

Backpropagating through Structured Argmax using a SPIGOT

Hao Peng[◇] Sam Thomson[♣] Noah A. Smith[◇]

[◇] Paul G. Allen School of Computer Science & Engineering, University of Washington

[♣] School of Computer Science, Carnegie Mellon University

{hapeng, nasmith}@cs.washington.edu, sthompson@cs.cmu.edu

Abstract

We introduce the **structured projection of intermediate gradients** optimization technique (SPIGOT), a new method for backpropagating through neural networks that include hard-decision structured predictions (e.g., parsing) in intermediate layers. SPIGOT requires no marginal inference, unlike structured attention networks (Kim et al., 2017) and some reinforcement learning-inspired solutions (Yogatama et al., 2017). Like so-called straight-through estimators (Hinton, 2012), SPIGOT defines gradient-like quantities associated with intermediate nondifferentiable operations, allowing backpropagation before and after them; SPIGOT’s proxy aims to ensure that, after a parameter update, the intermediate structure will remain well-formed.

We experiment on two structured NLP pipelines: syntactic-then-semantic dependency parsing, and semantic parsing followed by sentiment classification. We show that training with SPIGOT leads to a larger improvement on the downstream task than a modularly-trained pipeline, the straight-through estimator, and structured attention, reaching a new state of the art on semantic dependency parsing.

1 Introduction

Learning methods for natural language processing are increasingly dominated by end-to-end differentiable functions that can be trained using gradient-based optimization. Yet traditional NLP often assumed modular stages of processing that formed a pipeline; e.g., text was tokenized, then tagged with parts of speech, then parsed into a

phrase-structure or dependency tree, then semantically analyzed. Pipelines, which make “hard” (i.e., discrete) decisions at each stage, appear to be incompatible with neural learning, leading many researchers to abandon earlier-stage processing.

Inspired by findings that continue to see benefit from various kinds of linguistic or domain-specific preprocessing (He et al., 2017; Oepen et al., 2017; Ji and Smith, 2017), we argue that pipelines can be treated as layers in neural architectures for NLP tasks. Several solutions are readily available:

- Reinforcement learning (most notably the REINFORCE algorithm; Williams, 1992), and **structured attention** (SA; Kim et al., 2017). These methods replace argmax with a sampling or marginalization operation. We note two potential downsides of these approaches: (i) not all argmax-able operations have corresponding sampling or marginalization methods that are efficient, and (ii) inspection of intermediate outputs, which could benefit error analysis and system improvement, is more straightforward for hard decisions than for posteriors.
- The **straight-through estimator** (STE; Hinton, 2012) treats discrete decisions as if they were differentiable and simply passes through gradients. While fast and surprisingly effective, it ignores *constraints* on the argmax problem, such as the requirement that every word has exactly one syntactic parent. We will find, experimentally, that the quality of intermediate representations degrades substantially under STE.

This paper introduces a new method, the **structured projection of intermediate gradients** optimization technique (SPIGOT; §2), which defines a proxy for the gradient of a loss function with respect to the input to argmax. Unlike STE’s gradient proxy, SPIGOT aims to respect the constraints

in the argmax problem. SPIGOT can be applied with any intermediate layer that is expressible as a constrained maximization problem, and whose feasible set can be projected onto. We show empirically that SPIGOT works even when the maximization and the projection are done approximately.

We offer two concrete architectures that employ structured argmax as an intermediate layer: semantic parsing with syntactic parsing in the middle, and sentiment analysis with semantic parsing in the middle (§3). These architectures are trained using a joint objective, with one part using data for the intermediate task, and the other using data for the end task. The datasets are not assumed to overlap at all, but the parameters for the intermediate task are affected by both parts of the training data.

Our experiments (§4) show that our architecture improves over a state-of-the-art semantic dependency parser, and that SPIGOT offers stronger performance than a pipeline, SA, and STE. On sentiment classification, we show that semantic parsing offers improvement over a BiLSTM, more so with SPIGOT than with alternatives. Our analysis considers how the behavior of the intermediate parser is affected by the end task (§5). Our code is open-source and available at <https://github.com/Noahs-ARK/SPIGOT>.

2 Method

Our aim is to allow a (structured) argmax layer in a neural network to be treated almost like any other differentiable function. This would allow us to place, for example, a syntactic parser in the middle of a neural network, so that the forward calculation simply calls the parser and passes the parse tree to the next layer, which might derive syntactic features for the next stage of processing.

The challenge is in the *backward* computation, which is key to learning with standard gradient-based methods. When its output is discrete as we assume here, argmax is a piecewise constant function. At every point, its gradient is either zero or undefined. So instead of using the true gradient, we will introduce a *proxy* for the gradient of the loss function with respect to the inputs to argmax, allowing backpropagation to proceed through the argmax layer. Our proxy is designed as an improvement to earlier methods (discussed below) that completely ignore constraints on the argmax operation. It accomplishes this through a projec-

tion of the gradients.

We first lay out notation, and then briefly review max-decoding and its relaxation (§2.1). We define SPIGOT in §2.2, and show how to use it to back-propagate through NLP pipelines in §2.3.

Notation. Our discussion centers around two tasks: a structured *intermediate* task followed by an *end* task, where the latter considers the outputs of the former (e.g., syntactic-then-semantic parsing). Inputs are denoted as \mathbf{x} , and end task outputs as \mathbf{y} . We use \mathbf{z} to denote intermediate structures derived from \mathbf{x} . We will often refer to the intermediate task as “decoding”, in the structured prediction sense. It seeks an output $\hat{\mathbf{z}} = \operatorname{argmax}_{\mathbf{z} \in \mathcal{Z}} S$ from the feasible set \mathcal{Z} , maximizing a (learned, parameterized) scoring function S for the structured intermediate task. L denotes the loss of the end task, which may or may not also involve structured predictions. We use $\Delta^{k-1} = \{\mathbf{p} \in \mathbb{R}^k \mid \mathbf{1}^\top \mathbf{p} = 1, \mathbf{p} \geq \mathbf{0}\}$ to denote the $(k - 1)$ -dimensional simplex. We denote the domain of binary variables as $\mathbb{B} = \{0, 1\}$, and the unit interval as $\mathbb{U} = [0, 1]$. By projection of a vector \mathbf{v} onto a set \mathcal{A} , we mean the closest point in \mathcal{A} to \mathbf{v} , measured by Euclidean distance: $\operatorname{proj}_{\mathcal{A}}(\mathbf{v}) = \operatorname{argmin}_{\mathbf{v}' \in \mathcal{A}} \|\mathbf{v}' - \mathbf{v}\|_2$.

2.1 Relaxed Decoding

Decoding problems are typically decomposed into a collection of “parts”, such as arcs in a dependency tree or graph. In such a setup, each element of \mathbf{z} , z_i , corresponds to one possible part, and z_i takes a boolean value to indicate whether the part is included in the output structure. The scoring function S is assumed to decompose into a vector $\mathbf{s}(\mathbf{x})$ of part-local, input-specific scores:

$$\hat{\mathbf{z}} = \operatorname{argmax}_{\mathbf{z} \in \mathcal{Z}} S(\mathbf{x}, \mathbf{z}) = \operatorname{argmax}_{\mathbf{z} \in \mathcal{Z}} \mathbf{z}^\top \mathbf{s}(\mathbf{x}) \quad (1)$$

In the following, we drop \mathbf{s} ’s dependence on \mathbf{x} for clarity.

In many NLP problems, the output space \mathcal{Z} can be specified by linear constraints (Roth and Yih, 2004):

$$\mathbf{A} \begin{bmatrix} \mathbf{z} \\ \boldsymbol{\psi} \end{bmatrix} \leq \mathbf{b}, \quad (2)$$

where $\boldsymbol{\psi}$ are auxiliary variables (also scoped by argmax), together with integer constraints (typically, each $z_i \in \mathbb{B}$).

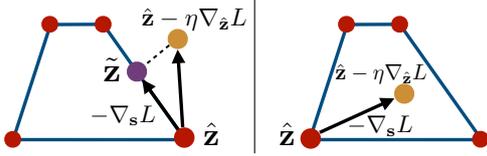


Figure 1: The original feasible set \mathcal{Z} (red vertices), is relaxed into a convex polytope \mathcal{P} (the area encompassed by blue edges). Left: making a gradient update to $\hat{\mathbf{z}}$ makes it step outside the polytope, and it is projected back to \mathcal{P} , resulting in the projected point $\tilde{\mathbf{z}}$. $\nabla_{\mathbf{s}}L$ is then along the edge. Right: updating $\hat{\mathbf{z}}$ keeps it within \mathcal{P} , and thus $\nabla_{\mathbf{s}}L = \eta \nabla_{\hat{\mathbf{z}}}L$.

The problem in Equation 1 can be NP-complete in general, so the $\{0, 1\}$ constraints are often relaxed to $[0, 1]$ to make decoding tractable (Martins et al., 2009). Then the discrete combinatorial problem over \mathcal{Z} is transformed into the optimization of a linear objective over a convex polytope $\mathcal{P} = \{\mathbf{p} \in \mathbb{R}^d \mid \mathbf{A}\mathbf{p} \leq \mathbf{b}\}$, which is solvable in polynomial time (Bertsimas and Tsitsiklis, 1997). This is not necessary in some cases, where the argmax can be solved exactly with dynamic programming.

2.2 From STE to SPIGOT

We now view structured argmax as an activation function that takes a vector of input-specific part-scores \mathbf{s} and outputs a solution $\hat{\mathbf{z}}$. For backpropagation, to calculate gradients for parameters of \mathbf{s} , the chain rule defines:

$$\nabla_{\mathbf{s}}L = \mathbf{J} \nabla_{\hat{\mathbf{z}}}L, \quad (3)$$

where the Jacobian matrix $\mathbf{J} = \frac{\partial \hat{\mathbf{z}}}{\partial \mathbf{s}}$ contains the derivative of each element of $\hat{\mathbf{z}}$ with respect to each element of \mathbf{s} . Unfortunately, argmax is a piecewise constant function, so its Jacobian is either zero (almost everywhere) or undefined (in the case of ties).

One solution, taken in *structured attention*, is to replace the argmax with marginal inference and a softmax function, so that $\hat{\mathbf{z}}$ encodes probability distributions over parts (Kim et al., 2017; Liu and Lapata, 2018). As discussed in §1, there are two reasons to avoid this modification. Softmax can only be used when marginal inference is feasible, by sum-product algorithms for example (Eisner, 2016; Friesen and Domingos, 2016); in general marginal inference can be #P-complete. Further, a soft intermediate layer will be less amenable to inspection by anyone wishing to understand and improve the model.

In another line of work, argmax is augmented with a strongly-convex penalty on the solutions (Martins and Astudillo, 2016; Amos and Kolter, 2017; Niculae and Blondel, 2017; Niculae et al., 2018; Mensch and Blondel, 2018). However, their approaches require solving a relaxation even when exact decoding is tractable. Also, the penalty will bias the solutions found by the decoder, which may be an undesirable conflation of computational and modeling concerns.

A simpler solution is the STE method (Hinton, 2012), which replaces the Jacobian matrix in Equation 3 by the identity matrix. This method has been demonstrated to work well when used to “backpropagate” through hard threshold functions (Bengio et al., 2013; Friesen and Domingos, 2018) and categorical random variables (Jang et al., 2016; Choi et al., 2017).

Consider for a moment what we would do if $\hat{\mathbf{z}}$ were a vector of parameters, rather than intermediate predictions. In this case, we are seeking points in \mathcal{Z} that minimize L ; denote that set of minimizers by \mathcal{Z}^* . Given $\nabla_{\hat{\mathbf{z}}}L$ and step size η , we would update $\hat{\mathbf{z}}$ to be $\hat{\mathbf{z}} - \eta \nabla_{\hat{\mathbf{z}}}L$. This update, however, might not return a value in the feasible set \mathcal{Z} , or even (if we are using a linear relaxation) the relaxed set \mathcal{P} .

SPIGOT therefore introduces a *projection* step that aims to keep the “updated” $\hat{\mathbf{z}}$ in the feasible set. Of course, we do not directly update $\hat{\mathbf{z}}$; we continue backpropagation through \mathbf{s} and onward to the parameters. But the projection step nonetheless alters the parameter updates in the way that our proxy for “ $\nabla_{\mathbf{s}}L$ ” is defined.

The procedure is defined as follows:

$$\hat{\mathbf{p}} = \hat{\mathbf{z}} - \eta \nabla_{\hat{\mathbf{z}}}L, \quad (4a)$$

$$\tilde{\mathbf{z}} = \text{proj}_{\mathcal{P}}(\hat{\mathbf{p}}), \quad (4b)$$

$$\nabla_{\mathbf{s}}L \triangleq \hat{\mathbf{z}} - \tilde{\mathbf{z}}. \quad (4c)$$

First, the method makes an “update” to $\hat{\mathbf{z}}$ as if it contained parameters (Equation 4a), letting $\hat{\mathbf{p}}$ denote the new value. Next, $\hat{\mathbf{p}}$ is projected back onto the (relaxed) feasible set (Equation 4b), yielding a feasible new value $\tilde{\mathbf{z}}$. Finally, the gradients with respect to \mathbf{s} are computed by Equation 4c.

Due to the convexity of \mathcal{P} , the projected point $\tilde{\mathbf{z}}$ will always be unique, and is guaranteed to be no farther than $\hat{\mathbf{p}}$ from any point in \mathcal{Z}^* (Luenberger and Ye, 2015).¹ Compared to STE, SPIGOT in-

¹Note that this property follows from \mathcal{P} ’s convexity, and we do not assume the convexity of L .

volves a projection and limits $\nabla_s L$ to a smaller space to satisfy constraints. See Figure 1 for an illustration.

When efficient exact solutions (such as dynamic programming) are available, they can be used. Yet, we note that SPIGOT does not assume the argmax operation is solved exactly.

2.3 Backpropagation through Pipelines

Using SPIGOT, we now devise an algorithm to “backpropagate” through NLP pipelines. In these pipelines, an intermediate task’s output is fed into an end task for use as features. The parameters of the complete model are divided into two parts: denote the parameters of the intermediate task model by ϕ (used to calculate s), and those in the end task model as θ .² As introduced earlier, the end-task loss function to be minimized is L , which depends on both ϕ and θ .

Algorithm 1 describes the forward and backward computations. It takes an end task training pair $\langle \mathbf{x}, \mathbf{y} \rangle$, along with the intermediate task’s feasible set \mathcal{Z} , which is determined by \mathbf{x} . It first runs the intermediate model and decodes to get intermediate structure $\hat{\mathbf{z}}$, just as in a standard pipeline. Then forward propagation is continued into the end-task model to compute loss L , using $\hat{\mathbf{z}}$ to define input features. Backpropagation in the end-task model computes $\nabla_\theta L$ and $\nabla_{\hat{\mathbf{z}}} L$, and $\nabla_s L$ is then constructed using Equations 4. Backpropagation then continues into the intermediate model, computing $\nabla_\phi L$.

Due to its flexibility, SPIGOT is applicable to many training scenarios. When there is no $\langle \mathbf{x}, \mathbf{z} \rangle$ training data for the intermediate task, SPIGOT can be used to induce latent structures for the end-task (Yogatama et al., 2017; Kim et al., 2017; Choi et al., 2017, *inter alia*). When intermediate-task training data is available, one can use SPIGOT to adopt joint learning by minimizing an interpolation of L (on end-task data $\langle \mathbf{x}, \mathbf{y} \rangle$) and an intermediate-task loss function \tilde{L} (on intermediate task data $\langle \mathbf{x}, \mathbf{z} \rangle$). This is the setting in our experiments; note that we do not assume any overlap in the training examples for the two tasks.

3 Solving the Projections

In this section we discuss how to compute approximate projections for the two intermediate tasks

²Nothing prohibits tying across pre-argmax parameters and post-argmax parameters; this separation is notationally convenient but not at all necessary.

Algorithm 1 Forward and backward computation with SPIGOT.

```

1: procedure SPIGOT( $\mathbf{x}, \mathbf{y}, \mathcal{Z}$ )
2:   Construct  $\mathbf{A}, \mathbf{b}$  such that  $\mathcal{Z} = \{\mathbf{p} \in \mathbb{Z}^d \mid \mathbf{A}\mathbf{p} \leq \mathbf{b}\}$ 
3:    $\mathcal{P} \leftarrow \{\mathbf{p} \in \mathbb{R}^d \mid \mathbf{A}\mathbf{p} \leq \mathbf{b}\}$  ▷ Relaxation
4:   Forwardprop and compute  $s_\phi(\mathbf{x})$ 
5:    $\hat{\mathbf{z}} \leftarrow \operatorname{argmax}_{\mathbf{z} \in \mathcal{Z}} \mathbf{z}^\top s_\phi(\mathbf{x})$  ▷ Intermediate decoding
6:   Forwardprop and compute  $L$  given  $\mathbf{x}, \mathbf{y}$ , and  $\hat{\mathbf{z}}$ 
7:   Backprop and compute  $\nabla_\theta L$  and  $\nabla_{\hat{\mathbf{z}}} L$ 
8:    $\tilde{\mathbf{z}} \leftarrow \operatorname{proj}_{\mathcal{P}}(\hat{\mathbf{z}} - \eta \nabla_{\hat{\mathbf{z}}} L)$  ▷ Projection
9:    $\nabla_s L \leftarrow \hat{\mathbf{z}} - \tilde{\mathbf{z}}$ 
10:  Backprop and compute  $\nabla_\phi L$ 
11: end procedure

```

considered in this work, arc-factored unlabeled dependency parsing and first-order semantic dependency parsing.

In early experiments we observe that for both tasks, projecting with respect to *all* constraints of their original formulations using a generic quadratic program solver was prohibitively slow. Therefore, we construct relaxed polytopes by considering only a subset of the constraints.³ The projection then decomposes into a series of singly constrained quadratic programs (QP), each of which can be efficiently solved in linear time.

The two approximate projections discussed here are used in backpropagation only. In the forward pass, we solve the decoding problem using the models’ original decoding algorithms.

Arc-factored unlabeled dependency parsing.

For unlabeled dependency trees, we impose $[0, 1]$ constraints and single-headedness constraints.⁴

Formally, given a length- n input sentence, excluding self-loops, an arc-factored parser considers $d = n(n - 1)$ candidate arcs. Let $i \rightarrow j$ denote an arc from the i th token to the j th, and $\sigma(i \rightarrow j)$ denote its index. We construct the relaxed feasible set by:

$$\mathcal{P}_{\text{DEP}} = \left\{ \mathbf{p} \in \mathbb{U}^d \mid \sum_{i \neq j} p_{\sigma(i \rightarrow j)} = 1, \forall j \right\}, \quad (5)$$

i.e., we consider each token j individually, and force single-headedness by constraining the number of arcs incoming to j to sum to 1. Algorithm 2 summarizes the procedure to project onto \mathcal{P}_{DEP} .

³A parallel work introduces an active-set algorithm to solve the same class of quadratic programs (Niculae et al., 2018). It might be an efficient approach to solve the projections in Equation 4b, which we leave to future work.

⁴It requires $O(n^2)$ auxiliary variables and $O(n^3)$ additional constraints to ensure well-formed tree structures (Martins et al., 2013).

Line 3 forms a singly constrained QP, and can be solved in $O(n)$ time (Brucker, 1984).

Algorithm 2 Projection onto the relaxed polytope \mathcal{P}_{DEP} for dependency tree structures. Let bold $\sigma(\cdot \rightarrow j)$ denote the index set of arcs incoming to j . For a vector \mathbf{v} , we use $\mathbf{v}_{\sigma(\cdot \rightarrow j)}$ to denote vector $[v_k]_{k \in \sigma(\cdot \rightarrow j)}$.

```

1: procedure DEPPROJ( $\hat{\mathbf{p}}$ )
2:   for  $j = 1, 2, \dots, n$  do
3:      $\tilde{\mathbf{z}}_{\sigma(\cdot \rightarrow j)} \leftarrow \text{proj}_{\Delta^{n-2}}(\hat{\mathbf{p}}_{\sigma(\cdot \rightarrow j)})$ 
4:   end for
5:   return  $\tilde{\mathbf{z}}$ 
6: end procedure

```

First-order semantic dependency parsing. Semantic dependency parsing uses labeled bilexical dependencies to represent sentence-level semantics (Oepen et al., 2014, 2015, 2016). Each dependency is represented by a labeled directed arc from a head token to a modifier token, where the arc label encodes broadly applicable semantic relations. Figure 2 diagrams a semantic graph from the DELPH-IN MRS-derived dependencies (DM), together with a syntactic tree.

We use a state-of-the-art semantic dependency parser (Peng et al., 2017) that considers three types of parts: heads, unlabeled arcs, and labeled arcs. Let $\sigma(i \xrightarrow{\ell} j)$ denote the index of the arc from i to j with semantic role ℓ . In addition to $[0, 1]$ constraints, we constrain that the predictions for labeled arcs sum to the prediction of their associated unlabeled arc:

$$\mathcal{P}_{\text{SDP}} \left\{ \mathbf{p} \in \mathbb{U}^d \left| \sum_{\ell} p_{\sigma(i \xrightarrow{\ell} j)} = p_{\sigma(i \rightarrow j)}, \forall i \neq j \right. \right\}. \quad (6)$$

This ensures that exactly one label is predicted if and only if its arc is present. The projection onto \mathcal{P}_{SDP} can be solved similarly to Algorithm 2. We drop the determinism constraint imposed by Peng et al. (2017) in the backward computation.

4 Experiments

We empirically evaluate our method with two sets of experiments: using syntactic tree structures in semantic dependency parsing, and using semantic dependency graphs in sentiment classification.

4.1 Syntactic-then-Semantic Parsing

In this experiment we consider an intermediate syntactic parsing task, followed by seman-

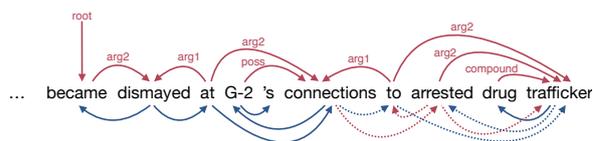


Figure 2: A development instance annotated with both gold DM semantic dependency graph (red arcs on the top), and gold syntactic dependency tree (blue arcs at the bottom). A pretrained syntactic parser predicts the same tree as the gold; the semantic parser backpropagates into the intermediate syntactic parser, and changes the dashed blue arcs into dashed red arcs (§5).

tic dependency parsing as the end task. We first briefly review the neural network architectures for the two models (§4.1.1), and then introduce the datasets (§4.1.2) and baselines (§4.1.3).

4.1.1 Architectures

Syntactic dependency parser. For intermediate syntactic dependencies, we use the unlabeled arc-factored parser of Kiperwasser and Goldberg (2016). It uses bidirectional LSTMs (BiLSTM) to encode the input, followed by a multilayer-perceptron (MLP) to score each potential dependency. One notable modification is that we replace their use of Chu-Liu/Edmonds’ algorithm (Chu and Liu, 1965; Edmonds, 1967) with the Eisner algorithm (Eisner, 1996, 2000), since our dataset is in English and mostly projective.

Semantic dependency parser. We use the basic model of Peng et al. (2017) (denoted as NEURBOPARSER) as the end model. It is a first-order parser, and uses local factors for heads, unlabeled arcs, and labeled arcs. NEURBOPARSER does not use syntax. It first encodes an input sentence with a two-layer BiLSTM, and then computes part scores with two-layer tanh-MLPs. Inference is conducted with AD³ (Martins et al., 2015). To add syntactic features to NEURBOPARSER, we concatenate a token’s contextualized representation to that of its syntactic head, predicted by the intermediate parser. Formally, given length- n input sentence, we first run a BiLSTM. We use the concatenation of the two hidden representations $\mathbf{h}_j = [\vec{\mathbf{h}}_j; \overleftarrow{\mathbf{h}}_j]$ at each position j as the contextualized token representations. We then concatenate

\mathbf{h}_j with the representation of its head $\mathbf{h}_{\text{HEAD}(j)}$ by

$$\tilde{\mathbf{h}}_j = [\mathbf{h}_j; \mathbf{h}_{\text{HEAD}(j)}] = \left[\mathbf{h}_j; \sum_{i \neq j} \hat{z}_{\sigma(i \rightarrow j)} \mathbf{h}_i \right], \quad (7)$$

where $\hat{z} \in \mathbb{B}^{n(n-1)}$ is a binary encoding of the tree structure predicted by the intermediate parser. We then use $\tilde{\mathbf{h}}_j$ anywhere \mathbf{h}_j would have been used in NEURBOPARSER. In backpropagation, we compute $\nabla_{\hat{z}} L$ with an automatic differentiation toolkit (DyNet; Neubig et al., 2017).

We note that this approach can be generalized to convolutional neural networks over graphs (Mou et al., 2015; Duvenaud et al., 2015; Kipf and Welling, 2017, *inter alia*), recurrent neural networks along paths (Xu et al., 2015; Roth and Lapata, 2016, *inter alia*) or dependency trees (Tai et al., 2015). We choose to use concatenations to control the model’s complexity, and thus to better understand which parts of the model work.

We refer the readers to Kiperwasser and Goldberg (2016) and Peng et al. (2017) for further details of the parsing models.

Training procedure. Following previous work, we minimize structured hinge loss (Tsochantaridis et al., 2004) for both models. We jointly train both models from scratch, by randomly sampling an instance from the union of their training data at each step. In order to isolate the effect of backpropagation, we do not share any parameters between the two models.⁵ Implementation details are summarized in the supplementary materials.

4.1.2 Datasets

- For semantic dependencies, we use the English dataset from SemEval 2015 Task 18 (Oepen et al., 2015). Among the three formalisms provided by the shared task, we consider DELPH-IN MRS-derived dependencies (DM) and Prague Semantic Dependencies (PSD).⁶ It includes §00–19 of the WSJ corpus as training data, §20 and §21 for development and in-domain test data, resulting in a 33,961/1,692/1,410 train/dev./test split, and

⁵ Parameter sharing has proved successful in many related tasks (Collobert and Weston, 2008; Søgaard and Goldberg, 2016; Ammar et al., 2016; Swayamdipta et al., 2016, 2017, *inter alia*), and could be easily combined with our approach.

⁶We drop the third (PAS) because its structure is highly predictable from parts-of-speech, making it less interesting.

Model	DM		PSD	
	UF	LF	UF	LF
NEURBOPARSER	–	89.4	–	77.6
FREDA3	–	90.4	–	78.5
PIPELINE	91.8	90.8	88.4	78.1
SA	91.6	90.6	87.9	78.1
STE	92.0	91.1	88.9	78.9
SPIGOT	92.4	91.6	88.6	78.9

(a) F_1 on in-domain test set.

Model	DM		PSD	
	UF	LF	UF	LF
NEURBOPARSER	–	84.5	–	75.3
FREDA3	–	85.3	–	76.4
PIPELINE	87.4	85.8	85.5	75.6
SA	87.3	85.6	84.9	75.9
STE	87.7	86.4	85.8	76.6
SPIGOT	87.9	86.7	85.5	77.1

(b) F_1 on out-of-domain test set.

Table 1: Semantic dependency parsing performance in both unlabeled (UF) and labeled (LF) F_1 scores. Bold font indicates the best performance. Peng et al. (2017) does not report UF.

1,849 out-of-domain test instances from the Brown corpus.⁷

- For syntactic dependencies, we use the Stanford Dependency (de Marneffe and Manning, 2008) conversion of the the Penn Treebank WSJ portion (Marcus et al., 1993). To avoid data leak, we depart from standard split and use §20 and §21 as development and test data, and the remaining sections as training data. The number of training/dev./test instances is 40,265/2,012/1,671.

4.1.3 Baselines

We compare to the following baselines:

- A pipelined system (PIPELINE). The pre-trained parser achieves 92.9 test unlabeled attachment score (UAS).⁸

⁷The organizers remove, e.g., instances with cyclic graphs, and thus only a subset of the WSJ corpus is included. See Oepen et al. (2015) for details.

⁸Note that this number is not comparable to the parsing literature due to the different split. As a sanity check, we found in preliminary experiments that the same parser archi-

- Structured attention networks (SA; Kim et al., 2017). We use the inside-outside algorithm (Baker, 1979) to populate z with arcs’ marginal probabilities, use log-loss as the objective in training the intermediate parser.
- The straight-through estimator (STE; Hinton, 2012), introduced in §2.2.

4.1.4 Empirical Results

Table 1 compares the semantic dependency parsing performance of SPIGOT to all five baselines. FREDAS (Peng et al., 2017) is a state-of-the-art variant of NEURBOPARSER that is trained using multitask learning to jointly predict three different semantic dependency graph formalisms. Like the basic NEURBOPARSER model that we build from, FREDAS does not use any syntax. Strong DM performance is achieved in a more recent work by using joint learning and an ensemble (Peng et al., 2018), which is beyond fair comparisons to the models discussed here.

We found that using syntactic information improves semantic parsing performance: using pipelined syntactic head features brings 0.5–1.4% absolute labeled F_1 improvement to NEURBOPARSER. Such improvements are smaller compared to previous works, where dependency path and syntactic relation features are included (Almeida and Martins, 2015; Ribeyre et al., 2015; Zhang et al., 2016), indicating the potential to get better performance by using more syntactic information, which we leave to future work.

Both STE and SPIGOT use hard syntactic features. By allowing backpropagation into the intermediate syntactic parser, they both consistently outperform PIPELINE. On the other hand, when marginal syntactic tree structures are used, SA outperforms PIPELINE only on the out-of-domain PSD test set, and improvements under other cases are not observed.

Compared to STE, SPIGOT outperforms STE on DM by more than 0.3% absolute labeled F_1 , both in-domain and out-of-domain. For PSD, SPIGOT achieves similar performance to STE on in-domain test set, but has a 0.5% absolute labeled F_1 improvement on out-of-domain data, where syntactic parsing is less accurate.

ture achieves 93.5 UAS when trained and evaluated with the standard split, close to the results reported by Kiperwasser and Goldberg (2016).

4.2 Semantic Dependencies for Sentiment Classification

Our second experiment uses semantic dependency graphs to improve sentiment classification performance. We are not aware of any efficient algorithm that solves marginal inference for semantic dependency graphs under determinism constraints, so we do not include a comparison to SA.

4.2.1 Architectures

Here we use NEURBOPARSER as the intermediate model, as described in §4.1.1, but with no syntactic enhancements.

Sentiment classifier. We first introduce a baseline that does not use any structural information. It learns a one-layer BiLSTM to encode the input sentence, and then feeds the sum of all hidden states into a two-layer ReLU-MLP.

To use semantic dependency features, we concatenate a word’s BiLSTM-encoded representation to the averaged representation of its heads, together with the corresponding semantic roles, similarly to that in Equation 7.⁹ Then the concatenation is fed into an affine transformation followed by a ReLU activation. The rest of the model is kept the same as the BiLSTM baseline.

Training procedure. We use structured hinge loss to train the semantic dependency parser, and log-loss for the sentiment classifier. Due to the discrepancy in the training data size of the two tasks (33K vs. 7K), we pre-train a semantic dependency parser, and then adopt joint training together with the classifier. In the joint training stage, we randomly sample 20% of the semantic dependency training instances each epoch. Implementations are detailed in the supplementary materials.

4.2.2 Datasets

For semantic dependencies, we use the DM dataset introduced in §4.1.2.

We consider a binary classification task using the Stanford Sentiment Treebank (Socher et al., 2013). It consists of roughly 10K movie review sentences from Rotten Tomatoes. The full dataset includes a rating on a scale from 1 to 5 for each constituent (including the full sentences), resulting in more than 200K instances. Following previous work (Iyyer et al., 2015), we only use full-sentence

⁹In a well-formed semantic dependency graph, a token may have multiple heads. Therefore we use average instead of the sum in Equation 7.

Model	Accuracy (%)
BiLSTM	84.8
PIPELINE	85.7
STE	85.4
SPIGOT	86.3

Table 2: Test accuracy of sentiment classification on Stanford Sentiment Treebank. Bold font indicates the best performance.

instances, with neutral instances excluded (3s) and the remaining four rating levels converted to binary “positive” or “negative” labels. This results in a 6,920/872/1,821 train/dev./test split.

4.2.3 Empirical Results

Table 2 compares our SPIGOT method to three baselines. Pipelined semantic dependency predictions brings 0.9% absolute improvement in classification accuracy, and SPIGOT outperforms all baselines. In this task STE achieves slightly worse performance than a fixed pre-trained PIPELINE.

5 Analysis

We examine here how the intermediate model is affected by the end-task training signal. Is the end-task signal able to “overrule” intermediate predictions?

We use the syntactic-then-semantic parsing model (§4.1) as a case study. Table 3 compares a pipelined system to one jointly trained using SPIGOT. We consider the development set instances where both syntactic and semantic annotations are available, and partition them based on whether the two systems’ syntactic predictions agree (SAME), or not (DIFF). The second group includes sentences with much lower syntactic parsing accuracy (91.3 vs. 97.4 UAS), and SPIGOT further reduces this to 89.6. Even though these changes hurt syntactic parsing accuracy, they lead to a 1.1% absolute gain in labeled F_1 for semantic parsing. Furthermore, SPIGOT has an overall less detrimental effect on the intermediate parser than STE: using SPIGOT, intermediate dev. parsing UAS drops to 92.5 from the 92.9 pipelined performance, while STE reduces it to 91.8.

We then take a detailed look and categorize the changes in intermediate trees by their correlations with the semantic graphs. Specifically, when a modifier m ’s head is changed from h to h' in the

Split	# Sent.	Model	UAS	DM
SAME	1011	PIPELINE	97.4	94.0
		SPIGOT	97.4	94.3
DIFF	681	PIPELINE	91.3	88.1
		SPIGOT	89.6	89.2

Table 3: Syntactic parsing performance (in unlabeled attachment score, UAS) and DM semantic parsing performance (in labeled F_1) on different groups of the development data. Both systems predict the same syntactic parses for instances from SAME, and they disagree on instances from DIFF (§5).

tree, we consider three cases: (a) h' is a head of m in the semantic graph; (b) h' is a modifier of m in the semantic graph; (c) h is the modifier of m in the semantic graph. The first two reflect modifications to the syntactic parse that rearrange semantically linked words to be neighbors. Under (c), the semantic parser removes a syntactic dependency that reverses the direction of a semantic dependency. These cases account for 17.6%, 10.9%, and 12.8%, respectively (41.2% combined) of the total changes. Making these changes, of course, is complicated, since they often require *other* modifications to maintain well-formedness of the tree. Figure 2 gives an example.

6 Related Work

Joint learning in NLP pipelines. To avoid cascading errors, much effort has been devoted to joint decoding in NLP pipelines (Habash and Rambow, 2005; Cohen and Smith, 2007; Goldberg and Tsarfaty, 2008; Lewis et al., 2015; Zhang et al., 2015, *inter alia*). However, joint inference can sometimes be prohibitively expensive. Recent advances in representation learning facilitate exploration in the joint learning of multiple tasks by sharing parameters (Collobert and Weston, 2008; Blitzer et al., 2006; Finkel and Manning, 2010; Zhang and Weiss, 2016; Hashimoto et al., 2017, *inter alia*).

Differentiable optimization. Gould et al. (2016) review the generic approaches to differentiation in bi-level optimization (Bard, 2010; Kunisch and Pock, 2013). Amos and Kolter (2017) extend their efforts to a class of subdifferentiable quadratic programs. However, they both require that the intermediate objective has an invertible Hessian, limiting their application

in NLP. In another line of work, the steps of a gradient-based optimization procedure are unrolled into a single computation graph (Stoyanov et al., 2011; Domke, 2012; Goodfellow et al., 2013; Brakel et al., 2013). This comes at a high computational cost due to the second-order derivative computation during backpropagation. Moreover, constrained optimization problems (like many NLP problems) often require projection steps within the procedure, which can be difficult to differentiate through (Belanger and McCallum, 2016; Belanger et al., 2017).

7 Conclusion

We presented SPIGOT, a novel approach to back-propagating through neural network architectures that include discrete structured decisions in intermediate layers. SPIGOT devises a proxy for the gradients with respect to argmax 's inputs, employing a projection that aims to respect the constraints in the intermediate task. We empirically evaluate our method with two architectures: a semantic parser with an intermediate syntactic parser, and a sentiment classifier with an intermediate semantic parser. Experiments show that SPIGOT achieves stronger performance than baselines under both settings, and outperforms state-of-the-art systems on semantic dependency parsing. Our implementation is available at <https://github.com/Noahs-ARK/SPIGOT>.

Acknowledgments

We thank the ARK, Julian Michael, Minjoon Seo, Eunsol Choi, and Maxwell Forbes for their helpful comments on an earlier version of this work, and the anonymous reviewers for their valuable feedback. This work was supported in part by NSF grant IIS-1562364.

References

Mariana S. C. Almeida and André F. T. Martins. 2015. Lisbon: Evaluating TurboSemanticParser on multiple languages and out-of-domain data. In *Proc. of SemEval*.

Waleed Ammar, George Mulcaire, Miguel Ballesteros, Chris Dyer, and Noah A. Smith. 2016. Many languages, one parser. *TACL* 4:431–444.

Brandon Amos and J. Zico Kolter. 2017. OptNet: Differentiable optimization as a layer in neural networks. In *Proc. of ICML*.

J. K. Baker. 1979. Trainable grammars for speech recognition. In *Speech Communication Papers for the 97th Meeting of the Acoustical Society of America*.

Jonathan F. Bard. 2010. *Practical Bilevel Optimization: Algorithms and Applications*. Springer.

David Belanger and Andrew McCallum. 2016. Structured prediction energy networks. In *Proc. of ICML*.

David Belanger, Bishan Yang, and Andrew McCallum. 2017. End-to-end learning for structured prediction energy networks. In *Proc. of ICML*.

Yoshua Bengio, Nicholas Leonard, and Aaron Courville. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. arXiv:1308.3432.

Dimitris Bertsimas and John Tsitsiklis. 1997. *Introduction to Linear Optimization*. Athena Scientific.

John Blitzer, Ryan McDonald, and Fernando Pereira. 2006. Domain adaptation with structural correspondence learning. In *Proc. of EMNLP*.

Philémon Brakel, Dirk Stroobandt, and Benjamin Schrauwen. 2013. Training energy-based models for time-series imputation. *Journal of Machine Learning Research* 14:2771–2797.

Peter Brucker. 1984. An $O(n)$ algorithm for quadratic knapsack problems. *Operations Research Letters* 3(3):163 – 166.

Jihun Choi, Kang Min Yoo, and Sang-goo Lee. 2017. Unsupervised learning of task-specific tree structures with tree-LSTMs. arXiv:1707.02786.

Yoeng-Jin Chu and Tseng-Hong Liu. 1965. On the shortest arborescence of a directed graph. *Science Sinica* 14:1396–1400.

Shay B. Cohen and Noah A. Smith. 2007. Joint morphological and syntactic disambiguation. In *Proc. of EMNLP*.

Ronan Collobert and Jason Weston. 2008. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proc. of ICML*.

Marie-Catherine de Marneffe and Christopher D. Manning. 2008. Stanford typed dependencies manual. Technical report, Stanford University.

Justin Domke. 2012. Generic methods for optimization-based modeling. In *Proc. of AIS-TATS*.

David K. Duvenaud, Dougal Maclaurin, Jorge Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alan Aspuru-Guzik, and Ryan P. Adams. 2015. Convolutional networks on graphs for learning molecular fingerprints. In *Proc. of NIPS*.

- Jack Edmonds. 1967. Optimum branchings. *Journal of Research of the National Bureau of Standards* 71B:233–240.
- Jason Eisner. 2000. Bilexical grammars and their cubic-time parsing algorithms. In *Advances in Probabilistic and Other Parsing Technologies*, Springer Netherlands, pages 29–61.
- Jason Eisner. 2016. Inside-outside and forward-backward algorithms are just backprop. In *Proceedings of the EMNLP Workshop on Structured Prediction for NLP*.
- Jason M. Eisner. 1996. Three new probabilistic models for dependency parsing: An exploration. In *Proc. of COLING*.
- Jenny Rose Finkel and Christopher D. Manning. 2010. Hierarchical joint learning: Improving joint parsing and named entity recognition with non-jointly labeled data. In *Proc. of ACL*.
- Abram L. Friesen and Pedro M. Domingos. 2016. The sum-product theorem: A foundation for learning tractable models. In *Proc. of ICML*.
- Abram L. Friesen and Pedro M. Domingos. 2018. Deep learning as a mixed convex-combinatorial optimization problem. In *Proc. of ICLR*.
- Yoav Goldberg and Reut Tsarfaty. 2008. A single generative model for joint morphological segmentation and syntactic parsing. In *Proc. of ACL*.
- Ian Goodfellow, Mehdi Mirza, Aaron Courville, and Yoshua Bengio. 2013. Multi-prediction deep Boltzmann machines. In *Proc. of NIPS*.
- Stephen Gould, Basura Fernando, Anoop Cherian, Peter Anderson, Rodrigo Santa Cruz, and Edison Guo. 2016. On differentiating parameterized argmin and argmax problems with application to bi-level optimization. arXiv:1607.05447.
- Nizar Habash and Owen Rambow. 2005. Arabic tokenization, part-of-speech tagging and morphological disambiguation in one fell swoop. In *Proc. ACL*.
- Kazuma Hashimoto, Caiming Xiong, Yoshimasa Tsuruoka, and Richard Socher. 2017. A joint many-task model: Growing a neural network for multiple NLP tasks. In *Proc. of EMNLP*.
- Luheng He, Kenton Lee, Mike Lewis, and Luke Zettlemoyer. 2017. Deep semantic role labeling: What works and whats next. In *Proc. of ACL*.
- Geoffrey Hinton. 2012. Neural networks for machine learning. *Coursera* video lectures.
- Mohit Iyyer, Varun Manjunatha, Jordan Boyd-Graber, and Hal Daumé III. 2015. Deep unordered composition rivals syntactic methods for text classification. In *Proc. of ACL*.
- Eric Jang, Shixiang Gu, and Ben Poole. 2016. Categorical reparameterization with Gumbel-Softmax. arXiv:1611.01144.
- Yangfeng Ji and Noah A. Smith. 2017. Neural discourse structure for text categorization. In *Proc. of ACL*.
- Yoon Kim, Carl Denton, Luong Hoang, and Alexander M. Rush. 2017. Structured attention networks. In *Proc. of ICLR*.
- Eliyahu Kiperwasser and Yoav Goldberg. 2016. Simple and accurate dependency parsing using bidirectional LSTM feature representations. *TACL* 4:313–327.
- Thomas N. Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *Proc. of ICLR*.
- Karl Kunisch and Thomas Pock. 2013. A bilevel optimization approach for parameter learning in variational models. *SIAM Journal on Imaging Sciences* 6(2):938–983.
- Mike Lewis, Luheng He, and Luke Zettlemoyer. 2015. Joint A* CCG parsing and semantic role labelling. In *Proc. of EMNLP*.
- Yang Liu and Mirella Lapata. 2018. Learning structured text representations. *TACL* 6:63–75.
- David G. Luenberger and Yinyu Ye. 2015. *Linear and Nonlinear Programming*. Springer.
- Mitchell P. Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. 1993. Building a large annotated corpus of english: The penn treebank. *Computational Linguistics* 19(2):313–330.
- Andre Martins and Ramon Astudillo. 2016. From softmax to sparsemax: A sparse model of attention and multi-label classification. In *Proc. of ICML*.
- André F. T. Martins, Miguel B. Almeida, and Noah A. Smith. 2013. Turning on the turbo: Fast third-order non-projective turbo parsers. In *Proc. of ACL*.
- André F. T. Martins, Mário A. T. Figueiredo, Pedro M. Q. Aguiar, Noah A. Smith, and Eric P. Xing. 2015. AD3: Alternating directions dual decomposition for map inference in graphical models. *Journal of Machine Learning Research* 16:495–545.
- André F. T. Martins, Noah A. Smith, and Eric P. Xing. 2009. Polyhedral outer approximations with application to natural language parsing. In *Proc. of ICML*.
- Arthur Mensch and Mathieu Blondel. 2018. Differentiable dynamic programming for structured prediction and attention. arXiv:1802.03676 .
- Lili Mou, Hao Peng, Ge Li, Yan Xu, Lu Zhang, and Zhi Jin. 2015. Discriminative neural sentence modeling by tree-based convolution. In *Proc. of EMNLP*.

- Graham Neubig, Chris Dyer, Yoav Goldberg, Austin Matthews, Waleed Ammar, Antonios Anastasopoulos, Miguel Ballesteros, David Chiang, Daniel Clothiaux, Trevor Cohn, Kevin Duh, Manaal Faruqui, Cynthia Gan, Dan Garrette, Yangfeng Ji, Lingpeng Kong, Adhiguna Kuncoro, Gaurav Kumar, Chaitanya Malaviya, Paul Michel, Yusuke Oda, Matthew Richardson, Naomi Saphra, Swabha Swayamdipta, and Pengcheng Yin. 2017. DyNet: The dynamic neural network toolkit. arXiv:1701.03980.
- Vlad Niculae and Mathieu Blondel. 2017. A regularized framework for sparse and structured neural attention. In *Proc. of NIPS*.
- Vlad Niculae, Andr F. T. Martins, Mathieu Blondel, and Claire Cardie. 2018. SparseMAP: Differentiable sparse structured inference. arXiv:1802.04223.
- Stephan Open, Marco Kuhlmann, Yusuke Miyao, Daniel Zeman, Silvie Cinkova, Dan Flickinger, Jan Hajič, and Zdenka Uresova. 2015. SemEval 2015 task 18: Broad-coverage semantic dependency parsing. In *Proc. of SemEval*.
- Stephan Open, Marco Kuhlmann, Yusuke Miyao, Daniel Zeman, Silvie Cinková, Dan Flickinger, Jan Hajič, Angelina Ivanova, and Zdeňka Urešová. 2016. Towards comparability of linguistic graph banks for semantic parsing. In *Proc. of LREC*.
- Stephan Open, Marco Kuhlmann, Yusuke Miyao, Daniel Zeman, Dan Flickinger, Jan Hajič, Angelina Ivanova, and Yi Zhang. 2014. SemEval 2014 task 8: Broad-coverage semantic dependency parsing. In *Proc. of SemEval*.
- Stephan Open, Lilja vrelid, Jari Björne, Richard Johansson, Emanuele Lapponi, Filip Ginter, and Erik Velldal. 2017. The 2017 shared task on extrinsic parser evaluation. towards a reusable community infrastructure. In *Proc. of the 2017 Shared Task on Extrinsic Parser Evaluation*.
- Hao Peng, Sam Thomson, and Noah A. Smith. 2017. Deep multitask learning for semantic dependency parsing. In *Proc. of ACL*.
- Hao Peng, Sam Thomson, Swabha Swayamdipta, and Noah A. Smith. 2018. Learning joint semantic parsers from disjoint data. In *Proc. of NAACL*.
- Corentin Ribeyre, Éric Villemonte De La Clergerie, and Djamel Seddah. 2015. Because syntax does matter: Improving predicate-argument structures parsing using syntactic features. In *Proc. of NAACL*.
- Dan Roth and Wen-tau Yih. 2004. A linear programming formulation for global inference in natural language tasks. In *Proc. of NAACL*.
- Michael Roth and Mirella Lapata. 2016. Neural semantic role labeling with dependency path embeddings. In *Proc. of ACL*.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proc. of EMNLP*.
- Anders Søgaard and Yoav Goldberg. 2016. Deep multi-task learning with low level tasks supervised at lower layers. In *Proc. of ACL*.
- Veselin Stoyanov, Alexander Ropson, and Jason Eisner. 2011. Empirical risk minimization of graphical model parameters given approximate inference, decoding, and model structure. In *Proc. of AISTATS*.
- Swabha Swayamdipta, Miguel Ballesteros, Chris Dyer, and Noah A. Smith. 2016. Greedy, joint syntactic-semantic parsing with stack LSTMs. In *Proc. of CoNLL*.
- Swabha Swayamdipta, Sam Thomson, Chris Dyer, and Noah A. Smith. 2017. Frame-semantic parsing with softmax-margin segmental RNNs and a syntactic scaffold. arXiv:1706.09528.
- Kai Sheng Tai, Richard Socher, and Christopher D. Manning. 2015. Improved semantic representations from tree-structured long short-term memory networks. In *Proc. of ACL*.
- Ioannis Tsochantaris, Thomas Hofmann, Thorsten Joachims, and Yasemin Altun. 2004. Support vector machine learning for interdependent and structured output spaces. In *Proc. of ICML*.
- Ronald J. Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning* 8(3-4):229–256.
- Yan Xu, Lili Mou, Ge Li, Yunchuan Chen, Hao Peng, and Zhi Jin. 2015. Classifying relations via long short term memory networks along shortest dependency paths. In *Proc. of EMNLP*.
- Dani Yogatama, Phil Blunsom, Chris Dyer, Edward Grefenstette, and Wang Ling. 2017. Learning to compose words into sentences with reinforcement learning. In *Proc. of ICLR*.
- Xun Zhang, Yantao Du, Weiwei Sun, and Xiaojun Wan. 2016. Transition-based parsing for deep dependency structures. *Computational Linguistics* 42(3):353–389.
- Yuan Zhang, Chengtao Li, Regina Barzilay, and Kareem Darwish. 2015. Randomized greedy inference for joint segmentation, POS tagging and dependency parsing. In *Proc. NAACL*.
- Yuan Zhang and David Weiss. 2016. Stack-propagation: Improved representation learning for syntax. In *Proc. of ACL*.