Cov} (\Lambda_i Y, \Lambda_j^2) &= 0. \end{aligned}$$

Proof. We prove each of the statements in turn. For the first statement,

$$\begin{aligned}\mathbf{E} [\Lambda_i Y] &= \mathbf{P} (\Lambda_i = Y) - \mathbf{P} (\Lambda_i = -Y) \\ &= \beta_i \frac{1 + \gamma_i}{2} - \beta_i \frac{1 - \gamma_i}{2} \\ &= \beta_i \gamma_i.\end{aligned}$$

For the second statement,

$$\begin{aligned}\mathbf{E} [\Lambda_i^2] &= \mathbf{P} (\Lambda_i = Y) + \mathbf{P} (\Lambda_i = -Y) \\ &= \beta_i \frac{1 + \gamma_i}{2} + \beta_i \frac{1 - \gamma_i}{2} \\ &= \beta_i.\end{aligned}$$

For the remaining statements, we derive the second moments; converting these to an expression of the covariance is trivial. For the third statement,

$$\mathbf{E} [(\Lambda_i Y)^2] = \mathbf{E} [\Lambda_i^2 Y^2] = \mathbf{E} [\Lambda_i^2] = \beta_i.$$

For the fourth statement,

$$\mathbf{E} [(\Lambda_i^2)^2] = \mathbf{E} [\Lambda_i^4] = \mathbf{E} [\Lambda_i^2] = \beta_i.$$

For subsequent statements, we first derive that

$$\mathbf{E} [\Lambda_i Y | Y] = \beta_i \frac{1 + \gamma_i}{2} - \beta_i \frac{1 - \gamma_i}{2} = \beta_i \gamma_i$$

and

$$\mathbf{E} [\Lambda_i^2 | Y] = \beta_i \frac{1 + \gamma_i}{2} + \beta_i \frac{1 - \gamma_i}{2} = \beta_i.$$

Now, for the fifth statement,

$$\mathbf{E} [(\Lambda_i Y)(\Lambda_j Y)] = \mathbf{E} [\mathbf{E} [\Lambda_i Y | Y] \mathbf{E} [\Lambda_j Y | Y]] = \beta_i \gamma_i \beta_j \gamma_j.$$

For the sixth statement,

$$\mathbf{E} [(\Lambda_i^2)(\Lambda_j^2)] = \mathbf{E} [\mathbf{E} [\Lambda_i^2 | Y] \mathbf{E} [\Lambda_j^2 | Y]] = \beta_i \beta_j.$$

Finally, for the seventh statement,

$$\mathbf{E} [(\Lambda_i Y)(\Lambda_j^2)] = \mathbf{E} [\mathbf{E} [\Lambda_i Y | Y] \mathbf{E} [\Lambda_j^2 | Y]] = \beta_i \gamma_i \beta_j.$$

This completes the proof. \square

Lemma E.2. *It holds that*

$$\gamma_i = \tanh(\psi_i).$$

Proof. From the definitions,

$$\beta_i = \frac{\exp(\psi_i + \phi_i) + \exp(-\psi_i + \phi_i)}{\exp(\psi_i + \phi_i) + \exp(-\psi_i + \phi_i) + 1}$$

and

$$\beta_i \gamma_i = \frac{\exp(\psi_i + \phi_i) - \exp(-\psi_i + \phi_i)}{\exp(\psi_i + \phi_i) + \exp(-\psi_i + \phi_i) + 1}.$$

Therefore,

$$\gamma_i = \frac{\exp(\psi_i + \phi_i) - \exp(-\psi_i + \phi_i)}{\exp(\psi_i + \phi_i) + \exp(-\psi_i + \phi_i)} = \tanh(\psi_i),$$

which is the desired result. \square

Lemma E.3. *For any feasible model, it will be the case that, for any other feasible parameter vector $\hat{\psi}$,*

$$\mathbf{P} \left(\hat{\psi}^T \Lambda Y \leq \frac{m}{2} \gamma_{\min} (\gamma \beta)_{\min} \right) \leq \exp \left(- \frac{m (\gamma \beta)_{\min} \gamma_{\min}^2}{9.34 \operatorname{artanh}(\gamma_{\max})} \right).$$

Proof. We start by noticing that

$$\hat{\psi}^T \Lambda Y = \sum_{i=1}^m \hat{\psi}_i \Lambda_i Y.$$

Since in this model, all the $\Lambda_i Y$ are independent of each other, we can bound this sum using a concentration bound. First, we note that

$$|\hat{\psi}_i \Lambda_i Y| \leq \hat{\psi}_i.$$

Second, we note that

$$\mathbf{E} [\hat{\psi}_i \Lambda_i Y] = \hat{\psi}_i \beta_i \gamma_i$$

and

$$\mathbf{Var} (\hat{\psi}_i \Lambda_i Y) = \hat{\psi}_i^2 (\beta_i - \beta_i^2 \gamma_i^2)$$

but

$$|\hat{\psi}_i \Lambda_i Y| \leq \hat{\psi}_i \leq \text{artanh}(\gamma_{\max}) \triangleq \hat{\psi}_{\max}$$

because, for feasible models, by definition

$$\gamma_{\min} \leq \text{artanh}(\gamma_{\min}) \leq \hat{\psi}_i \leq \text{artanh}(\gamma_{\max}).$$

Therefore, applying Bernstein's inequality gives us, for any t ,

$$\mathbf{P} \left(\sum_{i=1}^m \hat{\psi}_i \Lambda_i Y - \sum_{i=1}^m \hat{\psi}_i \beta_i \gamma_i \leq -t \right) \leq \exp \left(- \frac{3t^2}{6 \sum_{i=1}^m \hat{\psi}_i^2 \gamma_i \beta_i \gamma_i + 2 \hat{\psi}_{\max} t} \right).$$

It follows that, if we let

$$t = \frac{1}{2} \sum_{i=1}^m \hat{\psi}_i \beta_i \gamma_i,$$

then we get

$$\begin{aligned} \mathbf{P} \left(\sum_{i=1}^m \hat{\psi}_i \Lambda_i Y - \sum_{i=1}^m \hat{\psi}_i \beta_i \gamma_i \leq -t \right) &\leq \exp \left(- \frac{3 \left(\frac{1}{2} \sum_{i=1}^m \hat{\psi}_i \beta_i \gamma_i \right)^2}{6 \sum_{i=1}^m \hat{\psi}_i^2 \gamma_i \beta_i \gamma_i + 2 \hat{\psi}_{\max} \left(\frac{1}{2} \sum_{i=1}^m \hat{\psi}_i \beta_i \gamma_i \right)} \right) \\ &\leq \exp \left(- \frac{3 \sum_{i=1}^m \hat{\psi}_i \beta_i \gamma_i}{24 \gamma_{\max} \hat{\psi}_{\max} + 4 \hat{\psi}_{\max}} \right) \\ &\leq \exp \left(- \frac{3m(1 - \gamma_{\max})}{28 \hat{\psi}_{\max}} \right) \\ &\leq \exp \left(- \frac{3 \left(\sum_{i=1}^m \hat{\psi}_i \beta_i \gamma_i \right)^2}{24 \sum_{i=1}^m \hat{\psi}_i^2 \beta_i + 4 \hat{\psi}_{\max} \left(\sum_{i=1}^m \hat{\psi}_i \beta_i \gamma_i \right)} \right) \\ &\leq \exp \left(- \frac{3 \gamma_{\min} \left(\sum_{i=1}^m \hat{\psi}_i \beta_i \right) \left(\sum_{i=1}^m \hat{\psi}_i \beta_i \gamma_i \right)}{24 \hat{\psi}_{\max} \sum_{i=1}^m \hat{\psi}_i \beta_i + 4 \hat{\psi}_{\max} \left(\sum_{i=1}^m \hat{\psi}_i \beta_i \right)} \right) \\ &\leq \exp \left(- \frac{3 \gamma_{\min} \left(\sum_{i=1}^m \hat{\psi}_i \beta_i \gamma_i \right)}{28 \hat{\psi}_{\max}} \right) \\ &\leq \exp \left(- \frac{m \gamma_{\min}^2 (\gamma \beta)_{\min}}{9.34 \hat{\psi}_{\max}} \right). \end{aligned}$$

This is the desired expression. □

Lemma E.4. *The covariances of the sufficient statistics, conditioned on Λ , are for all $i \neq j$,*

$$\begin{aligned} \mathbf{Cov} (\Lambda_i Y, \Lambda_j Y | \Lambda) &= \Lambda_i \Lambda_j \text{sech}^2(\psi^T \Lambda) \\ \mathbf{Cov} (\Lambda_i^2, \Lambda_j^2 | \Lambda) &= 0. \end{aligned}$$

Proof. The second result is obvious, so it suffices to prove only the first result. Clearly,

$$\mathbf{Cov}(\Lambda_i Y, \Lambda_j Y | \Lambda) = \Lambda_i \Lambda_j \mathbf{Var}(Y | \Lambda) = \Lambda_i \Lambda_j (1 - \mathbf{E}[Y | \Lambda]^2).$$

Plugging into the distribution formula lets us conclude that

$$\mathbf{E}[Y | \Lambda] = \frac{\exp(\psi^T \Lambda + \phi^T \Lambda^2) - \exp(-\psi^T \Lambda + \phi^T \Lambda^2)}{\exp(\psi^T \Lambda + \phi^T \Lambda^2) + \exp(-\psi^T \Lambda + \phi^T \Lambda^2)} = \tanh^2(\psi^T \Lambda),$$

and so

$$\mathbf{Cov}(\Lambda_i Y, \Lambda_j Y | \Lambda) = \Lambda_i \Lambda_j (1 - \tanh^2(\psi^T \Lambda)) = \Lambda_i \Lambda_j \operatorname{sech}^2(\psi^T \Lambda),$$

which is the desired result. \square

Lemma E.5. *If θ and θ^* are two feasible models, then for any u ,*

$$\mathbf{E}_{\theta^*}[\mathbf{Var}_{\theta}(Y | \Lambda)] \leq 3 \exp\left(-\frac{m\beta_{\min}^2 \gamma_{\min}^3}{8 \operatorname{artanh}(\gamma_{\max})}\right).$$

Proof. First, we note that, by the result of Lemma E.4,

$$\mathbf{Var}_{\theta}(Y | \Lambda) = \operatorname{sech}^2(\psi^T \Lambda).$$

Therefore,

$$\mathbf{E}_{\theta^*}[\mathbf{Var}_{\theta}(Y | \Lambda)] = \mathbf{E}_{\theta^*}[\operatorname{sech}^2(\psi^T \Lambda)].$$

Applying Lemma E.3, we can bound this with

$$\begin{aligned} \mathbf{E}_{\theta^*}[\mathbf{Var}_{\theta}(u^T \Lambda Y | \Lambda)] &\leq \left(\operatorname{sech}^2\left(\frac{m}{2}(\gamma\beta)_{\min} \gamma_{\min}^2\right) + \exp\left(-\frac{m(\gamma\beta)_{\min} \gamma_{\min}^2}{9.34 \operatorname{artanh}(\gamma_{\max})}\right) \right) \\ &\leq \left(2 \exp\left(-\frac{m}{2}(\gamma\beta)_{\min} \gamma_{\min}^2\right) + \exp\left(-\frac{m(\gamma\beta)_{\min} \gamma_{\min}^2}{9.34 \operatorname{artanh}(\gamma_{\max})}\right) \right) \\ &\leq 3 \exp\left(-\frac{m(\gamma\beta)_{\min} \gamma_{\min}^2}{9.34 \operatorname{artanh}(\gamma_{\max})}\right). \end{aligned}$$

This is the desired expression. \square

G Additional Experimental Details

G.1 Relation Extraction Experiments

G.1.1 Systems

The original distantly-supervised experiments which we compare against as baselines—which we refer to as using the *if-then-return (ITR)* approach of distant or programmatic supervision—were implemented using DeepDive, an open-source system for building extraction systems.⁵ For our primary experiments, we adapted these programs to the framework and approach described in this paper, directly utilizing distant supervision rules as labeling functions.

In the disease tagging user experiments, we used an early version of our new lightweight extraction framework based around data programming, formerly called DDLite [12], now Snorkel.⁶ Snorkel is based around a Jupyter-notebook based interface, allowing users to iteratively develop labeling functions in Python for basic extraction tasks involving simple models. Details of the basic discriminative models used can be found in the Snorkel repository; in particular, Snorkel uses a simple logistic regression model with generic features defined in part over dependency paths⁷, and a basic LSTM model implemented using the Theano library.⁸ Snorkel is currently under continued development, and all versions are open-source.

⁵<http://deepdive.stanford.edu>

⁶<http://snorkel.stanford.edu>

⁷<https://github.com/HazyResearch/treedlib>

⁸<http://deeplearning.net/software/theano/>

G.1.2 Applications

We consider three primary applications which involve the extraction of binary relation mentions of some specific type from unstructured text input data. At a high level, all three system pipelines consist of an initial *candidate extraction* phase which leverages some upstream model or suite of models to extract mentions of involved entities, and then considers each pair of such mentions that occurs within the same local neighborhood in a document as a *candidate relation mention* to be potentially extracted. In each case, the discriminative model that we are aiming to train—and that we evaluate in this paper—is a binary classifier over these candidate relation mentions, which will decide which ones to output as final true extractions. In all tasks, we preprocessed raw input text with Stanford CoreNLP⁹, and then either used CoreNLP’s NER module or our own entity-extraction models to extract entity mentions. Further details of the basic information extraction pipeline utilized can be seen in the tutorials of the systems used, and in the referenced papers below.

In the 2014 TAC-KBP Slot Filling task, which we also refer to as the News application, we train a set of extraction models for a variety of relation types from news articles [30]. In reported results in this paper, we average over scores from each relation type. We utilized CoreNLP’s NER module for candidate extraction, and utilized CoreNLP outputs in developing the distant supervision rules / labeling functions for these tasks. We also considered a slightly simpler discriminative model than the one submitted in the 2014 competition, as reported in [2]: namely, we did not include any joint factors in our model in this paper.

In the Genomics application, our goal with our collaborators at Stanford Medicine was to extract mentions of genes that if mutated may cause certain phenotypes (symptoms) linked to Mendelian diseases, for use in a clinical diagnostic setting. The code for this project is online, although it remains partially under development and thus some material from our collaborators is private.¹⁰

In the Pharmacogenomics application, our goal was to extract interactions between genes for use in downstream pharmacogenomics research analyses; full results and system details are reported in [21].

In the Disease Tagging application, which we had our collaborators work on during a set of short hackathons as a user study, the goal was to tag mentions of human diseases in PubMed abstracts. We report results of this hackathon in [12], as well as in our Snorkel tutorial online.

G.1.3 Labeling Functions

In general, we saw two broad types of labeling functions in both prior applications (when they were referred to as “distant supervision rules”) and in our most recent user studies. The first type of labeling function leverages some weak supervision signal, such as an external knowledgebase (as in traditional distant supervision), very similar to the example illustrated in Fig. 1(a). All of the applications studied in this paper used some such labeling function or set of labeling functions.

The second type of labeling function uses simple heuristic patterns as positive or negative signals. For our text extraction examples, these heuristic patterns primarily consisted of regular expressions, also similar to the example pseudocode in Fig. 1(a). Further specific details of both types of labeling functions, as well as others used, can be seen in the linked code repositories and referenced papers.

G.2 Synthetic Experiments

In Fig. 3(a-b), we ran synthetic experiments with labeling functions having constant coverage $\beta = 0.1$, and accuracy drawn from $\alpha \sim \text{Uniform}(\mu_\alpha - 0.25, \mu_\alpha + 0.25)$ where $\mu_\alpha = 0.75$ in the above plots. In both cases we used 1000 normally-drawn features having mean correlation with the true label class of 0.5.

In this case we compare data programming (DP-Pipelined) against two baselines. First, we compare against an *if-then-return* setup where the ordering is optimal (ITR-Oracle). Second, we compare against simple majority vote (MV).

In Fig. 3(c), we show an experiment where we add dependent labeling functions to a set of $m_{ind} = 50$ independent labeling functions, and either provided this dependency structure (LDM-Aware) or did

⁹stanfordnlp.github.io/CoreNLP/

¹⁰<https://github.com/HazyResearch/dd-genomics>




Figure 3: Comparisons of data programming to two oracle methods on synthetic data.

not (Independent). In this case, the independent labeling functions had the same configurations as in (a-b), and the dependent labeling functions corresponded to “fixes” or “reinforces”-type dependent labeling functions.