

Automatic Accuracy Prediction for AMR Parsing

Juri Opitz and Anette Frank

Research Training Group AIPHES

Leibniz ScienceCampus “Empirical Linguistics and Computational Language Modeling”

Department for Computational Linguistics

69120 Heidelberg

{opitz, frank}@cl.uni-heidelberg.de

Abstract

Abstract Meaning Representation (AMR) represents sentences as directed, acyclic and rooted graphs, aiming at capturing their meaning in a machine readable format. AMR parsing converts natural language sentences into such graphs. However, evaluating a parser on new data by means of comparison to manually created AMR graphs is very costly. Also, we would like to be able to detect parses of questionable quality, or preferring results of alternative systems by selecting the ones for which we can assess good quality. We propose AMR accuracy prediction as the task of predicting several metrics of correctness for an automatically generated AMR parse – in absence of the corresponding gold parse. We develop a neural end-to-end multi-output regression model and perform three case studies: firstly, we evaluate the model’s capacity of predicting AMR parse accuracies and test whether it can reliably assign high scores to gold parses. Secondly, we perform parse selection based on predicted parse accuracies of candidate parses from alternative systems, with the aim of improving overall results. Finally, we predict system ranks for submissions from two AMR shared tasks on the basis of their predicted parse accuracy averages. All experiments are carried out across two different domains and show that our method is effective.

1 Introduction

Abstract Meaning Representation (AMR) (Banasescu et al., 2013) represents the semantic structure of a sentence, including concepts, semantic operators and relations, sense-disambiguated predicates and their arguments. As a machine readable representation of the meaning of a sentence, AMR is potentially useful for many NLP tasks. Among other applications it has been used in machine translation (Jones et al., 2012), text

```
(a / asbestos
 :polarity -
 :time (n / now)
 :location (t / thing
 :ARG1-of (p / produce-01
 :ARG0 (w / we))))
```

Figure 1: Humanly produced AMR for: *There is no asbestos in our products now.* Numbered predicates refer to PropBank senses (Palmer et al., 2005).

summarization (Liu et al., 2015; Dohare and Karnick, 2017) and question answering (Mittra and Baral, 2016). Since the introduction of AMR, many approaches to AMR parsing have been proposed: graph-based pipeline systems which rely on an alignment step (Flanigan et al., 2014, 2016) or transition-based parsers relying on dependency annotation (Wang et al., 2015b,a, 2016a). In the following we will denote the former by **JAMR** and the latter by **CAMR**. More recently, end-to-end neural systems have been proposed which produce linearized AMR graphs within character-based (van Noord and Bos, 2017b) or word-based (Konstas et al., 2017) encoding models. Both approaches greatly profit from large amounts of silver training data. The silver data is obtained with self-training (Konstas et al., 2017) or the aid of additional parsers, where only parses with considerable agreement are chosen to extend the training data (van Noord and Bos, 2017b). Lyu and Titov (2018) formulate a neural model that jointly predicts alignments, concepts and relations. Their system – henceforth called **GPLA** (Graph Prediction with Latent Alignments) – defines the current state-of-the-art in AMR parsing.

A system that can perform accuracy prediction for AMR parsing can be used in a variety of ways: (i) *estimating the quality of downstream tasks* that deploy AMR parses. E.g., in a document summarization scenario, we might expect lower qual-

ity of a summary if the estimated quality of AMR parses used as a basis for the summary is low; (ii) AMR parsing accuracy estimation can be used to produce *high-quality automatically parsed data*: by filtering the outputs of single parsing systems in self-training, by selecting high-quality outputs from different parsing systems in a tri-parsing setting, or else by predicting overall rankings over alternative parsing systems applied to in- or out-of-domain data; (iii) finally, AMR parse accuracy prediction could be used in the context of a *parser-supported treebank construction* process. E.g., in an active learning scenario, we can select useful targets for manual annotation based on their expected efficiency for parser improvement – the fine-grained evaluation measures predicted by our system can be used for targeted improvements. In the simplest case, we can provide the human annotator with automatic parses where only few flaws have to be mended. Hence, AMR accuracy prediction systems have the potential to tremendously reduce manual annotation cost and time.

Contributions We define AMR accuracy prediction as the task of predicting a rich suite of metrics to assess various subtasks covered by AMR parsing (e.g. negation detection or semantic role labeling). To approach this task, we use the AMR evaluation suite suggested by Damonte et al. (2017) and develop a hierarchical multi-output regression model for automatically performing evaluation of 12 different tasks involved in AMR parsing (Sections §3 and §4; our code is publicly accessible¹). We perform experiments in three different scenarios on unseen in-domain and out-of-domain data and show that our model (i) is able to predict scores with significant correlation to gold scores and (ii) can be used to rank *parses* on a *sentence-level* or to rank *parsers* on a *corpus-level* (§5).

2 Related Work

Automatic accuracy prediction for syntactic parsing comes closest to what we are doing. Ravi et al. (2008) propose a feature-based SVM regression model with RBF kernel that predicts syntactic parser performance on different domains. Like us, they aim at a cheap and effective means for estimating a parser’s performance. However, in contrast to their work, our method is *domain*

and *parser* agnostic: we do not take into account characteristics of the domains of interest and do not provide any performance statistics of the competing parsing systems as features to our regressor. Biici (2016) addresses the task without any domain-dependent features, which results in a lower correlation to gold scores – even if additional features from a background language model are incorporated. In contrast to the prior systems that predict a single score, we predict an ensemble of metrics suitable for assessing AMR parse quality with respect to different linguistic aspects. Also, our system does not rely on externally derived features or complex pre-processing. Moreover, an AMR graph differs in important ways from a syntactic tree. Nodes in AMR do not explicitly correspond to words (as in dependency trees) or phrases (as in constituency trees). AMR structure elements can exist without any alignment to words in the sentence. To our knowledge, we are the first to propose an accuracy prediction model for AMR parsing, and offer the first general end-to-end parse accuracy prediction model that predicts an ensemble of scores for different linguistic aspects.

Automatic accuracy prediction has also been researched for PoS-tagging (Van Asch and Daelemans, 2010) and in machine translation. For example, Soricut and Narsale (2012) predict BLEU scores for machine-produced translations. Under the umbrella of *quality estimation* researchers try to predict, i.a., the post-editing time or missing words in an automatic translation (Cai and Knight, 2013; Joshi et al., 2016; Chatterjee et al., 2018; Kim et al., 2017; Specia et al., 2013). The fact that manually creating an AMR graph is significantly more costly than a translation provides another compelling argument for investigating automatic AMR accuracy prediction techniques.²

In recent work, Dickinson and Smith (2011, 2017); Jain et al. (2015); Rehbein and Ruppenhofer (2018) detect annotation errors in automatically produced dependency parses. The latter approach uses active learning and ensemble parsing in combination with variational inference. They predict edge labelling and attachment errors and use a back-and-forth encoding mechanism from non-structured to structured tree data in order to provide the variational inference model with the

¹<https://gitlab.cl.uni-heidelberg.de/opitz/quamr>

²Creating an AMR graph requires trained linguists and takes on average 8 to 13 minutes, cf. Banarescu et al. (2013)

```

(a / asbestos
 :time (n / now)
 :polarity -
 :location (p / product
 :poss (w / we)))

(a / asbestos
 :polarity -
 :location (p / product)
 :time (n / now))

(a / asbesto
 :polarity -
 :ARG1 (w / we
 :ARG1-of (p / product
 :mod (n / now))))

```

metr. (F1)	GP	JA	CA
Smatch	70	30	67
SRL	0	14	0
Concepts	67	44	50
IgnVars	55	0	60

Figure 2: Three AMR parses for: *There is no asbestos in our products now*, generated by GPLA (top), JAMR (bottom), CAMR (right). Light and severe errors are found in GPLA and JAMR parses; CAMR fails to provide *we*, the manufacturer of the product. Bottom right: F1 for Smatch and three example subtasks from evaluation against the gold parse (given in Figure 1).

needed unstructured data. Their work differs from ours in three important aspects: firstly, they predict errors in specific edges or nodes, while we predict an accuracy score over the complete graph. Moreover, our model does not *need* several candidate parses as input – when several multiple parses are available, our model can be exploited for ranking (cf. Sections §5.2 & §5.3). Finally, our method is independent of live human feedback.

3 Accuracy Metrics for AMR Parsing

Automatic AMR parses are often deficient. Consider the examples in Figure 2. All parsers correctly detect the negation and its scope. The GPLA parse (top) provides a graph structure close to the gold annotation (Figure 1). However, it does not correctly analyze the possessive *our* (*product*), which in the gold parse is represented as an object produced by the speaker (*we*). Instead it recognizes a location in the speaker’s possession. JAMR (middle) fails to detect the concept in focus (*asbestos*), possibly due to a false-positive stemming mistake. Moreover, it fails to represent that *asbestos* is (not) in the product: it misses the *:location*-edge from *asbestos* to *product*.

AMR accuracy metrics Usually, a predicted AMR graph G is evaluated against a gold graph G' using triple matching based on a maximally scoring variable mapping. For finding the optimal variable mapping, Integer Linear Programming (ILP) can be used in the Smatch metric (Cai and Knight, 2013), which produces precision, recall and F1 score between G and G' . While it is important to obtain a global measure of parse accuracy, we may also be interested in a quality assessment

that focuses on specific subtasks or meaning aspects, such as entity linking, negation detection or word sense disambiguation (WSD). For example, if a parser commits a WSD error this might be less harmful than e.g., failing to capture negation, or missing or wrongly predicting a semantic role. However, the Smatch calculation would treat many of such errors with equal weight – a property which in some cases may be undesirable.

To alleviate this issue, Damonte et al. (2017) proposed an extended AMR evaluation suite which allows parser performance inspection with regard to 11 additional subtasks captured by AMR. In total, 36 metrics can be computed (precision, recall and F1 for 12 tasks). F1 scores for three example metrics are displayed in Figure 2 (bottom, right): *Smatch*, *SRL* (Smatch computed on arg-i roles), *IgnoreVars* (triple overlap after replacing variables with concepts) and *Concepts* (F1 for concept identification).³ GPLA produces the overall best parse but it is outperformed by the other systems in *SRL* (JAMR) and *IgnoreVars* (CAMR).

Task definition We adopt the proposed metrics by Damonte et al. (2017) and use them as target metrics for our task of AMR parse accuracy prediction. Given an automatic AMR graph G and a corresponding sentence S , we estimate precision, recall and F1 of the main task (Smatch) and of the subtasks, as they would emerge from comparing G with its gold counterpart G' .

One of our hypotheses is that predicting a wide range of accuracy metric scores for individual aspects of AMR structures will aid our model to better predict the global Smatch scores. We will therefore investigate a hierarchical model that builds on predicted subtask measures in order to predict the global smatch score. Being able to predict fine-grained quality aspects of AMR parses will also be useful to assess and exploit differences of alternative system outputs and provides a basis for guiding system development or targeted annotation in an active learning setting.

4 Neural Accuracy Prediction Model

We propose a neural hierarchical multi-output regression model for accuracy prediction of AMR

³The other subtasks are: *Unlabelled* (Smatch after edge label removal), *No WSD* (Smatch after PropBank sense removal), *NS frames* (PropBank frame identification without sense), *Wikification* (entity linking), *NER* (named entity recognition), *Reentrancy* (Smatch over re-entrant edges).

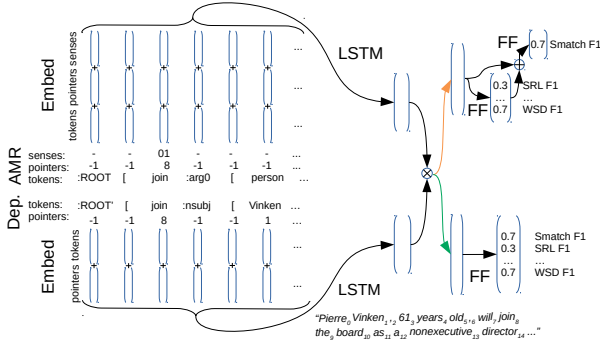


Figure 3: Our model: green: Evaluation metrics computed in a non-hierarchical fashion. orange: Main evaluation metric is computed on top of secondary metrics.

parses. Its architecture is outlined in Figure 3.

Inputs Our model takes the following inputs: (i) a linearized AMR and a linearized dependency graph (implementation details in §5). The motivation for feeding the dependency parse instead of the original sentence is due to the moderate similarity of dependency and AMR structures.⁴ We examine drawbacks and benefits of providing automatic dependency parses more closely in our ablation experiments (§5.4). In addition, (ii) we produce alignments between sentence tokens and tokens in the sequential AMR structure, as well as between sentence tokens and the linearized dependency structure, and feed these sequences of pointers to our accuracy prediction model. The intuition of using pointers is to provide the model with richer information via shallow alignment between AMR, dependencies and the sequence of sentence tokens (see Section §5 for implementation details). Finally, (iii) we feed a sequence of PropBank sense indicators for AMR predicates.

Joint encoding of AMR and dependency parses for metric prediction Embedding layers are shared between AMR/dependency pointers and AMR/dependency tokens. We embed the three sequences representing the AMR graph (tokens, pointers and senses) in three matrices and sum them up element-wise (indicated with + in Figure 3). The same procedure is applied to the linearized dependency graph (tokens and pointers). The resulting matrices are processed by two two-layered Bi-LSTMs to yield vectorized representations for (i) the AMR graph and (ii) the dependency tree (i.e., the last states of forward and backward reads are concatenated). Thereafter, we apply element-

⁴c.f. Groschwitz et al. (2018); Chen and Palmer (2017).

wise multiplication, subtraction and addition to both vector representations and concatenate the resulting vectors (\otimes in Figure 3). The joint AMR-dependency representation is further processed by a feed forward layer (FF) with sigmoid activation functions in order to predict, in total, 36 different metrics (green, Figure 3).

Hierarchical prediction of multiple metrics

The task naturally lends itself to be formulated in a hierarchical multi-task setup (orange, Figure 3). In this strand, we first compute the 33 fine-grained subtask metrics and on their basis we calculate the Smatch scores (precision, recall, F1) as our primary metrics. In order to accomplish this, we collect the outputs from the subtask metric prediction layer in a vector and concatenate it with the previous layer’s representation (\oplus in Figure 3). The resulting vector is fed through a last FF layer to predict the metrics for the task of main interest (Smatch). Our intuition is that the estimated quality of the parse with respect to the subtask metrics informs the model and allows it to better predict the overall quality.

Loss In the non-hierarchical case, we denote our full model with $f_\theta : X \rightarrow [0, 1]^d$ with parameters θ , where d describes the dimensionality of the score vector (one dimension represents one metric) and $D = \{(X_i, y_i)\}_{i=1}^N, y_i \in [0, 1]^d$ is our training data. In the non-hierarchical model, we minimize the mean squared error:

$$\ell(f_\theta) = \frac{1}{dN} \sum_{i=1}^N \sum_{j=1}^d (y_{i,j} - f_\theta(X_i)_j)^2 \quad (1)$$

For our hierarchical model, we have two functions, $f_\theta : X \rightarrow [0, 1]^{(d-k)}$ which returns the output vector for the $(d-k)$ subtask metrics and $f'_{\theta'} : X \rightarrow [0, 1]^k$ which returns the output vector for our k main metrics (in our experiments, $k = 3$ for Smatch recall, precision and F1). Then,

$$\begin{aligned} \ell'(f_\theta, f'_{\theta'}) &= \frac{\lambda_1}{(d-k)N} \sum_{i=1}^N \sum_{j=1}^{d-k} (y_{i,j} - f_\theta(X_i)_j)^2 \\ &+ \frac{\lambda_2}{kN} \sum_{i=1}^N \sum_{j=d-k+1}^d (y_{i,j} - f'_{\theta'}(X_i)_{j-(d-k)})^2 \end{aligned}$$

defines the total loss over the two entangled metric prediction models. Note that $\theta \subset \theta'$, which means that by optimizing the parameters of f' with gradient descent, we also concurrently optimize all

parser	training		development	
	Smatch (F1)	% def.	Smatch (F1)	% def.
JAMR	0.79	86.7	0.69	91.8
CAMR	0.75	93.6	0.66	95.7
GPLA	0.86	83.4	0.76	90.0

Table 1: Parser output evaluation on training and development partitions of LDC2015E86. Smatch F1: avg. over Smatch F1 per sentence, % def.: percentage of deficient parses (i.e., parses with Smatch F1 < 1).

parameters of f . By this construction, the hierarchical model instantiates a two-task model with shared parameters. For our experiments, we manually set the loss weights $\lambda_1 = 0.2$, $\lambda_2 = 1$.

5 Experiments

Data Since our goal is to predict the accuracy of an automatic parse, we need a data set containing automatically produced AMR parses and their scores, as they would emerge from comparison to gold parses. Our largest data set, LDC2015E86, comprises 19,572 sentences and comes in a pre-defined training, development and test split. We parse this data set with three parsers, JAMR (Flanigan et al., 2014, 2016), CAMR (Wang et al., 2015b,a, 2016a) and GPLA (Lyu and Titov, 2018). Since the three parsers have been trained on the training data partition, we naturally obtain more accurate parses for the training partition than for development and test data. Table 1, however, indicates that we still obtain a considerable amount of deficient parses for training. Based on the parser outputs we compute evaluations comparing the automatic parses with the gold parses by using *amr-evaluation-tool-enhanced*⁵, a bug-fixed version of the script that computes the metrics of Damonte et al. (2017). This allows us to create full-fledged training, development and test instances for our accuracy prediction task. Each instance consists of a sentence and an AMR parse as input and a vector of metric scores as target.

Our second data set, LDC2015R36, comprises submissions to the SemEval-2016 Task 8 (May, 2016). We have 1053 parses from each of the 11 team submissions (and 2 baseline systems).⁶ Our

⁵<https://github.com/ChunchuanLv/amr-evaluation-tool-enhanced>

⁶Riga (Barzdins and Gosko, 2016), CMU (equal to JAMR) (Flanigan et al., 2016), Brandeis (Wang et al., 2016b), UofR (Peng and Gildea, 2016), ICL-HD (Brandt et al., 2016), M2L (Puzikov et al., 2016), UMD (Rao et al., 2016),

data set	LDC2015E86	LDC2015R36	BioAMRTTest
domain	news	news	medical
nb. sentences	19,572	1,053	500
avg. sent. len.	21	22.35	36.52
nb. auto. parses	58,716	13,689	3,000
used as	train/dev/test	test	test

Table 2: Statistics of data sets used in this work.

third dataset, BioAMRTTest is used as the test set in the SemEval-2017 Task 9 (May and Priyadarshi, 2017) and consists of 500 parses from each of the 6 teams.⁷ The shared task organizers kindly made this data available for our experiments.

Preprocessing For dependency annotation, we parse all sentences with spacyV2.0⁸. For sequentializing the AMR and dependency graph representations we take intuitions from van Noord and Bos (2017b) & Konstas et al. (2017) and output tokens by performing a depth-first-search over the graph. We replace the AMR negation token ‘-’ and strings representing numbers with special tokens. The vocabularies (tokens, senses and pointers) are computed from our training partition of LDC2015E86 and comprise all tokens with a frequency ≥ 5 (tokens with lesser frequency are replaced by an *OOV*-token). PropBank senses of predicates are removed and collected in an extra list that is parallel to the tokens in the linearized AMR sequence. For each linearized AMR and dependency tree we generate a sequence with index pointers to tokens in the original sentence (-1 for tokens which do not explicitly refer to any token in the sentence, e.g. brackets, ‘subj’ or ‘arg0’ relations). Extraction of token-pointers from the dependency graph is trivial. For every concept in the linearized AMR we execute a search for the corresponding token in the sentence, looking for exact matches with surface tokens and lemmas.

Training For the optimization of the accuracy prediction model we use only the development and training sections of LDC2015E86 and the corresponding automatic parses together with the gold scores. Details on the training cycle can be found in the *Supplemental Material* §A (the loss is de-

DynamicPower (Butler, 2016), TMF (Bjerva et al., 2016), UCL+Sheffield (Goodman et al., 2016) and CU-NLP (Foland and Martin, 2016).

⁷TMF-1 and TMF-2 (van Noord and Bos, 2017a), DAN-GNT (Nguyen and Nguyen, 2017), Oxford (Buys and Blunsom, 2017), RIGOTRIO (Gruzitis et al., 2017) and JAMR (Flanigan et al., 2016)

⁸<https://spacy.io/>

	ρ LDC2015E86			ρ BioAMRTest		
	P	R	F1	P	R	F1
Smatch	0.74	0.79	0.78	0.54	0.41	0.47
Concepts	0.56	0.65	0.64	0.67	0.55	0.62
Frames	0.7	0.71	0.72	0.67	0.56	0.63
IgnoreVars	0.76	0.8	0.79	0.33	0.27	0.29
Named Ent.	0.81	0.81	0.81	0.5	0.48	0.5
Negations	0.87	0.87	0.87	0.33	0.32	0.32
No WSD	0.75	0.78	0.78	0.54	0.41	0.46
NS frames	0.76	0.75	0.77	0.72	0.59	0.67
Reentrancies	0.77	0.79	0.8	0.52	0.45	0.48
SRL	0.72	0.74	0.75	0.47	0.43	0.45
Unlabeled	0.71	0.75	0.75	0.6	0.45	0.51
Wikification	0.87	0.85	0.86	0.24	0.23	0.23

Table 3: Pearson correlation coefficient (ρ) over various metrics and across domains. Explanations of the metrics and AMR subtasks are in Section §3 and fn. 3




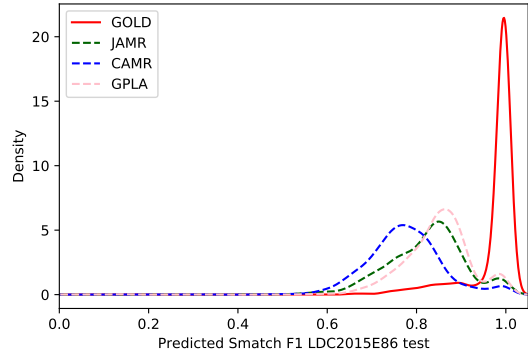
Figure 4: Predicted (y-axis) & gold (x-axis) Smatch F1.

scribed in §4). We use the same single (hierarchical) model for all three evaluation studies, proving its applicability across different scenarios (a non-hierarchical model is only instantiated for the ablation experiments in Section §5.4).

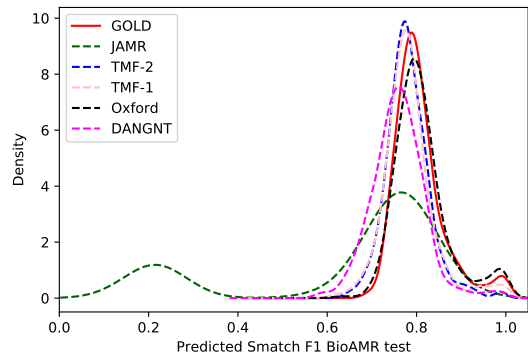
5.1 Correlation with Gold Accuracy

The primary goal in our first experiment is to test whether the system is able to differentiate good from bad parses. This capacity is expressed by a high correlation of predicted accuracies with true accuracies on unseen data and by the ability to assign high scores to gold parses. We evaluate on the test partition of LDC2015E86 and BioAMRTest.

Correlation results The results are displayed in Table 3. Over all metrics, in-domain and out-of-domain, we achieve significant correlations with the gold scores ($p < 0.005$ for every metric). While on LDC2015E86 the model has learned to predict the KB linking F1 ($\rho = 0.86$) and negation detection F1 with high correlation to the gold



(a) LDC2015E86



(b) BioAMRTest

Figure 5: Probability density function estimations for predicted F1 Smatch scores using Scott’s method (Scott, 2012) with respect to candidate parses from different systems.

scores ($\rho = 0.87$), Concept assessment poses the greatest challenge ($\rho = 0.64$). For the out-of-domain data BioAMRTest, these two facts seem almost reversed: here, the assessment of KB linking poses difficulties ($\rho = 0.23$) while the Concept F1 predictions are better ($\rho = 0.62$). The main metrics of interest (Smatch precision, recall and F1) can be predicted with high correlation on in-domain data ($\rho \geq 0.74$, cf. also Figure 4) and solid correlation for out-of-domain data ($\rho \geq 0.41$).

Find the Gold AMR! Now, we want to test our system’s capacity to reliably predict high Smatch F1 scores for unseen gold AMR parses. Ideally, the scores should be close or equal to 1. For in-domain data, it appears to work well: a large amount of Smatch predictions for gold AMR graphs are very close to one (Figure 5a).

Evidently, our system also gets the ranking of the parsing systems right: the distribution of the state-of-the-art (GPLA) is shifted right towards

dataset	percentile						
	5	25	75	90	95	97	99
LDC15E86	0.83	0.99	1.0	1.0	1.0	1.0	1.0
BioAMRTest	0.74	0.77	0.83	0.88	0.93	0.98	1.0

Table 4: Various percentiles of Smatch F1 predictions for gold graphs.

higher predicted F1 scores, whereas the distribution of CAMR is shifted left towards lower scores. Also, more than 75% of gold parses have a predicted Smatch score of more than 0.99 (Table 4).

On the other hand, finding gold parses in the BioAMRtest data is much harder: about 75% of Smatch scores get assigned a score of 0.83 or lower and only 1% of gold parses are predicted as perfect (Table 4). The estimated probability density function for gold parses (red solid line in Figure 5b) struggles to discriminate itself from the functions corresponding to the flawed parses of the automatic systems. Nevertheless, the prediction score density for gold parses is situated more on the right hand side than most others. In other words, we find that in the out-of-domain data gold parses tend to be assigned above-average scores.

To sum up, our observations for the out-of-domain data stand in some contrast to what we observe for the in-domain data. However, this outcome can be plausibly explained: assuming that the out-of-domain gold parses have some unfamiliar properties, a system that has never seen such parses cannot judge well whether they are gold or not. In fact, it can be interpreted positively that the system hesitates to assign maximum scores to gold parses from a domain in which the model is completely inexperienced. Additionally, bio-medical texts involve difficult concepts, naming conventions and complicated noun phrases which are hard to understand even for non-expert humans (e.g., “*TAK733 led to a decrease in pERK and G1 arrest in most of these melanoma cell lines regardless of their origin, driver oncogenic mutations and in vitro sensitivity to TAK733*”). Taking all this into account, the results for out-of-domain data may be not as bad as they perhaps appear at first glance.

5.2 Application Study: AMR Parse Ranking

Our automatic accuracy prediction method naturally lends itself for ranking parser outputs. For any sentence, provided automatic parses by competing systems can be ranked according to the

	Smatch LDC2015E86			Smatch BioAMRTest		
	P	R	F1	P	R	F1
lower-bound	64.9	57.9	60.5	41.7	31.3	34.3
random	72.4	67.0	69.1	60.3	50.3	54.0
ours	76.6	73.5	74.8	<u>64.9</u>	<u>56.0</u>	<u>59.2</u>
upper-bound	79.3	75.2	76.9	73.2	65.2	68.5
JAMR	71.4	66.5	68.4	48.4	39.7	42.9
CAMR	69.5	60.4	64.0	-	-	-
GPLA	76.3	73.4	74.6	-	-	-
TMF-1	-	-	-	56.0	46.5	49.3
TMF-2	-	-	-	70.0	54.5	60.5
DANGNT	-	-	-	70.2	58.6	63.1
Oxford	-	-	-	65.8	59.0	61.6
RIGOTRIO	-	-	-	65.0	50.8	56.4

Table 5: Results (sentence averages) of different AMR parsing (bottom part) and ranking (top part) systems on two test sets. Upper part: results when selecting from alternative parses: lower-bound (upper-bound): oracle selecting the worst (best) AMR parse; ours: results when selecting the best parse according to our models’ accuracy prediction (hierarchical model).

scores predicted by our system. This scenario arises, e.g., when we run several AMR parsers over a large corpus with the aim of selecting the best parse for each sentence in order to collect silver training data.⁹ In the worst case, we do not have any prior knowledge about a parser’s performance (we may not even know the source of a parse). We use the test partition from LDC2015E86 and BioAMRTest to rank, for each sentence, the automatic candidate parses provided by the different parsers. In LDC2015E86 we assume not to be agnostic about the parsers as their performances on the development data of this data set are known (in terms of their sentence-average F1 Smatch score). Consider that we are given a sentence and three automatic parses. We select the maximum-score parse, where the score is defined by predicted Smatch F1 plus the average Smatch F1 of the parse-producing parser on the development data. As baselines in this scenario we (i) randomly choose a parse from the three options or (ii) always choose the parse of GPLA. On BioAMRTest, however, we have no prior information about the submitted systems. We select from 6 automatic parses for each sentence. Since now we are completely parser agnostic, the baseline is to randomly select a parse from the candidate set.

Results The results are displayed in Table 5. For our in-domain test data, LDC2015E86, selecting

⁹In a self-training scenario, we also could set a threshold of minimum predicted accuracy to select confident parses.

	Smatch LDC2015E86		Smatch BioAMRTest	
	$\bar{\rho}$	%pos	$\bar{\rho}$	%pos
lower-bound	-1	0.0	-1	0.0
random	0.00	50.0	0.00	50.0
ours	0.54	77.0	0.22	70.4
upper-bound	1.00	100.0	1.00	100.0

Table 6: Results of different parse-ranking systems with respect to sentence-level parse rankings. $\bar{\rho}$: average Pearson-r on a sentence level. %pos: ratio of predicted rankings with positive ρ to gold ranking.

the best parse according to our model’s predicted accuracy score improves over *all* individual parser results: the obtained average Smatch F1 per sentence increases (i) slightly by 0.2 pp. compared to always choosing outputs from GPLA and (ii) observably by 5.7 pp. compared to randomly selecting a parse from the competing system outputs. The difference compared to always choosing GPLA seems negligible which perhaps can be explained by the fact that GPLA has been shown to be on par or better than doubly-blind human annotators.¹⁰ The oracle that always selects the best parse (*upper-bound* in Table 5) shows little room for improvement: it achieves 2.1 pp. Smatch F1 increase compared to our model. This margin is small and further success might also be hampered by peculiarities in the manual annotations. On BioAMRTest, no prior information about the systems is available. Using our model’s predicted scores to select from the alternative system outputs, we can boost Smatch F1 by 5.2 pp. compared to randomly selecting a parse. Compared to always selecting the parses of the best submitted system (in-hindsight), we lag behind by 3.9 pp.

Since our data comprises outputs from several parsers with varying performance, we can study the performance of our approach in combination with different parsers (Figure 6). When only choosing among CAMR and JAMR outputs, on LDC2015E86, our system boosts the F1 by 2.7 pp. compared to randomly selecting a parse, and by 0.6 pp. compared to always choosing the parse from the better system (determined on dev, here: JAMR). Choosing from CAMR and GPLA or JAMR and GPLA makes little difference: in most cases our system selects the GPLA parse and the difference to only choosing GPLA parses is

¹⁰GPLA (Lyu and Titov, 2018) achieves a high 74.4% corpus-level Smatch F1 (primarily news texts), while a prior annotation study (Banarescu et al., 2013) reported doubly blind annotation corpus-level F1 of 0.71 (for web texts).




Figure 6: Using our model to predict the best parse out of two candidate parses, each from a different system.

marginal. Moreover, across both test sets, the majority of rankings assigned by our method have positive correlations with the true rankings (Table 6): 77% of all assigned rankings have a positive correlation with the true ranking (70% for biomedical). In sum, we can draw two conclusions from this experiment: given a sentence, ranking AMR parser outputs using our accuracy prediction model, on in-domain and out-of-domain unseen data (i) clearly improves performance when non state-of-the-art parsers are applied or if we are not informed about the parsers’ performances and (ii) does not worsen results in other cases.

5.3 Application Study: Predict System Ranks

In our final case study, we use our accuracy prediction model to predict a ranking over systems. We use our model to rank the unseen submitted system parses of the SemEval-2017 Task 9 (evaluated on BioAMRTest) and SemEval-2016 Task 8 (evaluated on LDC2015R36) according to average predicted F1 Smatch scores. Again, we assume a parser-agnostic setting, meaning we have no prior knowledge of the submitted systems (i.e. we just consider their outputs). In this setting, we do not rank individual parses given a sentence, but rank the system outputs, according to estimated average Smatch F1 per sentence. We evaluate against the final team rankings of the two shared tasks.

Results The results are displayed in Table 7. On BioAMRTest we have a good, albeit non statistically significant correlation with the true team ranking. On the in-domain LDC2015R36 test set we see a significant correlation of $\rho = 0.645$ ($p_{1,2} < 0.05$). In this shared task, many teams were competitive and differences between the best teams were marginal. For example, in the true ranking, places 1 to 6 achieved between 0.60 and 0.62 Smatch F1. Notably, the first four teams ac-

	Rank LDC2015R36		Rank BioAMRTTest	
	rank r	rank \hat{r}	rank r	rank \hat{r}
DANGNT	-	-	1	3
Oxford	-	-	2	1
TMF-2	-	-	3	2
RIGOTRIO	-	-	4	5
TMF-1	-	-	5	4
JAMR	7	7	6	6
RIGA	1	4	-	-
Brandeis	2	3	-	-
CU-NLP	3	1	-	-
UCL+Sheffield	4	2	-	-
ICL-HD	5	8	-	-
M2L	6	10	-	-
JAMR-base	8	12	-	-
UofR	9	11	-	-
TMF	10	5	-	-
UMD	11	6	-	-
DynamicPower	12	13	-	-
det. baseline	13	9	-	-
ρ	0.645 ($p_1 = 0.017, p_2 = 0.011$)		0.771 ($p_1 = 0.072, p_2 = 0.051$)	

Table 7: True rank r (given corpus-Smatch) and predicted rank \hat{r} (based on sentence average Smatch computed using our model). p_1 : probability of non-correlation. p_2 : probability that a randomly produced ranking achieves equal or greater ρ (estimated over 10^6 random rankings). For team names, see fn. 6 & 7.

According to the true ranking and the first four teams according to our predicted ranking fall into the same group. This shows that our model successfully assigned high ranks to low error submissions.

5.4 Ablation Experiments

We finally perform ablation experiments to evaluate the impact of individual model components. We experiment with five different setups. (i) instead of stacking two Bi-LSTMs, we use only one Bi-LSTM (*one-lstm*, Table 8). (ii) instead of the dependency tree, we feed the words in the order as they occur in the sentence (*no-dep*). (iii) *no-pointers*: we remove the token-pointers from our model. (iv), instead of using the hierarchical setup, we predict all metrics on the same level (green in Figure 3, *no-HL* in Table 8) and (v), *no-HMTL*: we optimize the non-hierarchical model only with respect to Smatch, disregarding the AMR subtasks. Remarkably, the dependency tree greatly helps the model on in-domain data over all measures (-37 total Δ without dependencies) but hurts the model on out-of-domain data (+27 total Δ). A possible explanation is the degradation of the dependency parse quality: bio-medical data not only poses a challenge for our model, but also for the dependency parser. With special regard to the main AMR evaluation measure, Smatch F1, the learned pointer embeddings provide useful input on the in-domain test data (-4 Δ without pointers).

	LDC2015R36						BioAMRTTest					
	<i>complete</i>	<i>one-lstm</i>	<i>no-dep</i>	<i>no-pointers</i>	<i>no-HL</i>	<i>no-HMTL</i>	<i>complete</i>	<i>one-lstm</i>	<i>no-dep</i>	<i>no-pointers</i>	<i>no-HL</i>	<i>no-HMTL</i>
	ρ	Δ					ρ	Δ				
Smatch	78	-1	-1	-4	-3	-2	47	0	+5	+4	+2	-3
Concepts	64	-1	-4	-3	-4	-	62	0	+3	+2	0	-
Frames	72	0	-5	0	-1	-	63	0	+1	+1	-1	-
IgnoreVars	79	-1	-1	-1	-2	-	29	+5	+6	+5	+4	-
Named Ent.	81	+2	-3	+2	+3	-	50	-18	-9	-7	-9	-
Negations	87	-1	-1	+1	0	-	32	-16	+2	-1	-4	-
No WSD	78	-1	0	-1	-2	-	46	+2	+6	+5	+3	-
NS frames	77	0	-7	0	0	-	67	+1	+1	0	-1	-
Reentrancies	80	0	-9	+1	+2	-	48	0	+2	0	+1	-
SRL	75	-1	-4	0	+1	-	45	-4	+3	0	+2	-
Unlabeled	75	-1	0	-1	-1	-	51	0	+1	+4	+2	-
Wikification	86	0	-2	+1	+2	-	23	+5	+6	+6	+7	-
$\sum_i \Delta_i$	0	-5	-37	-5	-5	-2	0	-25	+27	+19	+6	-3

Table 8: ρ correlation (F1) differences over different setups (columns), test sets (out-of-domain, in-domain) and subtasks (rows). $\pm x$: plus and minus x pp. ρ .

6 Conclusion

AMR parser evaluation with human gold annotation is very costly. Our main contributions in this work are two-fold: Firstly, we introduced the concept of automatic AMR accuracy prediction. Given only an automatic parse and the sentence, from whence it was derived, the goal is to predict evaluation metrics cheaply and possibly at runtime. Secondly, we framed the task as a multiple-output regression task and developed a hierarchical neural model to predict a rich suite of AMR evaluation metrics. We presented three case studies proving (i) the feasibility of automatic AMR accuracy prediction in general (significant correlation with gold scores on unseen in-domain and out-of-domain data) and (ii) the applicability of our model in two use cases. In the first study, we ranked different automatic candidate parses per sentence, outperforming the random selection baseline by 5.7 pp. average Smatch F1 (in-domain) and 5.2 pp. (out-of-domain). In the second study, we ranked team submissions to two AMR shared tasks and our method was able to reproduce rankings similar to the true rankings.

Acknowledgments

This work has been supported by the German Research Foundation (grant no. GRK 1994/1) and the Leibniz Association (grant no. SAS-2015-IDS-LWC) and the Ministry of Science, Research, and Art of Baden-Württemberg. We are grateful to the NVIDIA corporation for donating the GPU used in this research.

References

- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. Abstract meaning representation for sembanking.
- Guntis Barzdins and Didzis Gosko. 2016. [Riga at semeval-2016 task 8: Impact of smatch extensions and character-level neural translation on amr parsing accuracy](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1143–1147. Association for Computational Linguistics.
- Ergun Biici. 2016. [Predicting the performance of parsing with referential translation machines](#). *The Prague Bulletin of Mathematical Linguistics*, 106(1):31 – 44.
- Johannes Bjerva, Johan Bos, and Hessel Haagsma. 2016. [The meaning factory at semeval-2016 task 8: Producing amrs with boxer](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1179–1184. Association for Computational Linguistics.
- Lauritz Brandt, David Grimm, Mengfei Zhou, and Yannick Versley. 2016. [Icl-hd at semeval-2016 task 8: Meaning representation parsing - augmenting amr parsing with a preposition semantic role labeling neural network](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1160–1166. Association for Computational Linguistics.
- Alastair Butler. 2016. [Dynamicpower at semeval-2016 task 8: Processing syntactic parse trees with a dynamic semantics core](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1148–1153. Association for Computational Linguistics.
- Jan Buys and Phil Blunsom. 2017. [Oxford at semeval-2017 task 9: Neural amr parsing with pointer-augmented attention](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 914–919. Association for Computational Linguistics.
- Shu Cai and Kevin Knight. 2013. Smatch: an evaluation metric for semantic feature structures. In *ACL (2)*, pages 748–752. The Association for Computer Linguistics.
- Rajen Chatterjee, Matteo Negri, Marco Turchi, Frédéric Blain, and Lucia Specia. 2018. Combining quality estimation and automatic post-editing to enhance machine translation output. In *AMTA (1)*, pages 26–38. Association for Machine Translation in the Americas.
- Wei-Te Chen and Martha Palmer. 2017. [Unsupervised amr-dependency parse alignment](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 558–567. Association for Computational Linguistics.
- Marco Damonte, Shay B. Cohen, and Giorgio Satta. 2017. [An incremental parser for abstract meaning representation](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 536–546. Association for Computational Linguistics.
- Markus Dickinson and Amber Smith. 2011. [Detecting dependency parse errors with minimal resources](#). In *Proceedings of the 12th International Conference on Parsing Technologies*, pages 241–252. Association for Computational Linguistics.
- Markus Dickinson and Amber Smith. 2017. [Simulating dependencies to improve parse error detection](#). In *Proceedings of the 15th International Workshop on Treebanks and Linguistic Theories (TLT15), Bloomington, IN, USA, January 20-21, 2017.*, volume 1779 of *CEUR Workshop Proceedings*, pages 76–88. CEUR-WS.org.
- Shibhansh Dohare and Harish Karnick. 2017. Text summarization using abstract meaning representation. *CoRR*, abs/1706.01678.
- Jeffrey Flanigan, Chris Dyer, Noah A. Smith, and Jaime G. Carbonell. 2016. CMU at semeval-2016 task 8: Graph-based AMR parsing with infinite ramp loss. In *SemEval@NAACL-HLT*, pages 1202–1206. The Association for Computer Linguistics.
- Jeffrey Flanigan, Sam Thomson, Jaime Carbonell, Chris Dyer, and Noah A. Smith. 2014. [A discriminative graph-based parser for the abstract meaning representation](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1426–1436. Association for Computational Linguistics.
- William Foland and James H. Martin. 2016. [Cu-nlp at semeval-2016 task 8: Amr parsing using lstm-based recurrent neural networks](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1197–1201. Association for Computational Linguistics.
- Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *AISTATS*, volume 9 of *JMLR Proceedings*, pages 249–256. JMLR.org.
- James Goodman, Andreas Vlachos, and Jason Naradowsky. 2016. [Ucl+sheffield at semeval-2016 task 8: Imitation learning for amr parsing with an alpha-bound](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1167–1172. Association for Computational Linguistics.

- Jonas Groschwitz, Matthias Lindemann, Meaghan Fowlie, Mark Johnson, and Alexander Koller. 2018. [AMR dependency parsing with a typed semantic algebra](#). *CoRR*, abs/1805.11465.
- Normunds Gruzitis, Didzis Gosko, and Guntis Barzdins. 2017. [Rigotrio at semeval-2017 task 9: Combining machine learning and grammar engineering for amr parsing and generation](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 924–928. Association for Computational Linguistics.
- Sambhav Jain, Naman Jain, Bhasha Agrawal, and Rajeve Sangal. 2015. Employing oracle confusion for parse quality estimation. In *Computational Linguistics and Intelligent Text Processing*, pages 213–226. Cham. Springer International Publishing.
- Bevan K. Jones, Jacob Andreas, Daniel Bauer, Karl Moritz Hermann, and Kevin Knight. 2012. Semantics-based machine translation with hyperedge replacement grammars. In *COLING*, pages 1359–1376. Indian Institute of Technology Bombay.
- Nisheeth Joshi, Iti Mathur, Hemant Darbari, and Ajai Kumar. 2016. [Quality estimation of english-hindi machine translation systems](#). In *Proceedings of the Second International Conference on Information and Communication Technology for Competitive Strategies, ICTCS '16*, pages 53:1–53:5, New York, NY, USA. ACM.
- Hyun Kim, Hun-Young Jung, Hongseok Kwon, Jong-Hyeok Lee, and Seung-Hoon Na. 2017. [Predictor-estimator: Neural quality estimation based on target word prediction for machine translation](#). *ACM Trans. Asian Low-Resour. Lang. Inf. Process.*, 17(1):3:1–3:22.
- Diederik P. Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *CoRR*, abs/1412.6980.
- Ioannis Konstas, Srinivasan Iyer, Mark Yatskar, Yejin Choi, and Luke Zettlemoyer. 2017. [Neural amr: Sequence-to-sequence models for parsing and generation](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 146–157. Association for Computational Linguistics.
- Fei Liu, Jeffrey Flanigan, Sam Thomson, Norman Sadeh, and Noah A. Smith. 2015. [Toward abstractive summarization using semantic representations](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1077–1086. Association for Computational Linguistics.
- Chunchuan Lyu and Ivan Titov. 2018. [Amr parsing as graph prediction with latent alignment](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 397–407. Association for Computational Linguistics.
- Jonathan May. 2016. [Semeval-2016 task 8: Meaning representation parsing](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1063–1073, San Diego, California. Association for Computational Linguistics.
- Jonathan May and Jay Priyadarshi. 2017. [Semeval-2017 task 9: Abstract meaning representation parsing and generation](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, page 536–545, Vancouver, Canada. Association for Computational Linguistics, Association for Computational Linguistics.
- Arindam Mitra and Chitta Baral. 2016. [Addressing a question answering challenge by combining statistical methods with inductive rule learning and reasoning](#). In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.*, pages 2779–2785. AAAI Press.
- Khoa Nguyen and Dang Nguyen. 2017. [UIT-DANGNT-CLNLP at semeval-2017 task 9: Building scientific concept fixing patterns for improving CAMR](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation, SemEval@ACL 2017, Vancouver, Canada, August 3-4, 2017*, pages 909–913. Association for Computational Linguistics.
- Rik van Noord and Johan Bos. 2017a. [The meaning factory at semeval-2017 task 9: Producing amrs with neural semantic parsing](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 929–933. Association for Computational Linguistics.
- Rik van Noord and Johan Bos. 2017b. Neural semantic parsing by character-based translation: Experiments with abstract meaning representations. *Computational Linguistics in the Netherlands Journal*, 7:93–108.
- Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. [The proposition bank: An annotated corpus of semantic roles](#). *Comput. Linguist.*, 31(1):71–106.
- Xiaochang Peng and Daniel Gildea. 2016. [Uofr at semeval-2016 task 8: Learning synchronous hyperedge replacement grammar for amr parsing](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1185–1189.
- Yevgeniy Puzikov, Daisuke Kawahara, and Sadao Kurohashi. 2016. [M2L at semeval-2016 task 8: AMR parsing with neural networks](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation, SemEval@NAACL-HLT 2016, San Diego, CA, USA, June 16-17, 2016*, pages 1154–1159. The Association for Computer Linguistics.

- Sudha Rao, Yogarshi Vyas, Hal Daumé III, and Philip Resnik. 2016. [Clip\\$@Sumd at semeval-2016 task 8: Parser for abstract meaning representation using learning to search](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1190–1196. Association for Computational Linguistics.
- Sujith Ravi, Kevin Knight, and Radu Soricut. 2008. Automatic prediction of parser accuracy. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 887–896. Association for Computational Linguistics.
- Ines Rehbein and Josef Ruppenhofer. 2018. [Sprucing up the trees – error detection in treebanks](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 107–118. Association for Computational Linguistics.
- David W Scott. 2012. Multivariate density estimation and visualization. In *Handbook of computational statistics*, pages 549–569. Springer.
- Radu Soricut and Sushant Narsale. 2012. Combining quality prediction and system selection for improved automatic translation output. In *Proceedings of the Seventh Workshop on Statistical Machine Translation*, pages 163–170. Association for Computational Linguistics.
- Lucia Specia, Kashif Shah, Jose G. C. De Souza, Trevor Cohn, and Fondazione Bruno Kessler. 2013. Quest - a translation quality estimation framework. In *Proceedings of the 51th Conference of the Association for Computational Linguistics (ACL), Demo Session*.
- Vincent Van Asch and Walter Daelemans. 2010. [Using domain similarity for performance estimation](#). In *Proceedings of the 2010 Workshop on Domain Adaptation for Natural Language Processing, DANLP 2010*, pages 31–36, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Chuan Wang, Sameer Pradhan, Xiaoman Pan, Heng Ji, and Nianwen Xue. 2016a. [Camr at semeval-2016 task 8: An extended transition-based amr parser](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1173–1178, San Diego, California. Association for Computational Linguistics.
- Chuan Wang, Sameer Pradhan, Xiaoman Pan, Heng Ji, and Nianwen Xue. 2016b. [Camr at semeval-2016 task 8: An extended transition-based amr parser](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1173–1178.
- Chuan Wang, Nianwen Xue, and Sameer Pradhan. 2015a. [Boosting transition-based amr parsing with refined actions and auxiliary analyzers](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 857–862, Beijing, China. Association for Computational Linguistics.
- Chuan Wang, Nianwen Xue, and Sameer Pradhan. 2015b. [A transition-based algorithm for amr parsing](#). In *Proceedings of the 2015 Conference for the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 366–375, Denver, Colorado. Association for Computational Linguistics.

A Supplemental Material

Hyper parameters and weights initialization

We initialize all parameters of the model randomly. Embedding vectors of dimension 128 are drawn from $U(0.05, 0.05)$ and the LSTM weights (neurons: 128) and weights of the feed forward output layers are sampled from a Glorot uniform distribution (Glorot and Bengio, 2010). For future work, initializing the embedding layer with pre-trained vectors could further increase the performance. In this work, however, we learn all parameters from the given data. We fit our model using Adam (Kingma and Ba, 2014) (learning rate: 0.001) on the training data over 20 epochs with mini batches of size 16. We apply early stopping according to the maximum Pearson’s ρ (with regard to Smatch F1) on the development data. $\rho = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^n (x_i - \bar{x})^2} \sqrt{\sum_{i=1}^n (y_i - \bar{y})^2}}$ quantifies the linear relationship between predicted scores (x_1, \dots, x_n) and true scores (y_1, \dots, y_n) .