
An Integer Linear Programming Framework for Mining Constraints from Data

Tao Meng¹ Kai-Wei Chang¹

Abstract

Structured output prediction problems (e.g., sequential tagging, hierarchical multi-class classification) often involve constraints over the output label space. These constraints interact with the learned models to filter infeasible solutions and facilitate in building an accountable system. However, although constraints are useful, they are often based on hand-crafted rules. This raises a question – *can we mine constraints and rules from data based on a learning algorithm?*

In this paper, we present a general framework for mining constraints from data. In particular, we consider the inference in structured output prediction as an integer linear programming (ILP) problem. Then, given the coefficients of the objective function and the corresponding solution, we mine the underlying constraints by estimating the outer and inner polytopes of the feasible set. We verify the proposed constraint mining algorithm in various synthetic and real-world applications and demonstrate that the proposed approach successfully identifies the feasible set at scale. In particular, we show that our approach can learn to solve 9x9 Sudoku puzzles and minimal spanning tree problems from examples without providing the underlying rules. Our algorithm can also integrate with a neural network model to learn the hierarchical label structure of a multi-label classification task. Besides, we provide a theoretical analysis about the tightness of the polytopes and the reliability of the mined constraints.

1. Introduction

A variety of machine learning problems involve making coherent decisions over a set of output variables, where the dependencies between them can be described by con-

straints (Punyakanok et al., 2005; Samdani & Roth, 2012; Nowozin & Lampert, 2011). For example, in part-of-speech tagging, a constraint specifying that every sentence should contain at least one verb and one noun can greatly improve the performance (Ganchev et al., 2010). Similarly, in hierarchical multi-label classification, a figure labeled ‘flower’ should also be labeled ‘plant’ as well (Dimitrovski et al., 2011). To incorporate constraints with learned models, one popular method is to formulate the inference problem into an integer linear programming (ILP) (Roth & Yih, 2004). This framework is general and can cope with constraints formed as propositional logics (Hooker, 1988; Richardson & Domingos, 2006). This approach has been widely used in natural language processing, computer vision, and many application areas. It has demonstrated great performance gains in various applications (e.g., Martins et al. (2010); Nowozin & Lampert (2011); Roth & Yih (2004); Goldwasser et al. (2012); Chang et al. (2008); Meng et al. (2019)).¹

In the literature, existing works most focus on how to utilize constraints to facilitate learning. They mostly assume constraints are given a priori. However, for some applications, manually identifying constraints is tedious. Besides, some constraints are obscure and cannot be easily identified by human experts.² Inspired by representation learning methods automate feature extraction, we envision that *an artificial intelligence system that could automatically recognize underlying constraints among output labels from data and incorporate them in the prediction time.*

To illustrate the goal, consider the following learning problem. We are given a set of input-output pairs as training data, where each input is an adjacency matrix of a graph representing the distances between nodes and the output is corresponding minimal spanning tree (MST). Our goal is to train a model to generate MST of a given graph (adjacency matrix) *without* telling the model that the output is a

¹Solving ILP is in general NP-hard. However, in practice, inference problems often can be solved efficiently using a commercial ILP solver or an approximation inference technique (e.g., LP-relaxation (Fromer & Globerson, 2009), loopy belief propagation (Murphy et al., 2013)). See discussion in (Finley & Joachims, 2008).

²For example, if we shuffle columns of all sudoku puzzles with the same order, the puzzles still follow a set of constraints. However, it is hard for humans to recognize these underlying rules.

¹Department of Computer Science, University of California, Los Angeles, USA. Correspondence to: Tao Meng <tmeng@cs.ucla.edu>, Kai-Wei Chang <kwchang@cs.ucla.edu>.

tree. Specifically, the model has to identify the underlying constraints that are satisfied by all training samples from a family of candidate constraints. The success in this problem has a great potential; however, there are limited prior works except some methods extending the basic Valiant’s algorithm (Valiant, 1984), such as inductive logic programming (Muggleton & Raedt, 1994; Riedel & Clarke, 2006) and constraint learning (Bessiere et al., 2013; 2016; 2017). Most of them use logic clauses to formulate the constraints and solve the satisfiability problem. However, it is unclear how to incorporate them with machine learning models.

Inspired by the great success of ILP in constrained output structure predictions, we propose a novel framework to formulate the constraint learning based on ILP. In particular, we estimate the feasible set defined by the constraints and explore three techniques: 1) mining inequality constraints to form a superset of the feasible set by constructing an outer polytope based on seen data; 2) mining equality constraints by dimension reduction of the superset; and 3) mining complex constraints with a latent variable method. We also propose an algorithm to induce the subset of the feasible set for evaluating the quality of the constraints. Note that although the constraint mining algorithm is designed under the ILP framework, our algorithm does not involve solving ILP when mining the constraints.

We evaluate the proposed framework on three tasks: MST, Sudoku, and hierarchical multi-label classification. The first two tasks demonstrate that our method is able to mine complex structures and deal with large label space. For example, in Sudoku, our model can perfectly learn the underlying rules and achieve 100% accuracy. We then incorporate the proposed approach with a learned neural network on hierarchical multi-label classification. Our framework helps models learn the structure in label space and improve the performance by over 10% compared with the baseline. Finally, we conduct a comprehensive analysis on MST. We verify the constraints learned by our approaches by comparing the corresponding feasible set with the ground truth. We also provide a theoretical estimation on the feasible set size and compare it with the empirical results and discuss the running time of the algorithm. The source code and data are available at <https://github.com/uclanlp/ILPLearning>.

2. Related Work

Constraints Formulated by Integer Linear Programming ILP is widely used in formulating constrained inference in machine learning tasks, including semantic role labeling (Punyakank et al., 2004), entity-relation extraction (Roth & Yih, 2005), sentence compression (Clarke & Lapata, 2008), dependency parsing (Martins et al., 2009), multi-lingual transfer (Meng et al., 2019), coreference res-

olution (Chang et al., 2013), relation extraction (Ye et al., 2020) and reducing bias amplification (Zhao et al., 2017). These works use pre-defined constraints to formulate ILPs. In contrast, we aim to mine constraints from data.

Mining Constraints From Data Raedt et al. (2018) summarize the milestones in constraint mining. Learning logical rules from data can be traced back to the Valiant’s algorithm (Valiant, 1984) that mines the hard constraints formulated as k -CNF. Inductive logic programming (Muggleton & Raedt, 1994; Riedel & Clarke, 2006), as an extension of Valiant’s Algorithm, is aiming to deal with general first-order logic. It has been used in both real world (Bratko & King, 1994) and mathematical applications (Colton & Muggleton, 2006). Constraint learning (Bessiere et al., 2013; 2016; 2017) combines these two approaches together. Besides, several efforts have been put on relaxing logical constraints such as soft constraint learning (Rossi & Sperduti, 2004). Wang et al. (2019) use semidefinite programming (SDP) (Wang & Kolter, 2019) to relax the maxSAT problem and cooperate with deep learning (see Sec. 4 for comparison). Another way to relax the logical constraints is to relax the Boolean variables to be continuous variables (Li et al., 2019; Li & Srikumar, 2019), or continuous random variables like probabilistic soft logic (Kimmig et al., 2012; Bach et al., 2015; Embar et al., 2018). Most of these previous works use logics to represent the constraints. In contrast, we design a framework based on ILP and use the linear form to formulate the constraints. This allows us directly incorporate constraints with inference in structured output predictions. Some concurrent works (Pan et al., 2020; Tan et al., 2020) learn constraints by initializing a set of constraints and updating them based on gradient, but it is not able to obtain the guarantee that the constraints are tight and converge to the ground truth as ours do (see discussion in Sec. 3 and Sec. 5.1).

3. Mining Constraints with Integer Linear Programming

We first review the constraint mining framework based on ILP. We then propose an approach to mine constraints by estimating the outer and inner polytopes of the feasible set. Finally, we discuss how to extend the framework to capture complex constraints.

ILP is a linear optimization problem with linear constraints and the values of variables are restricted to integers. Formally, the ILP problem can be formulated as

$$\max_{\mathbf{y} \in \mathbb{Z}^d} \mathbf{w} \cdot \mathbf{y} \quad \text{s.t.} \quad \mathbf{A}\mathbf{y} \leq \mathbf{b}, \quad (1)$$

where $\mathbf{w} \in \mathbb{R}^d$ is the coefficients of the objective function (a.k.a. weights) and \mathbf{y} is an integer vector that encodes




Figure 1: The pentagons (blue solid line) in Fig. 1a and 1b show the outer and the inner polytopes of 5 training samples $\{\mathbf{w}_i, \mathbf{y}_i\}_{i=1}^5$ (see Sec. 3.1). The dashed red line shows the boundary of the feasible set (yellow region). We also show w_i as the normal vector of the outer line.

the output label.³ The matrix A and vector \mathbf{b} specify the constraints. We use S^* to denote the feasible set defined by the constraints. Various structure prediction problems can be casted into the ILP formulation. For example, dependency parsing can be formulated as finding the maximum spanning tree in a directed graph (McDonald et al., 2005), where each node represents a word and the edge w_{ij} represents how likely the word i is the dependent of the word j predicted by a model. $\mathbf{y} = \{\mathbf{y}_{ij}\}, \mathbf{y}_{ij} \in \{0, 1\}$ is the indicator of the edges in the resulting tree. The objective in Eq. (1) then can be interpreted as the total score of edges in \mathbf{y} , and the constraints, described by (A, \mathbf{b}) , restrict \mathbf{y} to be a tree (Martins et al., 2009).

Prior works (see, e.g., (Martins et al., 2009)) mostly assume the constraints (A, \mathbf{b}) are given. However, in this paper, we assume (A, \mathbf{b}) are unknown and our goal is to identify the underlying feasible set S^* spanned by (A, \mathbf{b}) using a set of objective-solution pairs $\{(\mathbf{w}^{(i)}, \mathbf{y}^{(i)})\}_{i=1}^k$ that satisfy constraints defined by (A, \mathbf{b}) . For example, in MST, giving a set of weights $\mathbf{w}^{(i)}$ (adjacency matrix) with the corresponding optimal solution $\mathbf{y}^{(i)}$, our algorithm identifies the structure of the output \mathbf{y} form a tree structure.

In the following, we introduce algorithms to estimate the feasible set for mining the underlying constraints. These mined constraints define an superset of the feasible set. We also design an algorithm to get the subset of the feasible set to evaluate the estimation.

3.1. Mining Inequality Constraints

In the following, we discuss how to estimate the underlying feasible set S^* associated with inequality constraints. Our approach finds a convex hull S_O defined by a set of learned inequality constraints that is an outer polytope (i.e., super-

³In structure output prediction, usually each element of \mathbf{y} takes value 1 or 0, indicating if a specific value is assign to a specific output variable or not. \mathbf{w} are the scores of sub-components of output assigned by a model.

set) of the feasible set S^* . We also propose a method to get an inner polytope S_I that is a subset of S^* and use the gap between the S_O and S_I to estimate the quality of approximation. Figure 1 shows an example about S_I, S_O defined by 5 training samples in a 2-dimensional space. We denote $S_I^{(i)}, S_O^{(i)}$ as the inner and outer polytopes after considering the first i samples.

Outer polytope We first introduce how to identify S_O . Assume that we already know part of the constraints A', \mathbf{b}' . We initialize the outer polytope as $S_O^{(0)} = \{\mathbf{y} \in \mathbb{Z}^d \mid A'\mathbf{y} \leq \mathbf{b}'\}$ (if A', \mathbf{b}' are empty, $S_O^{(0)} = \mathbb{Z}^d$). For every training sample $(\mathbf{w}^{(i)}, \mathbf{y}^{(i)})$, we consider adding the following constraint to the outer polytope

$$\mathbf{w}^{(i)} \cdot \mathbf{y} \leq \mathbf{w}^{(i)} \cdot \mathbf{y}^{(i)}. \quad (2)$$

Since $\mathbf{y}^{(i)}$ is the optimal solution under weight $\mathbf{w}^{(i)}$, all the points in the feasible set must sit in the half-space defined by Eq. (2), otherwise $\mathbf{y}^{(i)}$ is not the optimal solution. We have $S_O^{(i)} = \{\mathbf{y} \in S_O^{(0)} \mid \mathbf{w}^{(j)} \cdot \mathbf{y} \leq \mathbf{w}^{(j)} \cdot \mathbf{y}^{(j)}, j = 1, 2, \dots, i\}$, and $S^* \subseteq S_O = S_O^{(k)} \subseteq \dots \subseteq S_O^{(1)} \subseteq S_O^{(0)}$.

The outer polytope $S_O^{(i)}$ (the upper bound of the feasible set S^*) is tight when we only observe the first i samples. That is, assuming $S_O^{(i)}$ is the feasible set, if we query $\mathbf{w}^{(1)}, \dots, \mathbf{w}^{(i)}$, we will find $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(i)}$ are (one of) the optimal solutions. Therefore, $S_O^{(i)}$ is a possible feasible set. Since $S^* \subseteq S_O^{(i)}$, this bound is tight. This shows that without any further assumption, we cannot do better than S_O for estimating the outer polytope of the feasible set S^* .

In the test time, we are requested to conduct inference with unseen input weight $\mathbf{w}^{(q)}$. Since all constraints in S_O are linear, we solve the following ILP problem

$$\begin{aligned} \max_{\mathbf{y} \in \mathbb{Z}^d} \quad & \mathbf{w}^{(q)} \cdot \mathbf{y} \\ \text{s.t.} \quad & [A'^T \quad \mathbf{w}^{(1)} \quad \dots \quad \mathbf{w}^{(k)}]^T \mathbf{y} \\ & \leq [\mathbf{b}'^T \quad \mathbf{w}^{(1)} \cdot \mathbf{y}^{(1)} \quad \dots \quad \mathbf{w}^{(k)} \cdot \mathbf{y}^{(k)}]^T. \end{aligned} \quad (3)$$

The objective value of the solution of Eq. (3) might be higher than the optimum as the solution might not satisfy all the underlying constraints. We will show that empirically the outer polytope can approximate the feasible set effectively in Sec. 4. Although the number of constraints grows linearly with the number of training samples, we find that empirically the inference time does not grow much.⁴

⁴In structure output prediction, constraints are often associated with only the problem structure. Therefore, all the inference instances share the same constraint set, and the overhead in solving ILPs is amortized (Srikumar et al., 2012; Kundu et al., 2013; Chang et al., 2015).

Inner polytope To understand the quality of S_O , we also construct the inner polytope S_I , then we can use the gap between S_O and S_I to estimate the quality of the approximation. We first initialize $S_I^{(0)} = \emptyset$. For every training sample i : $(\mathbf{w}^{(i)}, \mathbf{y}^{(i)})$, we set $S_I^{(i)} = \text{convex_hull}(\{\mathbf{y}^{(1)}, \mathbf{y}^{(2)}, \dots, \mathbf{y}^{(i)}\})$, and then $S_I^{(i-1)} \subseteq S_I^{(i)}$. Since all $\{\mathbf{y}^{(i)}\}$ are in the feasible set that is convex, all the convex hulls must be subsets of the feasible set. Therefore, we have $S_I^{(0)} \subseteq S_I^{(1)} \subseteq \dots \subseteq S_I^{(k)} = S_I \subseteq S^*$.

Similarly, we can prove that $S_I^{(i)}$ is a possible feasible set after observing the first i samples, which means as a lower bound, $S_I^{(i)}$ is also tight.

When we conduct inference with S_I , we examine every vertex of the convex hull and choose the one with the optimal objective. Since it is an inner polytope of the feasible set, the solution is guaranteed to satisfy all constraints, and the objective value can be lower than the optimum. Although inner polytope and outer polytope methods are two separate algorithms, the gap between their objective function value and the size of the feasible set provide an estimation of the tightness of the bound.

In Sec. 5.1 we will show that, empirically, this approach converges with reasonable number of training samples and running time is discussed in Sec. 5.2.

Dealing with predicted weights When we incorporate the proposed approach with a structured prediction model, the weights \mathbf{w} are predicted by a base model. In this situation, the predicted weights \mathbf{w} can be noisy and the corresponding label \mathbf{y} may not be the optimal solution to Eq. (3). As the result, the outer polytope may not contain some feasible solutions as they are filtered out later by the algorithm. To handle the noise, we adapt Eq. (2) to

$$\mathbf{w}^{(i)} \cdot \mathbf{y} \leq \mathbf{w}^{(i)} \cdot \mathbf{y}^{(i)} + \xi_i, \quad i \in [k], \quad (4)$$

where ξ_i is a slack variable to ensure every training point $\mathbf{y}_i \in S^*$ satisfies Eq. (4)

$$\xi_i = \min_{j \in [k]} \{\mathbf{w}^{(i)} \cdot \mathbf{y}^{(j)} - \mathbf{w}^{(i)} \cdot \mathbf{y}^{(i)}\}.$$

3.2. Mining Equality Constraints

When there are equality constraints in the output label space. Effectively, the dimension of the output space is reduced. However, the dimension of the set S_O is the same as that of \mathbf{w} and \mathbf{y} . Therefore, this inspires us to find the sub-space of S_O to further tighten the feasible set.

For example, in the MST problem the number of edges we select is exact $N - 1$ where N is the number of nodes. Formally, the linear constraint $\mathbf{1} \cdot \mathbf{y} = N - 1$ holds for every feasible point \mathbf{y} .

We denote this d' -dimensional affine sub-space as $S_D = \{\mathbf{y} \mid W_{eq} \cdot \mathbf{y} = \mathbf{c}\}$. We can obtain W_{eq}, \mathbf{c} by solving the kernel of $[\mathbf{y}^T, \mathbf{1}]$, which is

$$\begin{bmatrix} \mathbf{y}_1^T & \mathbf{y}_2^T & \dots & \mathbf{y}_n^T \\ 1 & 1 & \dots & 1 \end{bmatrix}^T \begin{bmatrix} W_{eq}^T \\ -\mathbf{c}^T \end{bmatrix} = \mathbf{0}. \quad (5)$$

The intersection of S_O and S_D is used to replace S_O as the outer polytope of the feasible set: $S_I \subseteq S^* \subseteq S_O \cap S_D = \{\mathbf{y} \in S_O \mid W_{eq} \cdot \mathbf{y} = \mathbf{c}\}$. For the reliability of this algorithm, we give two lemmas:

Lemma 1: Given training data $\{(\mathbf{w}^{(i)}, \mathbf{y}^{(i)})\}_{i=1}^M$ that satisfy the ILP constraints set defined in Eq.(1), if the ILP contains an equality constraint $W_{eq} \cdot \mathbf{y} = \mathbf{c}$, our equality constraint mining algorithm can identify it.

Proof Sketch: For any underlying equality constraint $W_{eq} \cdot \mathbf{y} = \mathbf{c}$, all the labels of training points \mathbf{y}_i should satisfy it and it must be in the kernel of Eq. (5).

Lemma 2: For an equality constraint given by this algorithm, the probability that this constraint does not hold for the optimal solution of a random query is less than $\frac{1}{eM}$, where M is the number of training points.

Proof Sketch: We use D_w to denote the domain of the query (weights), and $f^*(\mathbf{w})$ as the ground truth solution for the query \mathbf{w} . We denote the equality constraint learned by the algorithm is $g(\mathbf{y}) = 0$. Therefore, all the data \mathbf{y}_i in the training data satisfy $g(\mathbf{y}_i) = 0$. We let $p = Pr_{\mathbf{w} \sim D_w}(g(f^*(\mathbf{w})) = 0)$. The probability of all the training data satisfying $g(\mathbf{y}_i) = 0$ is p^M . Thus, the probability that this constraint does not hold for a random query solution is

$$p^M(1-p) \leq \frac{M^M}{(M+1)^{M+1}} < \frac{1}{eM}.$$

3.3. Latent Variables

Some prediction problems involve constraints with complex logics and require auxiliary variables to model the problem structure. Thanks to the flexibility of the ILP framework, we can introduce latent variables to extend the expressiveness of the constraint mining framework:

$$\max_{\mathbf{y} \in \mathbb{Z}^d} \quad \mathbf{w} \cdot \mathbf{y} \quad \text{s.t.} \quad A_{pre} \begin{bmatrix} \mathbf{y} \\ \mathbf{h} \end{bmatrix} \leq \mathbf{b}_{pre}, \quad A \begin{bmatrix} \mathbf{y} \\ \mathbf{h} \end{bmatrix} \leq \mathbf{b}, \quad (6)$$

where \mathbf{h} are latent variables, and they appear in the constraints but not in the objective function in Eq. (6). Despite that \mathbf{h} is not part of the output, it facilitates to formulate the ILP problem. In general, a set of pre-defined constraints $(A_{pre}, \mathbf{b}_{pre})$ are given to describe the relations between \mathbf{y}

and \mathbf{h} . Then, given a set of $\{(\mathbf{w}^{(i)}, \mathbf{y}^{(i)})\}_{i=1}^k$, our goal is to learn the constraints (A, \mathbf{b}) .

The latent variables can help us formulate the constraints better in the ILP framework. Specifically, with the help of the latent variables, some inequality constraints can be reformulated as the equality ones. As it is easier to identify equality constraints in our framework, this will make the constraints we learn more accurate. We adapt the method in Sec 3.2 to solve the kernel of the matrix $[\mathbf{y}^T, \mathbf{h}^T, \mathbf{1}]^T$. In this way, we can mine equality constraints with respect to \mathbf{h} and then derive the outer polytope S_O . Since \mathbf{h} is determined by the variables \mathbf{y} , adding constraints on \mathbf{h} also reduces the size of S_O .

For example, consider multi-label classification with output \mathbf{y} , where $y_i, i = 1 \dots m$ is a binary indicator of class i . If we would like to identify the constraints between pairs of labels from $\{(\mathbf{w}^{(i)}, \mathbf{y}^{(i)})\}$, we can introduce a set of latent variables $\{h_{i,j,b_1,b_2}\}_{i,j=1\dots m; b_1,b_2 \in \{0,1\}}$ with pre-defined equality constraints $h_{i,j,b_1,b_2} = (y_i = b_1) \wedge (y_j = b_2), \forall i, j, b_1, b_2$. They can be further formulated as

$$\forall i, j, \begin{cases} h_{i,j,1,0} + h_{i,j,1,1} = y_i \\ h_{i,j,0,1} + h_{i,j,1,1} = y_j \\ \sum_{b_1,b_2 \in \{0,1\}} h_{i,j,b_1,b_2} = 1. \end{cases}$$

By introducing h_{i,j,b_1,b_2} , we are able to capture some correlations between labels better. For example, label i and label j cannot be positive at the same time can be represented by an equality constraint $h_{i,j,1,1} = 0$. Without latent variables, it can only be represented by inequality constraint $y_i + y_j \leq 1$. In this case, the introducing of latent variables make the learned constraints more accurate.

4. Experiments

We experiment on two synthetic problems, 9×9 Sudoku and minimal spanning tree (MST), to show that the proposed methods can capture different kinds of constraints. We then incorporate the proposed technique with a feed-forward neural network in a hierarchical multi-label classification problem. For all the experiments, we use the Gurobi v8.1.1 (Gurobi Optimization, 2019) as the ILP solver.⁵

4.1. Sudoku

In Sudoku, given a 9×9 grid with numbers partially filled in, the player is requested to fill in the remaining of the grid with constraints that each 3×3 sub-grid, each column and each row must contain all the numbers of $1, 2, \dots, 9$. The size of the feasible set is 6.67×10^{21} (Felgenhauer &

⁵We configure the ILP solver such that it outputs optimal solution (i.e., set $MIPGap = 0$). The relaxed LP will be discussed in Sec. 5.2.

Table 1: Sudoku results. EQ stands for equality constraint mining. Performance is reported in entry-level accuracy. Our approaches can successfully identify the underlying Sudoku constraints.

(a) Original Sudoku		
Model	Train	Test
ConvNet (Park, 2018)	72.6%	0.04%
ConvNetMask (Park, 2018)	91.4%	15.1%
SATNet (Wang et al., 2019)	99.8%	98.3%
Outer + EQ (ours)	100%	100%
(b) Permuted Sudoku		
Model	Train	Test
ConvNet (Park, 2018)	0%	0%
ConvNetMask (Park, 2018)	0.01%	0%
SATNet (Wang et al., 2019)	99.7%	98.3%
Outer + EQ (ours)	100%	100%

Jarvis, 2005) which is extremely large. Our goal is to use the proposed method to solve Sudoku puzzles without telling the model the rules.

We follow the experiment setting in Wang et al. (2019) to represent the solution of Sudoku as a vector $\mathbf{y} \in \{0, 1\}^{729}$, where y_{ijk} denotes the i -th row j -th column is the number k or not. The partially filled entries $(r_i, c_i) = n_i$ (i.e., row r_i column c_i is number n_i) are encoded in $\mathbf{w} \in \{0, 1\}^{729}$, where we set the corresponding weight for $w_{r_i c_i n_i}$ to be 1 and the rest to be 0. In this way, maximizing the objective function $\mathbf{w} \cdot \mathbf{y}$ guarantees $y_{ijk} = 1$ if $w_{ijk} = 1$. Then given pairs of $\{\mathbf{w}, \mathbf{y}\}$, our methods mine the underlying constraints A and \mathbf{b} in Eq. (3).

We experiment on the dataset introduced in Wang et al. (2019). The dataset contains 9,000 training and 1,000 test samples, each of which has a unique solution. We also conduct experiments in the permuted setting (Wang et al., 2019), where a pre-defined permutation function is used to shuffle the 9×9 grid. In the permuted setting, it is almost impossible for humans to identify the underlying rules, despite the puzzle is still filled in a certain order. We follow the configuration in Wang et al. (2019) to compare our approach with a convolution neural network for Sudoku (Park, 2018) (ConvNet) and SATNet (Wang et al., 2019) and report our results along with their published results in Table 1.

As shown in the table, by using the equality constraints mining technique, our framework can realize the Sudoku rules and achieve 100% accuracy. The constraints we mine reduce the size of candidate solution space from 2^{729} to 6.67×10^{21} (i.e., the number of feasible Sudoku puzzles). Note that our approaches, as well as SATNet, do not utilize the position clues in the data. Therefore, it is not affected by permutations.

Table 2: MST results in exact match (EM) accuracy in Train and Test, the average edge accuracy (Edge) and the ratio of solutions that are feasible (Feasibility). Here, the feasibility means the solution forms a tree.

(a) 10,000 training samples				
Model	Train EM	Test EM	Test Edge	Test Feasibility
NN	8.3%	6.9%	89.3%	12.8%
Inner	100%	41.6%	93.7%	100%
Outer	100%	71.1%	96.1%	71.1%
Outer+EQ	100%	72.9%	96.1%	72.9%

(b) 20,000 training samples				
Model	Train EM	Test EM	Test Edge	Test Feasibility
NN	9.5%	10.4%	91.1%	13.4%
Inner	100%	69.2%	97.0%	100%
Outer	100%	87.2%	98.0%	87.2%
Outer+EQ	100%	91.8%	98.7%	91.8%

In fact, the equality constraints we learn are exactly the rules of Sudoku. The Sudoku rules can be represented linearly as

$$\sum_{i=1}^9 y_{kij} = 1, \sum_{i=1}^9 y_{jki} = 1, \forall j, k, \quad (7)$$

$$\sum_{i=1}^3 \sum_{j=1}^3 y_{(x+i)(y+j)k} = 1, \forall x, y \in \{0, 3, 6\} \forall k.$$

We use S^* to denote the space defined by Eq. (10) and \hat{S} as the space of our constraints. We verify the $S^* = \hat{S}$ by

1. Verifying that each constraint in S^* is indicated by the constraints we learn. This property guarantee that $S^* \subseteq \hat{S}$.
2. Comparing the dimension of S^* and S . We find that $d(S^*) = d(\hat{S}) = 249$.

In this way, we confirm that our framework can mine the underlying Sudoku rules successfully.

4.2. Minimal Spanning Tree

As discussed in Sec. 1, inference problems in many structured output prediction applications (e.g., dependency parsing) can be modeled as searching the minimal (or maximal) spanning tree (MST). In the following, we verify if the proposed approach can identify the structure of solution is a tree by merely providing pairs of adjacency matrices and the corresponding MST. We encode the MST problem as described in Sec. 3.

We generate a dataset with 7 nodes. The dataset contains 20,000 training and 500 test data. Every data point consists of an adjacency matrix serialized in a vector \mathbf{w} , where $w_{i,j}$ represents the distance between node i and node j , and its corresponding MST. The entry in the adjacency matrix is independently sampled from a uniform distribution in

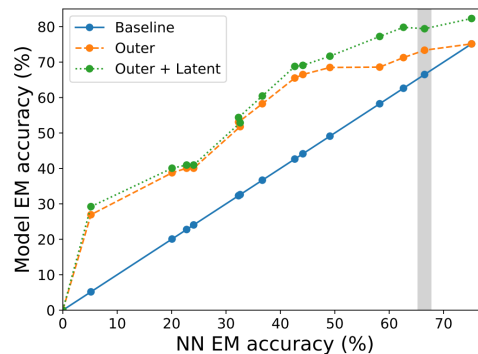


Figure 2: Results on ImCLEF07A in EM accuracy using base models with different performance levels.

$[-1, 1]$. We filtered out the adjacency matrix with identical values to ensure every sample has a unique optimal solution.

We test the inner polytope (Inner), outer polytope (Outer) and outer polytope with equality constraint mining (Outer+EQ) methods. We compare our approaches with fully connected feed-forward neural networks (NN) with 1, 2 or 3 layers, which directly learn the association between \mathbf{w} and \mathbf{y} and the hyper-parameters are given in the Appendix. We set the hidden dimension to be 50. Despite that our methods do not have hyper-parameters, to tune the neural network, we generate another 500 dev data points.

Table 2 shows the results in exact match (i.e., correct MST) and edge accuracies. The results show that Baseline-NN is unable to learn the tree structure from the given examples. We find that NN can learn reasonably well in each individual edge but is terrible to capture the output is a tree. In particular, in 87.2% and 86.6% of cases for 10,000, 20,000 training samples, respectively, the output by NN is not feasible (not a tree). Therefore, its exact match accuracy is low. In comparison, the proposed approaches Outer and Outer+EQ mostly produce feasible solutions⁶, resulting in much higher exact match accuracy. Comparing the results for 10,000 and 20,000 data points, we find that Outer+EQ is more effective than NN when doubling the training data as it improves 20% exact match accuracy.

4.3. Hierarchical Multi-label Classification

Finally, we apply the proposed approaches to a real-world problem and demonstrate its ability to cooperate with machine learning models. We conduct experiments on ImCLEF07A (Dimitrovski et al., 2011), which contains 10,000 training samples and 1,006 test samples. Each sam-

⁶Note that for Outer and Outer+EQ methods, the results of feasibility are the same as test EM because all feasible trees are in the outer polytope. Therefore, if a model outputs a feasible tree, the tree is guaranteed to be optimal.

Table 3: Detailed results (grey zone in Fig. 2) evaluated by exact match accuracy (EM), the ratio of solutions that are feasible (Feasibility), and average accuracy of label assignments (Label Acc.).

Model	Test EM	Test Feasibility	Test Label Acc.
Baseline	62.6%	73.8%	98.92%
Inner	79.8%	100%	98.98%
Outer	71.3%	89.7%	98.85%
Outer+Latent	79.8%	100%	98.98%

ple has 80 features and a set of labels selected from 96 classes. There is a hierarchy among the labels and the depth of the hierarchical structure is 4. A feasible label set forms a path from the root to a leaf node.

The base model is a 3-layer fully connected feed-forward neural network with hidden dimension 80. This model outputs a vector \mathbf{c} , where each component $c_i \in [0, 1]$ is predicted independently to the input instance. For the baseline model, if $c_i > 0.5$, then the label i is positive.

We take the base model as a sub-routine and use it to assign weight \mathbf{w} in Eq. (3). Specifically, $\mathbf{w} = \mathbf{c} - 0.5 \times \mathbf{1}$. Without constraints, solving the ILP in Eq. (3) is equivalent to make prediction by the baseline model. We evaluate 1) the inner polytope method (Inner), 2) the outer polytope method (Outer) and 3) the outer polytope method with latent variables (Outer+Latent). In Outer, as \mathbf{w} is generated by a predicted model, we use Eq. (4) to allow noise. In Outer+Latent, we use the label pairwise latent variables defined in Sec 3.3. To reduce the label spaces, we follow the convention to consider only induce latent variables to label pairs that occur in the training set.

Different from Sec. 4.1 and 4.2, the weight \mathbf{w} is a score vector predicted by the base model. To understand how the constraint mining approaches incorporate with base models with different performance levels, we train multiple versions of base models with different number of layers and training epochs then demonstrate the performance of our approach with these base models. The hyper-parameters for these models are given in the Appendix. The results are shown in Fig. 2.

Results show that Outer and Outer+Latent improve the base models in all cases. Even with a weak base model with only 10% in exact match accuracy, Outer and Outer+Latent are able to learn underlying constraints and improve the performance by more than 20%. The difference between Outer and Outer+Latent is not apparent when the base model is inaccurate. However, when the base model performance increases, Outer+Latent is capable of capturing more fine-grained constraints than Outer and achieves better performance. When the baseline achieves 0 loss in training data (the right-most column points), the constraints

learned by Outer can not filter out any point in the space. Therefore, Outer achieves the same performance as *Baseline*. However, Outer+Latent can still mine constraints related to the latent variable and improve the performance.

Table 3 highlights the detailed results with one base model.⁷ Outer improves Baseline about 12% in exact match accuracy and 16% in feasibility. This demonstrates Outer can successfully filter out many infeasible solutions and guide the model to find the correct ones. Inner learns exactly the feasible set as all pairs of labels appear in the test set also appear in the training. Similarly, Outer+Latent is able to identify all dependencies between labels and achieves high performance. The classes in this task have a tree structure and it has depth 4 including the root (root is a virtual concept, and it is not a real class). We use $p(x)$ to denote the parent class of class x , and L_i to denote the set of classes on layer i , $i \in \{1, 2, 3\}$. We verify the constraints with the same method in Sec. 4.1. The mined equality constraints are the linear transformation of the following constraints:

$$\begin{aligned} \sum_{j \in L_i} y_j &= 1, \forall i \in \{1, 2, 3\}, \\ h_{parent(x), x, 0, 1} &= 0, \forall x : parent(x) \neq root. \end{aligned} \quad (8)$$

5. Analysis and Discussion

5.1. Feasible Set Size Analysis

We provide a theoretical analysis about the convergence speed of the proposed approaches by estimating the cardinality of the outer polytope and inner polytopes. Our methods in Sec. 3 estimate the feasible set by squeezing the outer polytope and enlarging the inner polytope. We analyze how the sizes of outer polytope and inner polytope change with respect to the number of training samples.

We denote the size of ground truth feasible set S^* as M , the size of universal label space is N . The weights \mathbf{w} are drawn from the distribution $D_{\mathbf{w}}$. For each point i , we use p_i to denote the probability that given a randomly sampled weight $\mathbf{w} \sim D_{\mathbf{w}}$, $\mathbf{y}^{(i)}$ get higher score than all the feasible points. Formally, $p_i = P_{\mathbf{w} \sim D_{\mathbf{w}}} \{ \mathbf{w} \cdot \mathbf{y}^{(i)} \geq \mathbf{w} \cdot \mathbf{y}^{(j)}, \forall j \in S^* \}$. The following lemma bounds the expectation sizes of the outer and inner polygons. Full proof is in the Appendix.

Lemma 3: The expectation of the sizes of outer and inner polygon is given by

$$\begin{aligned} \mathbb{E}[|S_I|] &= M - \sum_{i \in S^*} (1 - p_i)^k; \\ \mathbb{E}[|S_O|] &= M + \sum_{j \notin S^*} (1 - p_j)^k. \end{aligned} \quad (9)$$

⁷We choose the second best baseline model since the best one gets 100% accuracy in training set, which causes the inequality constraints learned by Outer method filter out nothing.

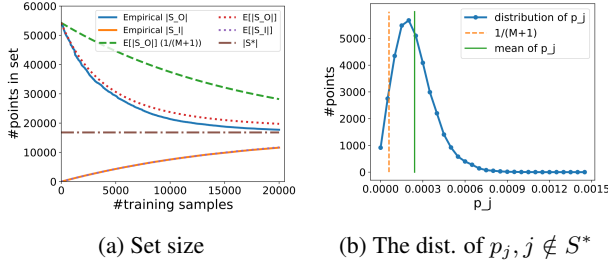


Figure 3: On the left Fig. 3a shows the empirical and expectation sizes of outer and inner polytope, comparing with ground truth. $E[|S_O|](1/(M+1))$ is the expectation of outer polytope size estimated by $p_j = \frac{1}{M+1}$ while $E[|S_O|]$ is estimated with empirical p_j . Note that the empirical inner polytope and its theoretical curve almost coincide. On the right Fig. 3b shows the distribution of p_j estimated from 20,000 training data comparing with $1/(M+1)$.

Proof sketch: We first consider the outer polygon. The point j out of the feasible set appears in the outer polygon if and only if it is not filtered out by any constraints, which is $(1 - p_j)^k$.

We then consider the inner polygon. The point i appears in the inner polygon if and only if at least one training sample takes it as the optimal solution, which is $1 - (1 - p_i)^k$.

Case study: MST We take MST discussed in Sec.4.2 as an example. According to Matrix-Tree Theorem (Chaiken & Kleitman, 1978), we know that the number of spanning trees is $M = 16,807$. The universal set size before mining equality constraints is $N_O = 2^{21} = 2,097,152$. However, with the equality constraints mining method, we can identify the constraint $\sum_{i=1}^{21} y_i = 6$, (i.e., number of edges is 6). With this constraint the size of the space is reduced to $N = \binom{21}{6} = 54,264$.

p_i in Eq. (9) is difficult to estimate directly; therefore, we approximate it by $p_i = \frac{1}{M}$ for $i \in S^*$ and $p_j = \frac{1}{M+1}$ for $j \notin S^*$. The approximation is exact if the following assumption holds (see details and proof in Appendix). *Data symmetric:* for k different points $\mathbf{y}^{(1)}, \mathbf{y}^{(2)}, \dots, \mathbf{y}^{(k)}$, $P_{\mathbf{w} \sim D_{\mathbf{w}}} \{\mathbf{w} \cdot \mathbf{y}^{(1)} \geq \mathbf{w} \cdot \mathbf{y}^{(i)}, i \in [k]\} = \frac{1}{k}$. MST only satisfies the part of the assumption, therefore, for $i \in S^*$, the approximation $p_i = \frac{1}{M}$ is close, and there is a gap between empirical $p_j (j \notin S^*)$ compared with $\frac{1}{M+1}$ (see Fig. 3b), this causes the gap between the empirical result and the estimated expectation about the outer polytope. Fig. 3a shows the empirical sizes of outer and inner polytope $|S_O|, |S_I|$, with their theoretical expectation in Eq. (9) and the ground truth $|S^*|$. The expectation of Inner perfectly fits the empirical results and the two curves are almost coincide. Both of the Outer and Inner methods eventually converge to the ground truth, and the Outer method is closer.

Table 4: Running time in seconds for experiments. The training time is computed by averaging in 3 runs, while inference time is computed by averaging in 100 samples.

Experiment	Model	Training	Inference
MST	Baseline NN	190	5e-4
	Outer	24.1	7.9
	Outer+EQ	26.5	3.8
Sudoku	ConvNet	636.4	5e-4
	Outer+EQ	5.2	1.2
HMC	Baseline NN	44.9	4e-4
	Outer	560.3	1.1
	Outer+Latent	599.3	7.1

5.2. Discussion about Running Time

There are two main factors affecting the running time: the number of data samples, the number of constraints K , and the size of the output variables D . Our approach contains two steps: identifying the feasible set (training) and solving ILPs (inference). In training, as shown in Sec 3, our approach is linear in K since we only needs to pass all samples once, and no worse than quadratic in D . In inference, we solve ILP which is generally NP-hard, and the main factor in complexity is D . However, for most structured prediction tasks, D is small. In our experiments, D is 21 and 729 and 96 in MST, Sudoku and HMC, respectively.

Table 4 shows the training and test running time of our approaches compared to neural network models. In training, our approach is more efficient than the baseline neural network in the Sudoku and MST experiments. For the HMC experiment, the training time of our approach includes updating the model parameters of the underlying neural networks. Therefore, the training time is longer compared to the Sudoku and MST cases.

For the inference time, we report the average time on solving one test sample. Despite ILP is NP-hard, a commercial solver (e.g., Gurobi) is capable of solving the problem within a reasonable time. Therefore, without carefully engineering to optimize the running time, the ILP solver can produce solutions within a few seconds.

To empirically understand the scalability of our approach in the inference, we test the inference time in the MST experiment with larger graph. Here we fix the number of constraints (number of training samples) as 20,000. The results are shown in Fig. 4. We find that despite some extreme cases, the inference time grows generally linearly in terms of the number of variables empirically.

Note that although the constraints are mined using the ILP framework, it does not mean that the inference has to be solved by an ILP solver. Once the constraints are identified, one can design a specific constraint solver to speed up the inference. Besides, the ILP inference can be accelerated




Figure 4: The inference time per instance on MST experiments with different number of variables. Here x-axis is the number of variables D , which equals to $N(N - 1)/2$. N is the number of nodes in the graph. y-axis is the inference time, computed by averaging 100 samples.




Figure 5: The trade-off between inference time and model accuracy when an approximate inference solver is used. Results are on MST. The x-axis is the allowed maximum relative gap between the returned solution and the optimum solution. On the left figure, the performance of our approach drops when the MIPGap is large, but our approach still significantly outperforms the neural network baseline (10.4%). The right figure shows that the inference time significantly reduces when MIPGap gets large.

by amortizing the computations when solving a batch test instances (Srikumar et al., 2012; Chang et al., 2015) or by applying approximate inference algorithms for solving ILP, e.g., LP relaxation methods (Kulesza & Pereira, 2007; Martins et al., 2015).

To demonstrate how the performance of our approach is affected by the approximate ILP solver. We show a trade-off curve in MST 20,000 training experiments in Fig. 5 by solving inference using Gurobi with different MIPGap, a parameter of the Gurobi solver controlling the quality of solutions. Specifically, MIPGap specifies the maximum gap of the objective function values between the returned solution and the optimum solution. We vary MIPGap from 0 (exact solutions are returned) to 0.15. The experimental results demonstrate that the inference can be accelerated

using an approximate inference solver with a trade-off of moderate performance loss.

6. Conclusion

We propose an integer linear programming framework for mining constraints from data. The framework is general and is able to identify underlying constraints in structured prediction problems. Experiments on synthetic problems and hierarchical classification show that the framework is capable of mining complex constraints over a label space, and it can cooperate with neural models. As the first paper to formulate the constraints mining as ILP, we focus on building the foundation for this potential area and understanding the properties of the proposed approach.

7. Acknowledgement

This work was supported by National Science Foundation Grant IIS 1927554 and a Facebook Research Award. We appreciate Cheng Ma and members of the UCLA-NLP lab for their inputs and feedback during this project. We also thank the anonymous reviewers their valuable comments.

References

Bach, S. H., Broecheler, M., Huang, B., and Getoor, L. Hinge-loss markov random fields and probabilistic soft logic. *arXiv preprint arXiv:1505.04406*, 2015.

Bessiere, C., Coletta, R., Hebrard, E., Katsirelos, G., Lazaar, N., Narodytska, N., Quimper, C., and Walsh, T. Constraint acquisition via partial queries. In *IJCAI*, 2013.

Bessiere, C., Daoudi, A., Hebrard, E., Katsirelos, G., Lazaar, N., Mechqrane, Y., Narodytska, N., Quimper, C.-G., and Walsh, T. New approaches to constraint acquisition. In *Data mining and constraint programming*, pp. 51–76. Springer, 2016.

Bessiere, C., Raedt, L. D., Guns, T., Kotthoff, L., Nanni, M., Nijssen, S., O’Sullivan, B., Paparrizou, A., Pedreschi, D., and Simonis, H. The inductive constraint programming loop. *IEEE Intelligent Systems*, 32(5):44–52, 2017.

Bratko, I. and King, R. D. Applications of inductive logic programming. *SIGART Bulletin*, 5(1):43–49, 1994.

Chaiken, S. and Kleitman, D. J. Matrix tree theorems. *Journal of combinatorial theory, Series A*, 24(3):377–381, 1978.

Chang, K., Sundararajan, S., and Keerthi, S. S. Tractable semi-supervised learning of complex structured prediction models. In *ECML/PKDD*, 2013.

- Chang, K., Upadhyay, S., Kundu, G., and Roth, D. Structural learning with amortized inference. In *AAAI*, 2015.
- Chang, M., Ratnoff, L., and Roth, D. Constraints as prior knowledge. In *ICML*, 2008.
- Clarke, J. and Lapata, M. Global inference for sentence compression: An integer linear programming approach. *J. Artif. Intell. Res.*, 31:399–429, 2008.
- Colton, S. and Muggleton, S. Mathematical applications of inductive logic programming. *Machine Learning*, 64 (1-3):25–64, 2006.
- Dimitrovski, I., Kocev, D., Loskovska, S., and Džeroski, S. Hierarchical annotation of medical images. *Pattern Recognition*, 44(10-11):2436–2449, 2011.
- Embar, V., Sridhar, D., Farnadi, G., and Getoor, L. Scalable structure learning for probabilistic soft logic. *arXiv preprint arXiv:1807.00973*, 2018.
- Felgenhauer, B. and Jarvis, F. Enumerating possible sudoku grids. *Preprint available at <http://www.afjarvis.staff.shef.ac.uk/sudoku/sudoku.pdf>*, 2005.
- Finley, T. and Joachims, T. Training structural svms when exact inference is intractable. In *ICML*, 2008.
- Fromer, M. and Globerson, A. An LP view of the M-best MAP problem. In *NIPS*, 2009.
- Ganchev, K., Graça, J., Gillenwater, J., and Taskar, B. Posterior regularization for structured latent variable models. *J. Mach. Learn. Res.*, 11:2001–2049, 2010.
- Goldwasser, D., Srikumar, V., and Roth, D. Predicting structures in NLP: constrained conditional models and integer linear programming in NLP. In *HLT-NAACL*, 2012.
- Gurobi Optimization, L. Gurobi optimizer reference manual, 2019. URL <http://www.gurobi.com>.
- Hooker, J. N. Generalized resolution and cutting planes. *Annals of Operations Research*, 12(1):217–239, 1988.
- Kimmig, A., Bach, S., Broecheler, M., Huang, B., and Getoor, L. A short introduction to probabilistic soft logic. In *Proceedings of the NIPS Workshop on Probabilistic Programming: Foundations and Applications*, 2012.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In Bengio, Y. and LeCun, Y. (eds.), *ICLR*, 2015.
- Kulesza, A. and Pereira, F. Structured learning with approximate inference. In *NIPS*, 2007.
- Kundu, G., Srikumar, V., and Roth, D. Margin-based decomposed amortized inference. In *ACL*, 2013.
- Li, T. and Srikumar, V. Augmenting neural networks with first-order logic. In *ACL*, 2019.
- Li, T., Gupta, V., Mehta, M., and Srikumar, V. A logic-driven framework for consistency of neural models. In *EMNLP/IJCNLP*, 2019.
- Martins, A. F. T., Smith, N. A., and Xing, E. P. Concise integer linear programming formulations for dependency parsing. In *ACL/IJCNLP*, 2009.
- Martins, A. F. T., Smith, N. A., Xing, E. P., Aguiar, P. M. Q., and Figueiredo, M. A. T. Turbo parsers: Dependency parsing by approximate variational inference. In *EMNLP*, 2010.
- Martins, A. F. T., Figueiredo, M. A. T., Aguiar, P. M. Q., Smith, N. A., and Xing, E. P. AD³: alternating directions dual decomposition for MAP inference in graphical models. *J. Mach. Learn. Res.*, 16:495–545, 2015.
- McDonald, R. T., Pereira, F., Ribarov, K., and Hajic, J. Non-projective dependency parsing using spanning tree algorithms. In *HLT/EMNLP*, 2005.
- Meng, T., Peng, N., and Chang, K. Target language-aware constrained inference for cross-lingual dependency parsing. In *EMNLP*, 2019.
- Muggleton, S. and Raedt, L. D. Inductive logic programming: Theory and methods. *J. Log. Program.*, 19/20: 629–679, 1994.
- Murphy, K., Weiss, Y., and Jordan, M. I. Loopy belief propagation for approximate inference: An empirical study. *arXiv preprint arXiv:1301.6725*, 2013.
- Nowozin, S. and Lampert, C. H. Structured learning and prediction in computer vision. *Foundations and Trends in Computer Graphics and Vision*, 6(3-4):185–365, 2011.
- Pan, X., Mehta, M., and Srikumar, V. Learning constraints for structured prediction using rectifier networks. In *ACL*, 2020.
- Park, K. Can convolutional neural networks crack sudoku puzzles? <https://github.com/Kyubyong/sudoku>, 2018.
- Punyakanok, V., Roth, D., Yih, W., and Zimak, D. Semantic role labeling via integer linear programming inference. In *COLING*, 2004.
- Punyakanok, V., Roth, D., Yih, W.-t., and Zimak, D. Learning and inference over constrained output. In *IJCAI*, 2005.

- Raedt, L. D., Passerini, A., and Teso, S. Learning constraints from examples. In *AAAI*, 2018.
- Richardson, M. and Domingos, P. M. Markov logic networks. *Machine Learning*, 62(1-2):107–136, 2006.
- Riedel, S. and Clarke, J. Incremental integer linear programming for non-projective dependency parsing. In *EMNLP*, 2006.
- Rossi, F. and Sperduti, A. Acquiring both constraint and solution preferences in interactive constraint systems. *Constraints*, 9(4):311–332, 2004.
- Roth, D. and Yih, W. A linear programming formulation for global inference in natural language tasks. In *CoNLL*, 2004.
- Roth, D. and Yih, W. Integer linear programming inference for conditional random fields. In *ICML*, 2005.
- Samdani, R. and Roth, D. Efficient decomposed learning for structured prediction. In *ICML*, 2012.
- Srikumar, V., Kundu, G., and Roth, D. On amortizing inference cost for structured prediction. In *EMNLP-CoNLL*, 2012.
- Tan, Y., Terekhov, D., and DeLong, A. Learning linear programs from optimal decisions. *arXiv preprint arXiv:2006.08923*, 2020.
- Valiant, L. G. A theory of the learnable. In *STOC*, 1984.
- Wang, P. and Kolter, J. Z. Low-rank semidefinite programming for the MAX2SAT problem. In *AAAI*, 2019.
- Wang, P., Donti, P. L., Wilder, B., and Kolter, J. Z. SATNet: Bridging deep learning and logical reasoning using a differentiable satisfiability solver. In *ICML*, 2019.
- Ye, Y., Feng, Y., Luo, B., Lai, Y., and Zhao, D. Integrating relation constraints with neural relation extractors. In *AAAI*, 2020.
- Zhao, J., Wang, T., Yatskar, M., Ordonez, V., and Chang, K. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. In *EMNLP*, 2017.

A. Mined Equality Constraints in Sudoku Experiments

The mined equality constraints are the linear transformation of the following constraints:

$$\begin{aligned}
 \sum_{i=1}^9 y_{ijk} &= 1, \forall j, k, \\
 \sum_{j=1}^9 y_{ijk} &= 1, \forall i, k, \\
 \sum_{k=1}^9 y_{ijk} &= 1, \forall i, j, \\
 \sum_{i=1}^3 \sum_{j=1}^3 y_{(x+i)(y+j)k} &= 1, \forall x, y \in \{0, 3, 6\} \forall k.
 \end{aligned} \tag{10}$$

These are the rules of Sudoku described by linear constraints. We denote the linear space defined by these constraints as S^* , and the linear space given by our mined constraints as \hat{S} . We verify $S^* = S$ by

1. For each constraint in Eq. (10), we verify it is indicated in \hat{S} . This property guarantee that $S^* \subseteq \hat{S}$.
2. Comparing the dimension of S^* and S . We find that $d(S^*) = d(\hat{S}) = 249$.

B. Mined Equality Constraints in Hierarchical Multiclass Classification Experiments

The classes in this task have a tree structure and it has depth 4 including the root (root is a virtual concept that it is not a real class). We use $p(x)$ to denote the parent class of class x , and L_i to denote the set of classes on layer i , $i \in \{1, 2, 3\}$. The mined equality constraints are the linear transformation of the following constraints:

$$\begin{aligned}
 \sum_{j \in L_i} y_j &= 1, \forall i \in \{1, 2, 3\}, \\
 h_{parent(x), x, 0, 1} &= 0, \forall x : parent(x) \neq root.
 \end{aligned} \tag{11}$$

We use the same method in Appendix. A to verify it.

C. The Expectation of the Size of Outer and Inner Polytopes

We define

$$p_i = P_{\mathbf{w} \sim D_{\mathbf{w}}} \{ \mathbf{w} \cdot \mathbf{y}^{(i)} \geq \mathbf{w} \cdot \mathbf{y}^{(j)}, \forall j \in S^* \}.$$

When $i \in S^*$, the condition in p_i means the under the given weight \mathbf{w} , i is the optimal point. So p_i is the probability of

i is the ground truth label for a random weight \mathbf{w} sampled from $D_{\mathbf{w}}$.

We then consider $E[S_I]$. The inner polytope is a convex hull of seen feasible points. For each feasible point $i \in S^*$, the probability that we have seen it after k samples is $1 - (1 - p_i)^k$. Thus, after k training samples, the expectation size of the inner polytope is given by

$$E[S_I] = \sum_{i \in S^*} (1 - (1 - p_i)^k) = M - \sum_{i \in S^*} (1 - p_i)^k,$$

where $M = |S^*|$.

When $i \notin S^*$, the condition in p_i means the given weight \mathbf{w} , i is better than all the feasible points. Given \mathbf{w} in training, we find the label is worse than i , which indicate i is infeasible and it will be filtered out. So p_i is the probability that infeasible point i is not filtered out in training for a random weight \mathbf{w} sampled from $D_{\mathbf{w}}$.

We then consider $E[S_O]$. The outer polytope is initialized as the whole space and filters out infeasible points in training. For each infeasible point $i \notin S^*$, the probability that it is not filtered out after k samples is $(1 - p_i)^k$. We also know that all the feasible points will not be filtered out. Thus, after k training samples, the expectation size of the outer polytope is given by

$$E[S_O] = M + \sum_{i \notin S^*} (1 - p_i)^k,$$

where N is the size of the universal set.

D. Expectation Approximation under Assumptions

The data symmetric assumption is: for k different points $\mathbf{y}^{(1)}, \mathbf{y}^{(2)}, \dots, \mathbf{y}^{(k)}$,

$$P_{\mathbf{w} \sim D_{\mathbf{w}}} \{ \mathbf{w} \cdot \mathbf{y}^{(1)} \geq \mathbf{w} \cdot \mathbf{y}^{(i)}, i \in [k] \} = \frac{1}{k}.$$

With this assumption, consider $p_i = P_{\mathbf{w} \sim D_{\mathbf{w}}} \{ \mathbf{w} \cdot \mathbf{y}^{(i)} \geq \mathbf{w} \cdot \mathbf{y}^{(j)}, \forall j \in S^* \}$. When $i \in S^*$, there are M points are taken into consideration. With the assumption, we can get $p_i = \frac{1}{M}, i \in S^*$. When $i \notin S^*$, there are $M + 1$ points are taken into consideration. With the assumption, we can get $p_i = \frac{1}{M+1}, i \notin S^*$.

E. Configurations for the Reproducibility

Data All the data and code can be found in <https://github.com/MtSomeThree/ILPLearning>.

Sudoku Experiments In the Sudoku experiments, we use the baseline following the settings in SATNet(Wang et al.,

2019)⁸.

MST Experiments In the MST experiments, we use the 3-layer feedforward neural network as the baseline model with ReLU activation. The hidden dimension is set to be 50. The input and output dimension is 21. We use the sigmoid function to regularize the output in $(-1, 1)$. We train the model for 300 epochs and we use the Adam optimizer(Kingma & Ba, 2015) to optimize the model. The learning rate is set to be 0.001.

Hierarchical Multi-label Classification Experiments

In the HMC experiments, we use multiple base models. We enumerate the number of layers in $\{1, 2, 3\}$, the number of training epochs in $\{1, 5, 50, 300\}$, and the learning rate in $\{0.001, 0.0003, 0.0001\}$. The hidden dimension is set to be 100. The input dimension is 80 and the output dimension is 96. In hidden layer we use ReLU as the activation and in output we use sigmoid function to regularize the output.

⁸The baseline models can be found in <https://github.com/locuslab/SATNet>.