



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Confidence Modeling for Neural Semantic Parsing

Citation for published version:

Dong, L, Quirk, C & Lapata, M 2018, Confidence Modeling for Neural Semantic Parsing. in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Melbourne, Australia, pp. 743-753, 56th Annual Meeting of the Association for Computational Linguistics, Melbourne, Australia, 15/07/18.
<<http://aclweb.org/anthology/P18-1069> >

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Confidence Modeling for Neural Semantic Parsing

Li Dong^{†*} and Chris Quirk[‡] and Mirella Lapata[†]

[†] School of Informatics, University of Edinburgh

[‡] Microsoft Research, Redmond

li.dong@ed.ac.uk chrisq@microsoft.com mlap@inf.ed.ac.uk

Abstract

In this work we focus on confidence modeling for neural semantic parsers which are built upon sequence-to-sequence models. We outline three major causes of uncertainty, and design various metrics to quantify these factors. These metrics are then used to estimate confidence scores that indicate whether model predictions are likely to be correct. Beyond confidence estimation, we identify which parts of the input contribute to uncertain predictions allowing users to interpret their model, and verify or refine its input. Experimental results show that our confidence model significantly outperforms a widely used method that relies on posterior probability, and improves the quality of interpretation compared to simply relying on attention scores.

1 Introduction

Semantic parsing aims to map natural language text to a formal meaning representation (e.g., logical forms or SQL queries). The neural sequence-to-sequence architecture (Sutskever et al., 2014; Bahdanau et al., 2015) has been widely adopted in a variety of natural language processing tasks, and semantic parsing is no exception. However, despite achieving promising results (Dong and Lapata, 2016; Jia and Liang, 2016; Ling et al., 2016), neural semantic parsers remain difficult to interpret, acting in most cases as a black box, not providing any information about what made them arrive at a particular decision. In this work, we explore ways to estimate and interpret the

model’s confidence in its predictions, which we argue can provide users with immediate and meaningful feedback regarding uncertain outputs.

An explicit framework for confidence modeling would benefit the development cycle of neural semantic parsers which, contrary to more traditional methods, do not make use of lexicons or templates and as a result the sources of errors and inconsistencies are difficult to trace. Moreover, from the perspective of application, semantic parsing is often used to build natural language interfaces, such as dialogue systems. In this case it is important to know whether the system understands the input queries with high confidence in order to make decisions more reliably. For example, knowing that some of the predictions are uncertain would allow the system to generate clarification questions, prompting users to verify the results before triggering unwanted actions. In addition, the training data used for semantic parsing can be small and noisy, and as a result, models do indeed produce uncertain outputs, which we would like our framework to identify.

A widely-used confidence scoring method is based on posterior probabilities $p(y|x)$ where x is the input and y the model’s prediction. For a linear model, this method makes sense: as more positive evidence is gathered, the score becomes larger. Neural models, in contrast, learn a complicated function that often overfits the training data. Posterior probability is effective when making decisions about model output, but is no longer a good indicator of confidence due in part to the nonlinearity of neural networks (Johansen and Socher, 2017). This observation motivates us to develop a confidence modeling framework for sequence-to-sequence models. We categorize the causes of uncertainty into three types, namely *model uncertainty*, *data uncertainty*, and *input uncertainty* and design different metrics to characterize them.

*Work carried out during an internship at Microsoft Research.

We compute these confidence metrics for a given prediction and use them as features in a regression model which is trained on held-out data to fit prediction F1 scores. At test time, the regression model’s outputs are used as confidence scores. Our approach does not interfere with the training of the model, and can be thus applied to various architectures, without sacrificing test accuracy. Furthermore, we propose a method based on backpropagation which allows to interpret model behavior by identifying which parts of the input contribute to uncertain predictions.

Experimental results on two semantic parsing datasets (IFTTT, Quirk et al. 2015; and DJANGO, Oda et al. 2015) show that our model is superior to a method based on posterior probability. We also demonstrate that thresholding confidence scores achieves a good trade-off between coverage and accuracy. Moreover, the proposed uncertainty backpropagation method yields results which are qualitatively more interpretable compared to those based on attention scores.

2 Related Work

Confidence Estimation Confidence estimation has been studied in the context of a few NLP tasks, such as statistical machine translation (Blatz et al., 2004; Ueffing and Ney, 2005; Soricut and Echiabi, 2010), and question answering (Gondek et al., 2012). To the best of our knowledge, confidence modeling for semantic parsing remains largely unexplored. A common scheme for modeling uncertainty in neural networks is to place distributions over the network’s weights (Denker and Lecun, 1991; MacKay, 1992; Neal, 1996; Blundell et al., 2015; Gan et al., 2017). But the resulting models often contain more parameters, and the training process has to be accordingly changed, which makes these approaches difficult to work with. Gal and Ghahramani (2016) develop a theoretical framework which shows that the use of dropout in neural networks can be interpreted as a Bayesian approximation of Gaussian Process. We adapt their framework so as to represent uncertainty in the encoder-decoder architectures, and extend it by adding Gaussian noise to weights.

Semantic Parsing Various methods have been developed to learn a semantic parser from natural language descriptions paired with meaning representations (Tang and Mooney, 2000; Zettlemoyer and Collins, 2007; Lu et al., 2008; Kwiatkowski

et al., 2011; Andreas et al., 2013; Zhao and Huang, 2015). More recently, a few sequence-to-sequence models have been proposed for semantic parsing (Dong and Lapata, 2016; Jia and Liang, 2016; Ling et al., 2016) and shown to perform competitively whilst eschewing the use of templates or manually designed features. There have been several efforts to improve these models including the use of a tree decoder (Dong and Lapata, 2016), data augmentation (Jia and Liang, 2016; Kočiský et al., 2016), the use of a grammar model (Xiao et al., 2016; Rabinovich et al., 2017; Yin and Neubig, 2017; Krishnamurthy et al., 2017), coarse-to-fine decoding (Dong and Lapata, 2018), network sharing (Susanto and Lu, 2017; Herzig and Berant, 2017), user feedback (Iyer et al., 2017), and transfer learning (Fan et al., 2017). Current semantic parsers will by default generate some output for a given input even if this is just a random guess. System results can thus be somewhat unexpected inadvertently affecting user experience. Our goal is to mitigate these issues with a confidence scoring model that can estimate how likely the prediction is correct.

3 Neural Semantic Parsing Model

In the following section we describe the neural semantic parsing model (Dong and Lapata, 2016; Jia and Liang, 2016; Ling et al., 2016) we assume throughout this paper. The model is built upon the sequence-to-sequence architecture and is illustrated in Figure 1. An *encoder* is used to encode natural language input $q = q_1 \cdots q_{|q|}$ into a vector representation, and a *decoder* learns to generate a logical form representation of its meaning $a = a_1 \cdots a_{|a|}$ conditioned on the encoding vectors. The encoder and decoder are two different recurrent neural networks with long short-term memory units (LSTMs; Hochreiter and Schmidhuber 1997) which process tokens sequentially. The probability of generating the whole sequence $p(a|q)$ is factorized as:

$$p(a|q) = \prod_{t=1}^{|a|} p(a_t | a_{<t}, q) \quad (1)$$

where $a_{<t} = a_1 \cdots a_{t-1}$.

Let $\mathbf{e}_t \in \mathbb{R}^n$ denote the hidden vector of the encoder at time step t . It is computed via $\mathbf{e}_t = \mathbf{f}_{\text{LSTM}}(\mathbf{e}_{t-1}, \mathbf{q}_t)$, where \mathbf{f}_{LSTM} refers to the LSTM unit, and $\mathbf{q}_t \in \mathbb{R}^n$ is the word embedding

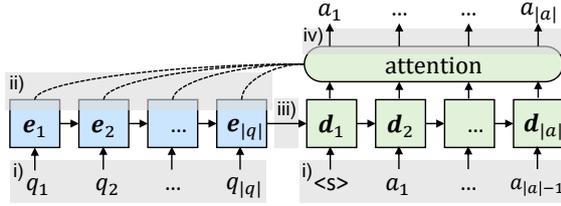


Figure 1: We use dropout as approximate Bayesian inference to obtain model uncertainty. The dropout layers are applied to i) token vectors; ii) the encoder’s output vectors; iii) bridge vectors; and iv) decoding vectors.

of q_t . Once the tokens of the input sequence are encoded into vectors, $\mathbf{e}_{|q|}$ is used to initialize the hidden states of the first time step in the decoder.

Similarly, the hidden vector of the decoder at time step t is computed by $\mathbf{d}_t = \text{f}_{\text{LSTM}}(\mathbf{d}_{t-1}, \mathbf{a}_{t-1})$, where $\mathbf{a}_{t-1} \in \mathbb{R}^n$ is the word vector of the previously predicted token. Additionally, we use an attention mechanism (Luong et al., 2015a) to utilize relevant encoder-side context. For the current time step t of the decoder, we compute its attention score with the k -th hidden state in the encoder as:

$$r_{t,k} \propto \exp\{\mathbf{d}_t \cdot \mathbf{e}_k\} \quad (2)$$

where $\sum_{j=1}^{|q|} r_{t,j} = 1$. The probability of generating a_t is computed via:

$$\mathbf{c}_t = \sum_{k=1}^{|q|} r_{t,k} \mathbf{e}_k \quad (3)$$

$$\mathbf{d}_t^{\text{att}} = \tanh(\mathbf{W}_1 \mathbf{d}_t + \mathbf{W}_2 \mathbf{c}_t) \quad (4)$$

$$p(a_t | a_{<t}, q) = \text{softmax}_{a_t}(\mathbf{W}_o \mathbf{d}_t^{\text{att}}) \quad (5)$$

where $\mathbf{W}_1, \mathbf{W}_2 \in \mathbb{R}^{n \times n}$ and $\mathbf{W}_o \in \mathbb{R}^{|V_a| \times n}$ are three parameter matrices.

The training objective is to maximize the likelihood of the generated meaning representation a given input q , i.e., maximize $\sum_{(q,a) \in \mathcal{D}} \log p(a|q)$, where \mathcal{D} represents training pairs. At test time, the model’s prediction for input q is obtained via $\hat{a} = \arg \max_{a'} p(a'|q)$, where a' represents candidate outputs. Because $p(a|q)$ is factorized as shown in Equation (1), we can use beam search to generate tokens one by one rather than iterating over all possible results.

4 Confidence Estimation

Given input q and its predicted meaning representation a , the confidence model estimates

Algorithm 1 Dropout Perturbation

Input: q, a : Input and its prediction

\mathcal{M} : Model parameters

1: **for** $i \leftarrow 1, \dots, F$ **do**

2: $\hat{\mathcal{M}}^i \leftarrow$ Apply dropout layers to \mathcal{M} \triangleright Figure 1

3: Run forward pass and compute $\hat{p}(a|q; \hat{\mathcal{M}}^i)$

4: Compute variance of $\{\hat{p}(a|q; \hat{\mathcal{M}}^i)\}_{i=1}^F$ \triangleright Equation (6)

score $s(q, a) \in (0, 1)$. A large score indicates the model is confident that its prediction is correct. In order to gauge confidence, we need to estimate “what we do not know”. To this end, we identify three causes of uncertainty, and design various metrics characterizing each one of them. We then feed these metrics into a regression model in order to predict $s(q, a)$.

4.1 Model Uncertainty

The model’s parameters or structures contain uncertainty, which makes the model less confident about the values of $p(a|q)$. For example, noise in the training data and the stochastic learning algorithm itself can result in model uncertainty. We describe metrics for capturing uncertainty below:

Dropout Perturbation Our first metric uses dropout (Srivastava et al., 2014) as approximate Bayesian inference to estimate model uncertainty (Gal and Ghahramani, 2016). Dropout is a widely used regularization technique during training, which relieves overfitting by randomly masking some input neurons to zero according to a Bernoulli distribution. In our work, we use dropout at *test time*, instead. As shown in Algorithm 1, we perform F forward passes through the network, and collect the results $\{\hat{p}(a|q; \hat{\mathcal{M}}^i)\}_{i=1}^F$ where $\hat{\mathcal{M}}^i$ represents the perturbed parameters. Then, the uncertainty metric is computed by the variance of results. We define the metric on the sequence level as:

$$\text{var}\{\hat{p}(a|q; \hat{\mathcal{M}}^i)\}_{i=1}^F. \quad (6)$$

In addition, we compute uncertainty u_{a_t} at the token-level a_t via:

$$u_{a_t} = \text{var}\{\hat{p}(a_t | a_{<t}, q; \hat{\mathcal{M}}^i)\}_{i=1}^F \quad (7)$$

where $\hat{p}(a_t | a_{<t}, q; \hat{\mathcal{M}}^i)$ is the probability of generating token a_t (Equation (5)) using perturbed model $\hat{\mathcal{M}}^i$. We operationalize token-level uncertainty in two ways, as the average score $\text{avg}\{u_{a_t}\}_{t=1}^{|a|}$ and the maximum score

$\max\{u_{a_t}\}_{t=1}^{|a|}$ (since the uncertainty of a sequence is often determined by the most uncertain token). As shown in Figure 1, we add dropout layers in i) the word vectors of the encoder and decoder $\mathbf{q}_t, \mathbf{a}_t$; ii) the output vectors of the encoder \mathbf{e}_t ; iii) bridge vectors $\mathbf{e}_{|q|}$ used to initialize the hidden states of the first time step in the decoder; and iv) decoding vectors \mathbf{d}_t^{att} (Equation (4)).

Gaussian Noise Standard dropout can be viewed as applying noise sampled from a Bernoulli distribution to the network parameters. We instead use Gaussian noise, and apply the metrics in the same way discussed above. Let \mathbf{v} denote a vector perturbed by noise, and \mathbf{g} a vector sampled from the Gaussian distribution $\mathcal{N}(0, \sigma^2)$. We use $\hat{\mathbf{v}} = \mathbf{v} + \mathbf{g}$ and $\hat{\mathbf{v}} = \mathbf{v} + \mathbf{v} \odot \mathbf{g}$ as two noise injection methods. Intuitively, if the model is more confident in an example, it should be more robust to perturbations.

Posterior Probability Our last class of metrics is based on posterior probability. We use the log probability $\log p(a|q)$ as a sequence-level metric. The token-level metric $\min\{p(a_t|a_{<t}, q)\}_{t=1}^{|a|}$ can identify the most uncertain predicted token. The perplexity per token $-\frac{1}{|a|} \sum_{t=1}^{|a|} \log p(a_t|a_{<t}, q)$ is also employed.

4.2 Data Uncertainty

The coverage of training data also affects the uncertainty of predictions. If the input q does not match the training distribution or contains unknown words, it is difficult to predict $p(a|q)$ reliably. We define two metrics:

Probability of Input We train a language model on the training data, and use it to estimate the probability of input $p(q|\mathcal{D})$ where \mathcal{D} represents the training data.

Number of Unknown Tokens Tokens that do not appear in the training data harm robustness, and lead to uncertainty. So, we use the number of unknown tokens in the input q as a metric.

4.3 Input Uncertainty

Even if the model can estimate $p(a|q)$ reliably, the input itself may be ambiguous. For instance, the input *the flight is at 9 o'clock* can be interpreted as either `flight.time(9am)` or `flight.time(9pm)`. Selecting between these predictions is difficult, especially if they are both highly likely. We use the

following metrics to measure uncertainty caused by ambiguous inputs.

Variance of Top Candidates We use the variance of the probability of the top candidates to indicate whether these are similar. The sequence-level metric is computed by:

$$\text{var}\{p(a^i|q)\}_{i=1}^K$$

where $a^1 \dots a^K$ are the K -best predictions obtained by the beam search during inference (Section 3).

Entropy of Decoding The sequence-level entropy of the decoding process is computed via:

$$H[a|q] = - \sum_{a'} p(a'|q) \log p(a'|q)$$

which we approximate by Monte Carlo sampling rather than iterating over all candidate predictions. The token-level metrics of decoding entropy are computed by $\text{avg}\{H[a_t|a_{<t}, q]\}_{t=1}^{|a|}$ and $\max\{H[a_t|a_{<t}, q]\}_{t=1}^{|a|}$.

4.4 Confidence Scoring

The sentence- and token-level confidence metrics defined in Section 4 are fed into a gradient tree boosting model (Chen and Guestrin, 2016) in order to predict the overall confidence score $s(q, a)$. The model is wrapped with a logistic function so that confidence scores are in the range of (0, 1).

Because the confidence score indicates whether the prediction is likely to be correct, we can use the prediction's F1 (see Section 6.2) as target value. The training loss is defined as:

$$\sum_{(q,a) \in \mathcal{D}} \ln(1+e^{-\hat{s}(q,a)})^{y_{q,a}} + \ln(1+e^{\hat{s}(q,a)})^{(1-y_{q,a})}$$

where \mathcal{D} represents the data, $y_{q,a}$ is the target F1 score, and $\hat{s}(q, a)$ the predicted confidence score. We refer readers to Chen and Guestrin (2016) for mathematical details of how the gradient tree boosting model is trained. Notice that we learn the confidence scoring model on the held-out set (rather than on the training data of the semantic parser) to avoid overfitting.

5 Uncertainty Interpretation

Confidence scores are useful in so far they can be traced back to the inputs causing the uncertainty in the first place. For semantic parsing, identifying

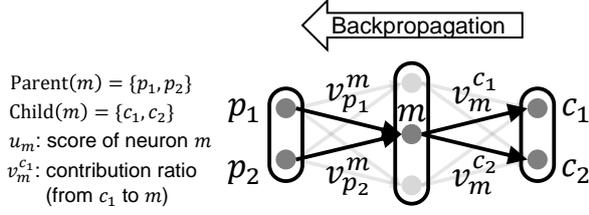


Figure 2: Uncertainty backpropagation at the neuron level. Neuron m 's score u_m is collected from child neurons c_1 and c_2 by $u_m = v_m^{c_1} u_{c_1} + v_m^{c_2} u_{c_2}$. The score u_m is then redistributed to its parent neurons p_1 and p_2 , which satisfies $v_{p_1}^m + v_{p_2}^m = 1$.

which input words contribute to uncertainty would be of value, e.g., these could be treated explicitly as special cases or refined if they represent noise.

In this section, we introduce an algorithm that backpropagates token-level uncertainty scores (see Equation (7)) from predictions to input tokens, following the ideas of Bach et al. (2015) and Zhang et al. (2016). Let u_m denote neuron m 's uncertainty score, which indicates the degree to which it contributes to uncertainty. As shown in Figure 2, u_m is computed by the summation of the scores backpropagated from its child neurons:

$$u_m = \sum_{c \in \text{Child}(m)} v_m^c u_c$$

where $\text{Child}(m)$ is the set of m 's child neurons, and the non-negative contribution ratio v_m^c indicates how much we backpropagate u_c to neuron m . Intuitively, if neuron m contributes more to c 's value, ratio v_m^c should be larger.

After obtaining score u_m , we redistribute it to its parent neurons in the same way. Contribution ratios from m to its parent neurons are normalized to 1:

$$\sum_{p \in \text{Parent}(m)} v_p^m = 1$$

where $\text{Parent}(m)$ is the set of m 's parent neurons.

Given the above constraints, we now define different backpropagation rules for the operators used in neural networks. We first describe the rules used for fully-connected layers. Let \mathbf{x} denote the input. The output is computed by $\mathbf{z} = \sigma(\mathbf{W}\mathbf{x} + \mathbf{b})$, where σ is a nonlinear function, $\mathbf{W} \in \mathbb{R}^{|\mathbf{z}| \times |\mathbf{x}|}$ is the weight matrix, $\mathbf{b} \in \mathbb{R}^{|\mathbf{z}|}$ is the bias, and neuron \mathbf{z}_i is computed via $\mathbf{z}_i = \sigma(\sum_{j=1}^{|\mathbf{x}|} \mathbf{W}_{i,j} \mathbf{x}_j + \mathbf{b}_i)$. Neuron \mathbf{x}_k 's uncertainty score u_{x_k} is gath-

Algorithm 2 Uncertainty Interpretation

Input: q, a : Input and its prediction

Output: $\{\hat{u}_{q_t}\}_{t=1}^{|q|}$: Interpretation scores for input tokens

Function: TokenUnc: Get token-level uncertainty

- 1: \triangleright Get token-level uncertainty for predicted tokens
- 2: $\{u_{a_t}\}_{t=1}^{|a|} \leftarrow \text{TokenUnc}(q, a)$
- 3: \triangleright Initialize uncertainty scores for backpropagation
- 4: **for** $t \leftarrow 1, \dots, |a|$ **do**
- 5: Decoder classifier's output neuron $\leftarrow u_{a_t}$
- 6: \triangleright Run backpropagation
- 7: **for** $m \leftarrow$ neuron in backward topological order **do**
- 8: \triangleright Gather scores from child neurons
- 9: $u_m \leftarrow \sum_{c \in \text{Child}(m)} v_m^c u_c$
- 10: \triangleright Summarize scores for input words
- 11: **for** $t \leftarrow 1, \dots, |q|$ **do**
- 12: $u_{q_t} \leftarrow \sum_{c \in \mathbf{q}_t} u_c$
- 13: $\{\hat{u}_{q_t}\}_{t=1}^{|q|} \leftarrow \text{normalize} \{u_{q_t}\}_{t=1}^{|q|}$

ered from the next layer:

$$u_{x_k} = \sum_{i=1}^{|\mathbf{z}|} v_{x_k}^{z_i} u_{z_i} = \sum_{i=1}^{|\mathbf{z}|} \frac{|\mathbf{W}_{i,k} \mathbf{x}_k|}{\sum_{j=1}^{|\mathbf{x}|} |\mathbf{W}_{i,j} \mathbf{x}_j|} u_{z_i}$$

ignoring the nonlinear function σ and the bias \mathbf{b} . The ratio $v_{x_k}^{z_i}$ is proportional to the contribution of \mathbf{x}_k to the value of \mathbf{z}_i .

We define backpropagation rules for element-wise vector operators. For $\mathbf{z} = \mathbf{x} \pm \mathbf{y}$, these are:

$$u_{x_k} = \frac{|\mathbf{x}_k|}{|\mathbf{x}_k| + |\mathbf{y}_k|} u_{z_k} \quad u_{y_k} = \frac{|\mathbf{y}_k|}{|\mathbf{x}_k| + |\mathbf{y}_k|} u_{z_k}$$

where the contribution ratios $v_{x_k}^{z_k}$ and $v_{y_k}^{z_k}$ are determined by $|\mathbf{x}_k|$ and $|\mathbf{y}_k|$. For multiplication, the contribution of two elements in $\frac{1}{3} * 3$ should be the same. So, the propagation rules for $\mathbf{z} = \mathbf{x} \odot \mathbf{y}$ are:

$$u_{x_k} = \frac{|\log |\mathbf{x}_k||}{|\log |\mathbf{x}_k|| + |\log |\mathbf{y}_k||} u_{z_k} \quad u_{y_k} = \frac{|\log |\mathbf{y}_k||}{|\log |\mathbf{x}_k|| + |\log |\mathbf{y}_k||} u_{z_k}$$

where the contribution ratios are determined by $|\log |\mathbf{x}_k||$ and $|\log |\mathbf{y}_k||$.

For scalar multiplication, $\mathbf{z} = \lambda \mathbf{x}$ where λ denotes a constant. We directly assign \mathbf{z} 's uncertainty scores to \mathbf{x} and the backpropagation rule is $u_{x_k} = u_{z_k}$.

As shown in Algorithm 2, we first initialize uncertainty backpropagation in the decoder (lines 1–5). For each predicted token a_t , we compute its uncertainty score u_{a_t} as in Equation (7). Next, we find the dimension of a_t in the decoder's softmax classifier (Equation (5)), and initialize the neuron with the uncertainty score u_{a_t} . We then backpropagate these uncertainty scores through

Dataset	Example
IFTTT	<i>turn android phone to full volume at 7am monday to friday</i> <code>date_time=every_day_of_the_week_at-((time_of_day (07)(:)(00)) (days_of_the_week (1)(2)(3)(4)(5))) THEN android_device-set_ringtones-volume-(volume ({'volume_level':1.0,'name':'100%'}))</code>
DJANGO	<i>for every key in sorted list of user_settings</i> <code>for key in sorted(user_settings):</code>

Table 1: Natural language descriptions and their meaning representations from IFTTT and DJANGO.

the network (lines 6–9), and finally into the neurons of the input words. We summarize them and compute the token-level scores for interpreting the results (line 10–13). For input word vector \mathbf{q}_t , we use the summation of its neuron-level scores as the token-level score:

$$\hat{u}_{q_t} \propto \sum_{c \in \mathbf{q}_t} u_c$$

where $c \in \mathbf{q}_t$ represents the neurons of word vector \mathbf{q}_t , and $\sum_{t=1}^{|\mathbf{q}_t|} \hat{u}_{q_t} = 1$. We use the normalized score \hat{u}_{q_t} to indicate token q_t 's contribution to prediction uncertainty.

6 Experiments

In this section we describe the datasets used in our experiments and various details concerning our models. We present our experimental results and analysis of model behavior. Our code is publicly available at <https://github.com/donglixp/confidence>.

6.1 Datasets

We trained the neural semantic parser introduced in Section 3 on two datasets covering different domains and meaning representations. Examples are shown in Table 1.

IFTTT This dataset (Quirk et al., 2015) contains a large number of if-this-then-that programs crawled from the IFTTT website. The programs are written for various applications, such as home security (e.g., “*email me if the window opens*”), and task automation (e.g., “*save instagram photos to dropbox*”). Whenever a program’s trigger is satisfied, an action is performed. Triggers and actions represent functions with arguments; they are selected from different channels (160 in total) representing various services (e.g., Android). There are 552 trigger functions and 229 action functions. The original split contains 77, 495 training, 5, 171 development, and 4, 294 test instances. The

subset that removes non-English descriptions was used in our experiments.

DJANGO This dataset (Oda et al., 2015) is built upon the code of the Django web framework. Each line of Python code has a manually annotated natural language description. Our goal is to map the English pseudo-code to Python statements. This dataset contains diverse use cases, such as iteration, exception handling, and string manipulation. The original split has 16, 000 training, 1, 000 development, and 1, 805 test examples.

6.2 Settings

We followed the data preprocessing used in previous work (Dong and Lapata, 2016; Yin and Neubig, 2017). Input sentences were tokenized using NLTK (Bird et al., 2009) and lowercased. We filtered words that appeared less than four times in the training set. Numbers and URLs in IFTTT and quoted strings in DJANGO were replaced with place holders. Hyperparameters of the semantic parsers were validated on the development set. The learning rate and the smoothing constant of RMSProp (Tieleman and Hinton, 2012) were 0.002 and 0.95, respectively. The dropout rate was 0.25. A two-layer LSTM was used for IFTTT, while a one-layer LSTM was employed for DJANGO. Dimensions for the word embedding and hidden vector were selected from {150, 250}. The beam size during decoding was 5.

For IFTTT, we view the predicted trees as a set of productions, and use balanced F1 as evaluation metric (Quirk et al., 2015). We do not measure accuracy because the dataset is very noisy and there rarely is an exact match between the predicted output and the gold standard. The F1 score of our neural semantic parser is 50.1%, which is comparable to Dong and Lapata (2016). For DJANGO, we measure the fraction of exact matches, where F1 score is equal to accuracy. Because there are unseen variable names at test time, we use attention scores as alignments to replace unknown to-

Method	IFTTT	DJANGO
POSTERIOR	0.477	0.694
CONF	0.625	0.793
– MODEL	0.595	0.759
– DATA	0.610	0.787
– INPUT	0.608	0.785

Table 2: Spearman ρ correlation between confidence scores and F1. Best results are shown in **bold**. All correlations are significant at $p < 0.01$.

kens in the prediction with the input words they align to (Luong et al., 2015b). The accuracy of our parser is 53.7%, which is better than the result (45.1%) of the sequence-to-sequence model reported in Yin and Neubig (2017).

To estimate model uncertainty, we set dropout rate to 0.1, and performed 30 inference passes. The standard deviation of Gaussian noise was 0.05. The language model was estimated using KenLM (Heafield et al., 2013). For input uncertainty, we computed variance for the 10-best candidates. The confidence metrics were implemented in batch mode, to take full advantage of GPUs. Hyperparameters of the confidence scoring model were cross-validated. The number of boosted trees was selected from $\{20, 50\}$. The maximum tree depth was selected from $\{3, 4, 5\}$. We set the subsample ratio to 0.8. All other hyperparameters in XGBoost (Chen and Guestrin, 2016) were left with their default values.

6.3 Results

Confidence Estimation We compare our approach (CONF) against confidence scores based on posterior probability $p(a|q)$ (POSTERIOR). We also report the results of three ablation variants (–MODEL, –DATA, –INPUT) by removing each group of confidence metrics described in Section 4. We measure the relationship between confidence scores and F1 using Spearman’s ρ correlation coefficient which varies between -1 and 1 (0 implies there is no correlation). High ρ indicates that the confidence scores are high for correct predictions and low otherwise.

As shown in Table 2, our method CONF outperforms POSTERIOR by a large margin. The ablation results indicate that model uncertainty plays the most important role among the confidence metrics. In contrast, removing the metrics of data uncertainty affects performance less, because most examples in the datasets are in-domain. Improve-

	F1	Dout	Noise	PR	PPL	LM	#UNK	Var
Dout	0.59							
Noise	0.59	0.90						
PR	0.52	0.84	0.82					
PPL	0.48	0.78	0.78	0.89				
LM	0.30	0.26	0.32	0.27	0.25			
#UNK	0.27	0.31	0.33	0.29	0.25	0.32		
Var	0.49	0.83	0.78	0.88	0.79	0.25	0.27	
Ent	0.53	0.78	0.78	0.80	0.75	0.27	0.30	0.76

Table 3: Correlation matrix for F1 and individual confidence metrics on the IFTTT dataset. All correlations are significant at $p < 0.01$. Best predictors are shown in **bold**. Dout is short for dropout, PR for posterior probability, PPL for perplexity, LM for probability based on a language model, #UNK for number of unknown tokens, Var for variance of top candidates, and Ent for Entropy.

	F1	Dout	Noise	PR	PPL	LM	#UNK	Var
Dout	0.76							
Noise	0.78	0.94						
PR	0.73	0.89	0.90					
PPL	0.64	0.80	0.81	0.84				
LM	0.32	0.41	0.40	0.38	0.30			
#UNK	0.27	0.28	0.28	0.26	0.19	0.35		
Var	0.70	0.87	0.87	0.89	0.87	0.37	0.23	
Ent	0.72	0.89	0.90	0.92	0.86	0.38	0.26	0.90

Table 4: Correlation matrix for F1 and individual confidence metrics on the DJANGO dataset. All correlations are significant at $p < 0.01$. Best predictors are shown in **bold**. Same shorthands apply as in Table 3.

ments for each group of metrics are significant with $p < 0.05$ according to bootstrap hypothesis testing (Efron and Tibshirani, 1994).

Tables 3 and 4 show the correlation matrix for F1 and individual confidence metrics on the IFTTT and DJANGO datasets, respectively. As can be seen, metrics representing model uncertainty and input uncertainty are more correlated to each other compared with metrics capturing data uncertainty. Perhaps unsurprisingly metrics of the same group are highly inter-correlated since they model the same type of uncertainty. Table 5 shows the relative importance of individual metrics in the regression model. As importance score we use the average gain (i.e., loss reduction) brought by the confidence metric once added as feature to the branch of the decision tree (Chen and Guestrin, 2016). The results indicate that model uncertainty (Noise/Dropout/Posterior/Perplexity) plays

Metric	Dout	Noise	PR	PPL	LM	#UNK	Var	Ent
IFTTT	0.39	1.00	0.89	0.27	0.26	0.46	0.43	0.34
DJANGO	1.00	0.59	0.22	0.58	0.49	0.14	0.24	0.25

Table 5: Importance scores of confidence metrics (normalized by maximum value on each dataset). Best results are shown in **bold**. Same shorthands apply as in Table 3.

the most important role. On IFTTT, the number of unknown tokens (#UNK) and the variance of top candidates (var(K-best)) are also very helpful because this dataset is relatively noisy and contains many ambiguous inputs.

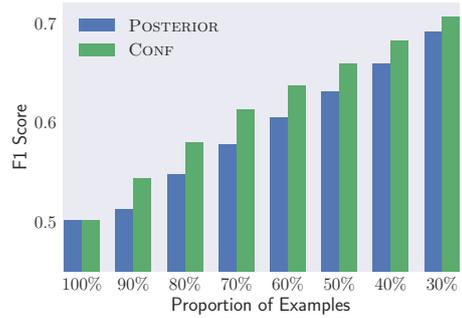
Finally, in real-world applications, confidence scores are often used as a threshold to trade-off precision for coverage. Figure 3 shows how F1 score varies as we increase the confidence threshold, i.e., reduce the proportion of examples that we return answers for. F1 score improves monotonically for POSTERIOR and our method, which, however, achieves better performance when coverage is the same.

Uncertainty Interpretation We next evaluate how our backpropagation method (see Section 5) allows us to identify input tokens contributing to uncertainty. We compare against a method that interprets uncertainty based on the attention mechanism (ATTENTION). As shown in Equation (2), attention scores $r_{t,k}$ can be used as soft alignments between the time step t of the decoder and the k -th input token. We compute the normalized uncertainty score \hat{u}_{q_t} for a token q_t via:

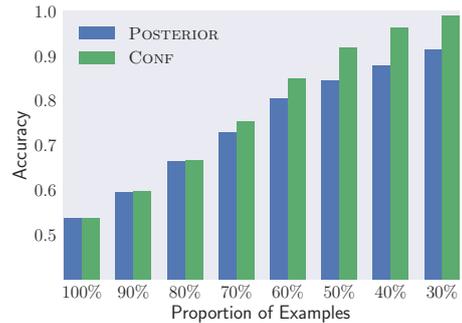
$$\hat{u}_{q_t} \propto \sum_{t=1}^{|a|} r_{t,k} u_{a_t} \quad (8)$$

where u_{a_t} is the uncertainty score of the predicted token a_t (Equation (7)), and $\sum_{t=1}^{|q|} \hat{u}_{q_t} = 1$.

Unfortunately, the evaluation of uncertainty interpretation methods is problematic. For our semantic parsing task, we do not a priori know which tokens in the natural language input contribute to uncertainty and these may vary depending on the architecture used, model parameters, and so on. We work around this problem by creating a proxy gold standard. We inject noise to the vectors representing tokens in the encoder (see Section 4.1) and then estimate the uncertainty caused by each token q_t (Equation (6)) under the assumption that



(a) IFTTT



(b) DJANGO

Figure 3: Confidence scores are used as threshold to filter out uncertain test examples. As the threshold increases, performance improves. The horizontal axis shows the proportion of examples beyond the threshold.

addition of noise should only affect genuinely uncertain tokens. Notice that here we inject noise to one token at a time¹ instead of all parameters (see Figure 1). Tokens identified as uncertain by the above procedure are considered gold standard and compared to those identified by our method. We use Gaussian noise to perturb vectors in our experiments (dropout obtained similar results).

We define an evaluation metric based on the overlap ($overlap@K$) among tokens identified as uncertain by the model and the gold standard. Given an example, we first compute the interpretation scores of the input tokens according to our method, and obtain a list τ_1 of K tokens with highest scores. We also obtain a list τ_2 of K tokens with highest ground-truth scores and measure the degree of overlap between these two lists:

$$overlap@K = \frac{|\tau_1 \cap \tau_2|}{K}$$

¹Noise injection as described above is used for evaluation purposes only since we need to perform forward passes multiple times (see Section 4.1) for each token, and the running time increases linearly with the input length.

Method	IFTT		DJANGO	
	@2	@4	@2	@4
ATTENTION	0.525	0.737	0.637	0.684
BACKPROP	0.608	0.791	0.770	0.788

Table 6: Uncertainty interpretation against inferred ground truth; we compute the overlap between tokens identified as contributing to uncertainty by our method and those found in the gold standard. Overlap is shown for top 2 and 4 tokens. Best results are in **bold**.

google_calendar-any_event_starts THEN facebook -create_a_status_message-(status_message ({description}))	ATT post calendar event to facebook BP post calendar event to facebook
feed-new_feed_item-(feed_url(url_sports.espn.go.com)) THEN ...	ATT espn mlb headline to readability BP espn mlb headline to readability
weather-tomorrow's_low_drops_below-((temperature(0)) (degrees_in(c))) THEN ...	ATT warn me when it's going to be freezing tomorrow BP warn me when it's going to be freezing tomorrow
if str_number[0] == '_STR_':	ATT if first element of str_number equals a string _STR_. BP if first element of str_number equals a string _STR_.
start = 0	ATT start is an integer 0. BP start is an integer 0.
if name.startswith('_STR_'):	ATT if name starts with an string _STR_. BP if name starts with an string _STR_.

Table 7: Uncertainty interpretation for ATTENTION (ATT) and BACKPROP (BP). The first line in each group is the model prediction. Predicted tokens and input words with large scores are shown in red and blue, respectively.

where $K \in \{2, 4\}$ in our experiments. For example, the overlap@4 metric of the lists $\tau_1 = [q_7, q_8, q_2, q_3]$ and $\tau_2 = [q_7, q_8, q_3, q_4]$ is $3/4$, because there are three overlapping tokens.

Table 6 reports results with overlap@2 and overlap@4. Overall, BACKPROP achieves better interpretation quality than the attention mechanism. On both datasets, about 80% of the top-4 tokens identified as uncertain agree with the ground truth. Table 7 shows examples where our method has identified input tokens contributing to the uncertainty of the output. We highlight token a_t if its uncertainty score u_{a_t} is greater than $0.5 * \text{avg}\{u_{a_{t'}}\}_{t'=1}^{|a|}$. The results illustrate that the parser tends to be uncertain about tokens which are

function arguments (e.g., URLs, and message content), and ambiguous inputs. The examples show that BACKPROP is qualitatively better compared to ATTENTION; attention scores often produce inaccurate alignments while BACKPROP can utilize information flowing through the LSTMs rather than only relying on the attention mechanism.

7 Conclusions

In this paper we presented a confidence estimation model and an uncertainty interpretation method for neural semantic parsing. Experimental results show that our method achieves better performance than competitive baselines on two datasets. Directions for future work are many and varied. The proposed framework could be applied to a variety of tasks (Bahdanau et al., 2015; Schmalz et al., 2017) employing sequence-to-sequence architectures. We could also utilize the confidence estimation model within an active learning framework for neural semantic parsing.

Acknowledgments

We would like to thank Pengcheng Yin for sharing with us the preprocessed version of the DJANGO dataset. We gratefully acknowledge the financial support of the European Research Council (award number 681760; Dong, Lapata) and the Adept-Mind Scholar Fellowship program (Dong).

References

- Jacob Andreas, Andreas Vlachos, and Stephen Clark. 2013. [Semantic parsing as machine translation](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, pages 47–52, Sofia, Bulgaria.
- Sebastian Bach, Alexander Binder, Grgoire Montavon, Frederick Klauschen, Klaus-Robert Miller, and Wojciech Samek. 2015. [On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation](#). *PLOS ONE*, 10(7):1–46.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *Proceedings of the 3rd International Conference on Learning Representations*, San Diego, California.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. [Natural Language Processing with Python](#). O'Reilly Media.
- John Blatz, Erin Fitzgerald, George Foster, Simona Gandrabur, Cyril Goutte, Alex Kulesza, Alberto

- Sanchis, and Nicola Ueffing. 2004. [Confidence estimation for machine translation](#). In *Proceedings of the 20th International Conference on Computational Linguistics*, pages 315–321, Geneva, Switzerland.
- Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. 2015. [Weight uncertainty in neural networks](#). In *Proceedings of the 32nd International Conference on International Conference on Machine Learning*, pages 1613–1622, Lille, France.
- Tianqi Chen and Carlos Guestrin. 2016. [XGBoost: A scalable tree boosting system](#). In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 785–794, San Francisco, California.
- John S Denker and Yann Lecun. 1991. [Transforming neural-net output levels to probability distributions](#). In *Advances in neural information processing systems*, pages 853–859, Denver, Colorado.
- Li Dong and Mirella Lapata. 2016. [Language to logical form with neural attention](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 33–43, Berlin, Germany.
- Li Dong and Mirella Lapata. 2018. Coarse-to-fine decoding for neural semantic parsing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, Melbourne, Australia.
- Bradley Efron and Robert J Tibshirani. 1994. *An Introduction to the Bootstrap*. CRC press.
- Xing Fan, Emilio Monti, Lambert Mathias, and Markus Dreyer. 2017. [Transfer learning for neural semantic parsing](#). In *Proceedings of the 2nd Workshop on Representation Learning for NLP*, pages 48–56, Vancouver, Canada.
- Yarin Gal and Zoubin Ghahramani. 2016. [Dropout as a bayesian approximation: Representing model uncertainty in deep learning](#). In *Proceedings of the 33rd International Conference on Machine Learning*, pages 1050–1059, New York City, NY.
- Zhe Gan, Chunyuan Li, Changyou Chen, Yunchen Pu, Qinliang Su, and Lawrence Carin. 2017. [Scalable bayesian learning of recurrent neural networks for language modeling](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 321–331, Vancouver, Canada.
- D. C. Gondek, A. Lally, A. Kalyanpur, J. W. Murdock, P. A. Duboue, L. Zhang, Y. Pan, Z. M. Qiu, and C. Welty. 2012. [A framework for merging and ranking of answers in DeepQA](#). *IBM Journal of Research and Development*, 56(3.4):14:1–14:12.
- Kenneth Heafield, Ivan Pouzyrevsky, Jonathan H. Clark, and Philipp Koehn. 2013. [Scalable modified Kneser-Ney language model estimation](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, pages 690–696, Sofia, Bulgaria.
- Jonathan Herzig and Jonathan Berant. 2017. [Neural semantic parsing over multiple knowledge-bases](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 623–628, Vancouver, Canada.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural Computation*, 9:1735–1780.
- Srinivasan Iyer, Ioannis Konstas, Alvin Cheung, Jayant Krishnamurthy, and Luke Zettlemoyer. 2017. [Learning a neural semantic parser from user feedback](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 963–973, Vancouver, Canada.
- Robin Jia and Percy Liang. 2016. [Data recombination for neural semantic parsing](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 12–22, Berlin, Germany.
- Alexander Johansen and Richard Socher. 2017. [Learning when to skim and when to read](#). In *Proceedings of the 2nd Workshop on Representation Learning for NLP*, pages 257–264, Vancouver, Canada.
- Tomáš Kočiský, Gábor Melis, Edward Grefenstette, Chris Dyer, Wang Ling, Phil Blunsom, and Karl Moritz Hermann. 2016. [Semantic parsing with semi-supervised sequential autoencoders](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1078–1087, Austin, Texas.
- Jayant Krishnamurthy, Pradeep Dasigi, and Matt Gardner. 2017. [Neural semantic parsing with type constraints for semi-structured tables](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1517–1527, Copenhagen, Denmark.
- Tom Kwiatkowski, Luke Zettlemoyer, Sharon Goldwater, and Mark Steedman. 2011. [Lexical generalization in CCG grammar induction for semantic parsing](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 1512–1523, Edinburgh, Scotland.
- Wang Ling, Phil Blunsom, Edward Grefenstette, Karl Moritz Hermann, Tomáš Kočiský, Fumin Wang, and Andrew Senior. 2016. [Latent predictor networks for code generation](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 599–609, Berlin, Germany.
- Wei Lu, Hwee Tou Ng, Wee Sun Lee, and Luke Zettlemoyer. 2008. [A generative model for parsing natural language to meaning representations](#). In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 783–792, Honolulu, Hawaii.

- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015a. [Effective approaches to attention-based neural machine translation](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421, Lisbon, Portugal.
- Thang Luong, Ilya Sutskever, Quoc Le, Oriol Vinyals, and Wojciech Zaremba. 2015b. [Addressing the rare word problem in neural machine translation](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, pages 11–19, Beijing, China.
- David J. C. MacKay. 1992. [A practical bayesian framework for backpropagation networks](#). *Neural Computation*, 4(3):448–472.
- Radford M Neal. 1996. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media.
- Yusuke Oda, Hiroyuki Fudaba, Graham Neubig, Hideaki Hata, Sakriani Sakti, Tomoki Toda, and Satoshi Nakamura. 2015. [Learning to generate pseudo-code from source code using statistical machine translation](#). In *Proceedings of the 2015 30th IEEE/ACM International Conference on Automated Software Engineering*, pages 574–584, Washington, DC.
- Chris Quirk, Raymond Mooney, and Michel Galley. 2015. [Language to code: Learning semantic parsers for if-this-then-that recipes](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, pages 878–888, Beijing, China.
- Maxim Rabinovich, Mitchell Stern, and Dan Klein. 2017. [Abstract syntax networks for code generation and semantic parsing](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 1139–1149, Vancouver, Canada.
- Allen Schmalz, Yoon Kim, Alexander Rush, and Stuart Shieber. 2017. [Adapting sequence models for sentence correction](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2797–2803, Copenhagen, Denmark.
- Radu Soricut and Abdessamad Echihabi. 2010. [Trustrank: Inducing trust in automatic translations via ranking](#). In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 612–621, Uppsala, Sweden.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. [Dropout: A simple way to prevent neural networks from overfitting](#). *Journal of Machine Learning Research*, 15:1929–1958.
- Raymond Hendy Susanto and Wei Lu. 2017. [Neural architectures for multilingual semantic parsing](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 38–44, Vancouver, Canada.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. [Sequence to sequence learning with neural networks](#). In *Advances in Neural Information Processing Systems*, pages 3104–3112, Montreal, Canada.
- Lappoon R. Tang and Raymond J. Mooney. 2000. [Automated construction of database interfaces: Integrating statistical and relational learning for semantic parsing](#). In *2000 Joint SIGDAT Conference on Empirical Methods in Natural Language Processing and Very Large Corpora*, pages 133–141, Hong Kong, China.
- T. Tieleman and G. Hinton. 2012. Lecture 6.5—RMSProp: Divide the gradient by a running average of its recent magnitude. Technical report.
- Nicola Ueffing and Hermann Ney. 2005. [Word-level confidence estimation for machine translation using phrase-based translation models](#). In *Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing*, pages 763–770, Vancouver, Canada.
- Chunyang Xiao, Marc Dymetman, and Claire Gardent. 2016. [Sequence-based structured prediction for semantic parsing](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 1341–1350, Berlin, Germany.
- Pengcheng Yin and Graham Neubig. 2017. [A syntactic neural model for general-purpose code generation](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 440–450, Vancouver, Canada.
- Luke Zettlemoyer and Michael Collins. 2007. [Online learning of relaxed CCG grammars for parsing to logical form](#). In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 678–687, Prague, Czech Republic.
- Jianming Zhang, Zhe Lin, Jonathan Brandt, Xiaohui Shen, and Stan Sclaroff. 2016. [Top-down neural attention by excitation backprop](#). In *European Conference on Computer Vision*, pages 543–559, Amsterdam, Netherlands.
- Kai Zhao and Liang Huang. 2015. [Type-driven incremental semantic parsing with polymorphism](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1416–1421, Denver, Colorado.