

AMR Quality Rating with a Lightweight CNN

Juri Opitz

Dept. of Computational Linguistics
Heidelberg University
69120 Heidelberg
opitz@cl.uni-heidelberg.de

Abstract

Structured semantic sentence representations such as Abstract Meaning Representations (AMRs) are potentially useful in various NLP tasks. However, the quality of automatic parses can vary greatly and jeopardizes their usefulness. This can be mitigated by models that can accurately rate AMR quality in the absence of costly gold data, allowing us to inform downstream systems about an incorporated parse’s trustworthiness or select among different candidate parses.

In this work, we propose to transfer the AMR graph to the domain of images. This allows us to create a simple convolutional neural network (CNN) that imitates a human judge tasked with rating graph quality. Our experiments show that the method can rate quality more accurately than strong baselines, in several quality dimensions. Moreover, the method proves to be efficient and reduces the incurred energy consumption.

1 Introduction

The goal of sentence meaning representations is to capture the meaning of sentences in a well-defined format. One of the most prominent frameworks for achieving this is *Abstract Meaning Representation* (AMR) (Banarescu et al., 2013). In AMR, sentences are represented as directed acyclic and rooted graphs. An example is displayed in Figure 1, where we see three equivalent displays of an AMR that represents the meaning of the sentence “*The baby is sleeping well*”. In AMR, nodes are variables or concepts, while (labeled) edges express their relations. Among other phenomena, this allows AMR to capture coreference (via re-entrant structures) and semantic roles (via $:arg_n$ relation). Furthermore, AMR links sentences to KBs: e.g., predicates are mapped to PropBank (Palmer et al., 2005; Kingsbury and Palmer, 2002), while named

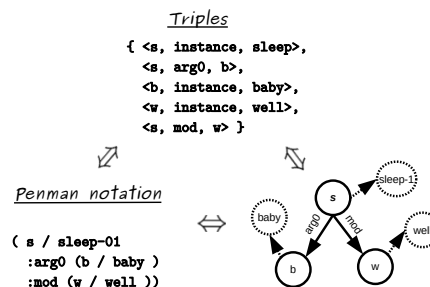


Figure 1: Equivalent representations of the AMR for “*The baby is sleeping well*”.

```
(p4 / possible-01
:arg1 (d5 / destabilize-01
+arg0 [:arg1] (c3 / country
:quant (w2 / whole)))
:condition (e1 / economy
[:poss c3]
:arg0-of (f0 / function-01
[:pol -] )))
```

Figure 2: Parse of *Without a functioning economy, the whole country may destabilize with errors outlined*.

entities are linked to Wikipedia. From a logical perspective, AMR is closely related to first-order logic (FOL, see Bos (2016, 2019) for translation mechanisms).

Currently, AMRs are leveraged to enhance a variety of natural language understanding tasks. E.g., they have enhanced commonsense reasoning and question answering (Mitra and Baral, 2016), machine translation (Song et al., 2019), text summarization (Liao et al., 2018; Dohare et al., 2017) and paraphrasing (Issa et al., 2018). However, there is a critical issue with automatically generated AMRs (parses): they are often deficient.

These deficiencies can be quite severe, even when high-performance parsers are used. For example, in Figure 2, a neural parser (Lyu and Titov, 2018) conducts several errors when parsing *Without a functioning economy the whole country may destabilize*. E.g., it misses a negative polarity and classifies a patient argument as the agent by failing

graph representation	computer processing	human understanding	well-defined
triples	✓(e.g., GNN)	✗	✓
graph visualization	✗	✓(short sentences)	✗
PENMAN, linearized string	✓(e.g., LSTM)	✗	✓
PENMAN, indents	✓(this work)	✓	✓

Table 1: Four (equivalent) AMR representations and their accessibility with respect to human or computer (✓: ‘okay’, ✗: ‘perhaps possible, but difficult’).

to see that *destabilize* here functions as an ergative verb (parser: the country is the causer of destabilize; correct: the country is the object that is destabilized). In sum, the parse has misrepresented the sentence’s meaning.¹ However, assessing such deficiencies via comparison against a gold reference (as in classical parser evaluation) is often infeasible in practice: it takes a trained annotator and appr. 10 minutes to manually create one AMR (Banarescu et al., 2013).

To mitigate these issues, we would like to automatically rate the quality of AMRs without the costly gold graphs. This would allow us to signal downstream task systems the incorporated graphs’ trustworthiness or select among different candidate graphs from different parsing systems. To achieve this, we propose a method that imitates a human rater, who is inspecting the graphs. We show that the method can efficiently rate the quality of the AMRs in the absence of gold graphs.

The remainder of the paper is structured as follows: in Section 2, we outline our idea to exploit the textual multi-line string representation of AMRs, allowing for efficient and simple AMR processing while preserving vital graph structure. In Section 2.2, we instantiate this idea in a lightweight CNN that predicts the quality of AMR graphs along multiple dimensions of interest. In our experiments (Section 3), we show that this framework is efficient and performs better than strong baselines. Our code is available at <https://github.com/flipz357/amr-quality-rater>.

2 AMR as image with latent channels

In this section, we first motivate to treat AMRs as images with latent channels in order to rate them efficiently. Second, we briefly describe the task at hand: Rating the quality of AMR graphs in the absence of gold graphs. Finally, to solve this task, we create a lightweight CNN that evaluates AMR quality in multiple dimensions of interest.

¹?With a functioning economy, the whole country may cause something to destabilize.




Figure 3: Different displays for an AMR structure of a sentence that has medium length (left: PENMAN notation, right: graphical visualization)

The PENMAN notation and its (hidden) advantages

The native AMR notation is called PENMAN-notation or *Sentence Plan Language* (Kasper, 1989; Mann, 1983). Provably, an advantage of this notation is that it allows for secure AMR storage in text files. However, we argue that it has more advantages. For example, due to its clear structure, it allows humans a fairly quick understanding even of medium-sized to large AMR structures (Figure 3, left). On the other hand, we argue that a graphical visualization of such medium-sized to large AMRs (Figure 3, right) could hamper intuitive understanding, since the abundant visual signals (circles, arrows, etc.) may more easily overwhelm humans. Moreover, in every display, one would depend on an algorithm that needs to determine a suitable (and spacious) arrangement of the nodes, edges and edge labels. It may be for these reasons, that in the AMR annotation tool², the graph that is under construction is always shown in PENMAN notation to the human user.

In sum, we find that the indented multi-line PENMAN form possesses three key advantages (Table 1): (i) it enables fairly easy human understanding, (ii), it is well-defined and (iii), which is what we will show next, it can be computationally exploited to better rate AMR quality.

AMR as image to preserve graph structure

Figure 4 describes our proposed sentence representation treatment. After non-degenerate AMR graph simplification (more details in *Preprocessing*, 3.1), we first project the PENMAN representation onto a small grid (‘image’). Each AMR token (e.g., a node or an edge) is represented as a ‘categorical pixel’. Second, Φ adds latent ‘channels’ to the categorical pixels, which can be learned incrementally in an application. In other words, every AMR token is represented by a fixed-sized vector of real

²<https://www.isi.edu/cgi-bin/div3/mt/amr-editor/login-gen-v1.7.cgi>




Figure 4: We transform the (simplified) PENMAN representation to an image and use Φ to add latent channels.

numbers. These vectors are arranged such that the original graph structure is fully preserved.

2.1 Task: Rating the quality of AMR graphs

We aim at rating the quality of AMR graphs (‘parses’) in the absence of gold graphs. This boils down to answering the following question: *how well does a candidate AMR graph capture a given natural language sentence?* Therefore, the exact goal in this task is to learn a mapping

$$f : \mathcal{S} \times \mathcal{G} \rightarrow \mathbb{R}^d, \quad (1)$$

that maps a sentence $s \in \mathcal{S}$ together with a candidate AMR graph $g \in \mathcal{G}$ onto d scores, which describe the AMR with regard to d quality dimensions of interest. A successful mapping function should strongly correlate with the gold scores as they would emerge from evaluation against gold graphs. We proceed by describing the targeted dimensions in more detail.

Main AMR quality dimensions The main quality dimensions that we desire our model to predict are estimated **Smatch F1/recall/precision**. Smatch is the canonical AMR metric, assessing the triple overlap between two graphs, after an alignment step (Cai and Knight, 2013).

AMR sub-task quality dimensions However, we predict also other quality dimensions to assess various AMR aspects (Damonte et al., 2017). In this place, we can merely provide a brief overview: (i) **Unlabeled**: Smatch F1 when disregarding edge-labels. (ii) **No WSD**: Smatch F1 when ignoring ProbBank senses. (iii) **Frames**: PropBank frame identification F1 (iii) **Wikification**: KB linking F1 (iii) **Wikification**: KB linking F1 score on *:wiki* relations. (iv) **Negations**: negation detection F1. (v) **NamedEnt**: NER F1. (vi) **NS frames**: F1 score for ProbBank frame identification when disregarding the sense. (vii) **Concepts** F score for concept identification (viii) **SRL**: Smatch F1 computed on arg-i roles only. (ix) **Reentrancy**: Smatch F1 computed on re-entrant edges only. (x)

IgnoreVars: F1 when variable nodes are ignored. (xi) **Concepts**: F1 for concept detection.

2.2 A lightweight CNN to rate AMR quality

We want to model f (Eq. 1) in order to estimate a suite of quality scores $y \in \mathbb{R}^d$ for any automatically generated AMR graph, given only the graph and the sentence from whence it is derived. Following Opitz and Frank (2019), we will contrast the AMR against the sentence’s dependency parse, exploiting observed structural similarities between these two types of information (Wang et al., 2015). Our proposed method allows this in a simple way by processing dependency and AMR graphs in parallel. The architecture is outlined in Figure 5.

Symbol embedding The latent channels of AMR and dependency ‘pixels’ represent the embeddings of the ‘tokens’ or ‘symbols’ contained in the AMR and dependency vocabulary. These symbols represent nodes or edges. We use two special tokens: the $\langle \text{tab} \rangle$ token, which represents the indentation level, and the $\langle \text{pad} \rangle$ token, which fills the remaining empty ‘pixels’. By embedding lookup, we obtain AMR and dependency images with 128 latent channels and 45×15 ‘pixels’ (Φ in Figure 5; the amount of pixels is chosen such that more than 95% of training AMRs can be fully captured).

Encoding local graph regions Given AMR and dependency images with 128 latent channels and 45×15 pixels, we apply to each of the two images 256 filters of size 3×3 , which is a standard type of kernel in CNNs. This converts both graphs to 256 feature maps each $\in \mathbb{R}^{45 \times 15}$ (same-padding), obtaining two three-dimensional tensors $L_{amr}^1, L_{dep}^1 \in \mathbb{R}^{45 \times 15 \times 256}$. From here, we construct our first joint representation, which matches local dependency regions with local AMR regions:

$$j_{res} = GPF(L_{amr}^1 \otimes L_{dep}^1), \quad (2)$$

where $x \otimes y = [x \odot y; x \ominus y]$ denotes the concatenation of element-wise multiplication and element-wise subtraction. GPF is an operation that performs global pooling and vectorization (‘flattening’) of any input tensor. This means that $j_{res} \in \mathbb{R}^{512}$ is a joint representation of the locally matched dependency and AMR graph regions. This intermediate process is outlined in Figure 5 by \otimes (left) and GPF . Finally, we reduce the dimensions of the two intermediate three-dimensional representations L_{amr}^1 and L_{dep}^1 with 3×3 max-pooling and

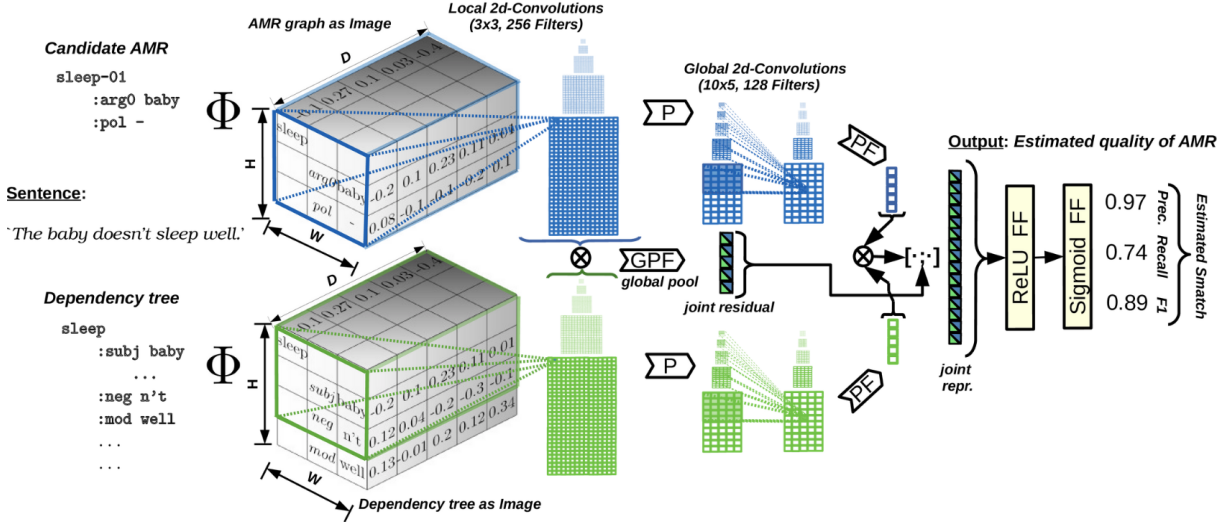


Figure 5: Our proposed architecture for efficient AMR quality assessment.

obtain L_{amr}^2 and $L_{dep}^2 \in \mathbb{R}^{15 \times 5 \times 256}$

Encoding global graph regions For a moment, we put the *joint residual* (j_{res}) aside and proceed by processing the locally convolved feature maps with larger filters. While the first convolutions allowed us to obtain abstract *local* graph regions L_{amr}^2 and L_{dep}^2 , we now aim at matching more *global* regions. More precisely, we use 128 2D filters of shape 10×5 , followed by a 5×5 max-pooling operations on L_{amr}^2 and L_{dep}^2 . Thus, we have obtained vectorized abstract global graph representations $g_{amr}, g_{dep} \in \mathbb{R}^{384}$. Then, we construct a joint representation (right \otimes , Figure 5):

$$j_{glob} = g_{amr} \otimes g_{dep}. \quad (3)$$

At this point, together with the joint residual representation from the local region matching, we have arrived at two joint vector representations j_{glob} and j_{res} . We concatenate them ($[\cdot; \cdot]$ in Figure 5) to form one joint representation $\mathbf{j} \in \mathbb{R}^{1280}$:

$$\mathbf{j} = [j_{res}; j_{glob}] \quad (4)$$

Quality prediction The shared representation \mathbf{j} is further processed by a feed-forward layer with ReLU activation functions (FF_{+ReLU} , Figure 5) and a consecutive feed-forward layer with sigmoid activation functions (FF_{+sigm} , Figure 5):

$$\mathbf{y} = \text{sigm}(\text{ReLU}(\mathbf{j}^T A) B), \quad (5)$$

where $A \in \mathbb{R}^{1280 \times h}$, $B \in \mathbb{R}^{h \times \dim(out)}$ are parameters of the model and $\text{sigm}(x) = (\frac{1}{1+e^{-x_1}}, \dots, \frac{1}{1+e^{-x_{\dim(out)}}})$ projects x onto

$[0, 1]^{\dim(out)}$. When estimating the main AMR metric scores we instantiate three output neurons ($\dim(out) = 3$) that represent estimated Smatch precision, Smatch recall and Smatch F1. In the case where we are interested in a more fine-grained assessment of AMR quality (e.g., knowledge-base linking quality), we have 33 output neurons representing expected scores for various semantic aspects involved in AMR parsing (we predict precision, recall and F1 of 11 aspects, as outlined in §2.1).

To summarize, the residual joint representation should capture local similarities. On the other hand, the second joint representation aims to capture the more global and structural properties of the two graphs. Both types of information inform the final quality assessment of our model in the last layer.

3 Experiments

In this section, we first describe the data, changes to the data that target the reduction of biases, and the baseline. After discussing our main results, we conduct further analyses. (i), we study the effects of our data-debiasing steps. (ii), we assess the performance of our model in a classification task (distinguishing good from bad parses). (iii), we assess the model performance when we only provide the candidate AMR and the sentence (dependency tree ablation). (iv), we provide detailed measurements of the method's computational cost.

3.1 Experimental setup

Data We use the data from Opitz and Frank (2019). The data set consists of more than 15,000

sentences with more than 60,000 corresponding parses, by three different automatic parsing systems and a human. More precisely, the data set $\mathcal{D} = \{(s_i, g_i, y_i)\}_{i=1}^N$ consists of tuples (s_i, g_i, y_i) , where $s_i \in \mathcal{S}$ is a natural language sentence, $g_i \in \mathcal{G}$ is a ‘candidate’ AMR graph and $y_i \in \mathbb{R}^d$ is a 36-dimensional vector containing scores which represent the quality of the AMR graph in terms of precision, recall and F1 with respect to 12 different tasks captured by AMR (as outlined in §2.1).

Debiasing of the data We observe three biases in the data. First, the graphs in the training section of our data are less deficient than in the development and testing data, because the parsers were trained on *(sentence, gold graph)* pairs from the training section. For our task, this means that the training section’s target scores are higher, on average, than the target scores in the other data partitions. To achieve more balance in this regard, we re-split the data randomly on the sentence-id level (such that a sentence does not appear in more than one partition with different parses).

Second, we observe that the data contains some superficial hidden clues that could give away the parse’s source. This bears the danger that a model does not learn to assess the *parse quality*, but to assess the *source of the parse*. And since some parsers are better or worse than others, the model could exploit this bias. For example, consider that one parser prefers to write *(r / run-01 :arg1 (c / cat) :polarity -)*, while the other parser prefers to write *(r / run-01 :polarity - :arg1 (c / cat))*. These two structures are semantically equivalent but differ on the surface. Hence, the arrangement of the output may provide unwanted clues on the source of the parse. To alleviate this issue, we randomly re-arrange all parses on the surface, keeping their semantics.³⁴

A third bias stems from a design choice in the metric scripts used to calculate the target scores. More precisely, the extended *Smatch*-metric script, per default, assigns a parse that does not contain a certain edge-type (e.g., $:\text{arg}_n$) the score 0 with respect to the specific quality dimension (in this case, SRL: 0.00 Precision/Recall/F1). However, if the gold parse also does not contain an

³Technically, this is achieved by reformatting the parses such that in the depth-first writing-traversal at node n the out-going edges of n will be traversed in random order.

⁴Different variable names, e.g., *(r / run-01)* and *(x / run-01)* are not an issue in this work since the variables are handled via (van Noord and Bos, 2017a). See also *Preprocessing*, §3.1

edge of this type (i.e., $:\text{arg}_n$), then we believe that the correct default score should be 1, since the parse is, in the specific dimension, in perfect agreement with the gold (i.e., SRL: 1.00 Precision/Recall/F1). Therefore, we set all sub-task scores, where the predicted graph agrees with the gold graph in the absence of a feature, from 0 to 1.

Preprocessing Same as prior work, we dependency-parse and tokenize the sentences with *spacy* (Honnibal and Montani, 2017) and replace variables with corresponding concepts (e.g., *(j / jump-01 :arg0 (g / girl))* is translated to *(jump-01 :arg0 (girl))*. Re-entrancies are handled with pointers according to van Noord and Bos (2017a), which ensures non-degenerate AMR simplification.⁵ Furthermore, we lower-case all tokens, remove quotation marks and join sub-structures that represent names.⁶ The vocabulary encompasses all tokens of frequency ≥ 5 , remaining ones are set to `<unk>`.

Training All parameters are initialized randomly. We train for 5 epochs and select the parameters θ from the epoch where maximum development scores were achieved (with respect to average Pearson’s ρ over the quality dimensions). In training, we reduce the squared error with gradient descent (Adam rule (Kingma and Ba, 2019), learning rate = 0.001, mini batch size = 64):

$$\theta^* = \arg \min_{\theta} \sum_{i=1}^{|\mathcal{D}|} \sum_{j=1}^{|\mathcal{M}|} (y_{i,j} - f_{\theta}(s_i, g_i)_j)^2, \quad (6)$$

where \mathcal{M} is the set of target metrics.

Baseline Our main baseline is the model of previous work, henceforth denoted by **LG-LSTM**. The method works in the following steps: first, it uses a depth-first graph traversal to linearize the automatic AMR graph and the corresponding dependency tree of the sentence. Second, it constructs a joint representation and predicts the score estimations. To further improve its performance, the baseline uses some extra-features (e.g., a shallow alignment from

⁵For example, consider the sentence *The cat scratches itself* and its graph *(x / scratch-01 :arg0 (y / cat) :arg1 y)*. Replacing the variables with concepts would come at the cost of an information loss w.r.t. to coreference: *(scratch-01 :arg0 cat :arg1 cat)* — does the cat scratch itself or another cat? Hence, pointers are used to translate the graph into *(scratch-01 :arg0 *0* cat :arg1 *0*)*.

⁶E.g., *:name (name :op1 ‘Barack’ :op2 ‘Obama’)* is translated to *:name barack obama*.

	Smatch	Ridge	GNN	LG-LSTM	ours	change %
P's ρ	F1	0.428	0.659	0.662 \pm 0.00	0.696 \pm 0.00	+5.14 †‡
	Precision	0.348	0.601	0.600 \pm 0.00	0.623 \pm 0.01	+3.83 †
	Recall	0.463	0.667	0.676 \pm 0.00	0.719 \pm 0.00	+6.36 †‡
RMSE	F1	0.155	0.132	0.130 \pm 0.00	0.128 \pm 0.00	-1.54
	Precision	0.146	0.127	0.126 \pm 0.00	0.126 \pm 0.00	+0.0
	Recall	0.169	0.141	0.142 \pm 0.00	0.136 \pm 0.00	-4.23

Table 2: Main results. Pearson’s corr. coefficient (row 1-3) is better if higher; root mean square error (RMSE, row 4-6) is better if lower. The quality dimensions are explained in §2.1. † (‡): $p < 0.05$ ($p < 0.005$), significant difference in the correlations with two-tailed test using Fisher ρ to z transformation (Fisher, 1915).

dependency tokens to AMR tokens).⁷ Generally speaking, the baseline is a model that works based on graph linearizations. Such type of model, despite its apparent simplicity, has proven to be an effective baseline or state-of-the-art method in various works about converting texts into graphs (Konstas et al., 2017; van Noord and Bos, 2017b), or converting graphs into texts (Bastings et al., 2017; Beck et al., 2018; Song, 2019; Pourdamghani et al., 2016; Song et al., 2018; Vinyals et al., 2015; Mager et al., 2020), or performing mathematically complex tasks modeled as graph-to-graph problems, such as symbolic integration (Lample and Charton, 2020). However, in our main results, we also display the results of two additional baselines: **GNN** (Song et al., 2018), where we encode the dependency tree and the AMR with a graph-recurrent encoder and perform regression on the joint averaged node embedding vectors.⁸ And **Ridge**, an l2-regularized linear regression that is based on shallow graph statistics.⁹

3.2 Results

Main AMR quality dimensions The main quality of an AMR graph is estimated in expected triple match ratios (Smatch F1, Precision and Recall). The results, averaged over 10 runs, are displayed in Table 2. With regard to estimated Smatch F1, we

⁷Furthermore, the baseline uses auxiliary losses to achieve a slight performance gain in predicting the Smatch metrics. For the sake of simplicity, we do not use these auxiliary losses, except in one experiment, where we show that our method achieves a similar small gain with the auxiliary losses.

⁸ $\left[\frac{1}{|V_A|} \sum_{v \in V_A} emb(v) \right] \otimes \left[\frac{1}{|V_D|} \sum_{v \in V_D} emb(v) \right]$.

⁹For the dependency graph (D) and the AMR graph (A) we both compute $\phi(A|D) = [\text{density, avg. node degree, node count, edge count, (arg0|subj) count, (arg1|obj) count}]$, the final feature vector then is defined as $\Phi(x) = [\phi(A) - \phi(D); \phi(D); \phi(A); \frac{|\text{lemmas}(D) \cap \text{concepts}(A)|}{|\text{lemmas}(D) \cup \text{concepts}(A)|}]$

	Quality Dim.	LG-LSTM	ours	change %
F1 Pearson's ρ	Concepts	0.508 \pm 0.01	0.545 \pm 0.01	+7.28 †
	Frames	0.420 \pm 0.01	0.488 \pm 0.01	+16.19 †‡
	IgnoreVars	0.627 \pm 0.01	0.665 \pm 0.00	+6.06 †‡
	NamedEnt.	0.429 \pm 0.02	0.460 \pm 0.01	+7.23 †
	Negations	0.685 \pm 0.02	0.746 \pm 0.01	+8.91 †‡
	NoWSD	0.640 \pm 0.01	0.680 \pm 0.00	+6.25 †‡
	NS-frames	0.419 \pm 0.02	0.505 \pm 0.01	+20.53 †‡
	Reentrancies	0.508 \pm 0.01	0.602 \pm 0.00	+18.50 †‡
	SRL	0.519 \pm 0.01	0.581 \pm 0.01	+11.95 †‡
	Unlabeled	0.628 \pm 0.01	0.663 \pm 0.00	+5.57 †‡
Wikification	0.901 \pm 0.00	0.904 \pm 0.00	+0.33	
F1 RMSE	Concepts	0.117 \pm 0.00	0.114 \pm 0.00	-2.56
	Frames	0.186 \pm 0.00	0.182 \pm 0.00	-2.15
	IgnoreVars	0.195 \pm 0.00	0.186 \pm 0.00	-4.62
	NamedEnt.	0.159 \pm 0.00	0.156 \pm 0.00	-1.89
	Negations	0.197 \pm 0.00	0.180 \pm 0.00	-8.63
	NoWSD	0.132 \pm 0.00	0.126 \pm 0.00	-4.55
	NS-frames	0.157 \pm 0.00	0.155 \pm 0.00	-1.27
	Reentrancies	0.285 \pm 0.00	0.265 \pm 0.00	-7.02
	SRL	0.189 \pm 0.00	0.181 \pm 0.00	-4.23
	Unlabeled	0.124 \pm 0.00	0.121 \pm 0.00	-2.42
Wikification	0.165 \pm 0.00	0.162 \pm 0.00	-1.82	

Table 3: Results for AMR quality rating w.r.t. various sub-tasks. † (‡): significance (c.f. caption Table 2).

achieve a correlation with the gold scores of 0.695 Pearson’s ρ . This constitutes a significant improvement of appr. 5% over LG-LSTM. Similarly, recall and precision correlations improve by 6.36% and 3.83 % (from 0.676 to 0.719 and 0.600 to 0.623). While the improvement in predicted recall is significant at $p < 0.05$ and $p < 0.005$, the improvement in predicted precision is significant at $p < 0.05$. When we consider the root mean square error (RMSE), we find that the method improves over the best baseline by -1.54% in estimated Smatch F1 and -4.23% in estimated Smatch recall. On the other hand, the RMS error in estimated precision remains unchanged.

AMR subtask quality Our model can also rate the quality of an AMR graph in a more fine-grained way. The results are displayed in Table 3. Over almost every dimension we see considerable improvements. For instance, a considerable improvement in Pearson’s ρ is achieved for assessment of *frame prediction quality* (‘NSFrames’ in Table 3, +20.5% ρ) and *coreference quality* (‘Reentrancies’ in Table 3, +18.5%).

A substantial error reduction is achieved in *polarity* (‘Negations’, Table 3), where we reduce the RMSE of the estimated F1 score by -8.6%. When rating the SRL-quality of an AMR parse, our model reduces the RMSE by appr. 4%. In general,

data	method	Pearson’s ρ			error
		P	R	F1	RMSE (F1)
$\frac{0}{2}$	LG-LSTM	0.72	0.78	0.77	0.138
	LG-LSTM _{+aux}	0.74	0.79	0.78	0.137
	ours	0.75	0.80	0.79	0.133
	ours _{+aux}	0.76	0.81	0.80	0.132
$\frac{1}{2}$	LG-LSTM	0.67	0.73	0.72	0.120
	ours	0.68	0.75	0.74	0.117
$\frac{2}{2}$	LG-LSTM	0.60	0.68	0.66	0.130
	ours	0.62	0.72	0.70	0.128

Table 4: Performance-effects of data debiasing steps. _{+aux} indicates a model variant that is trained using auxiliary losses that incorporate information about the other AMR aspects in the training process (see Fn.7).

improvements are obtained over almost all tested quality dimensions, both in RMSE reduction and increased correlation with the gold scores.

3.3 Analysis

Effect of data debiasing We want to study the effect of the data set cleaning steps by analyzing the performance of our method and the baseline on three different versions of the data, with respect to estimated Smatch scores. The three versions are (i) $\frac{0}{2}$ = **AMRQUALITY**, which is the original data; (ii) $\frac{1}{2}$, which is the data after the random re-split and score correction; (iii) $\frac{2}{2}$ = **AMRQUALITYCLEAN** which is our main data after the final debiasing step (shallow structure debiasing) has been applied.

The results are shown in Table 4. We can make three main observations: (i) from the first to the second debiasing step, the baseline and our model have in common that Pearson’s ρ and the error decrease. While we cannot exactly explain why ρ decreases, it is somewhat in line with recent research that observed performance drops when data was re-split (Gorman and Bedrick, 2019). On the other hand, the error decrease can be explained by the random re-split that balances the target scores. (ii) The second debiasing step leads to a decrease in ρ and an increase in error, for both models. This indicates that we have successfully removed shallow biases from the data that can give away the parse’s source. (iii) On all considered versions of the data, the method performs better than the baseline.

AMRs: telling the good from the bad In this experiment, we want to see how well the model can discriminate between good and bad graphs. To this aim, we create a five-way classification task: graphs are assigned the label ‘very bad’ (Smatch F1 < 0.25), ‘bad’ (0.25 \geq Smatch F1 < 0.5), ‘good’

	majority	random	LG-LSTM	ours
avg. F1	0.13	0.20	0.40	0.44 ^{†‡}
quadr. kappa	0.0	0.03	0.53	0.60 ^{†‡}

Table 5: Graph quality classification task. † (‡) significance with paired t-test at p<0.05 (p<0.005) over 10 random initializations.

	Quality Dim.	LG-LSTM	ours	ours (no dep.)
P’s ρ	Smatch F1	0.662 \pm 0.00	0.696 \pm 0.00	0.682 \pm 0.01
	Smatch precision	0.600 \pm 0.00	0.623 \pm 0.01	0.614 \pm 0.01
	Smatch recall	0.676 \pm 0.00	0.719 \pm 0.00	0.702 \pm 0.01
RMSE	Smatch F1	0.130 \pm 0.00	0.128 \pm 0.00	0.128 \pm 0.00
	Smatch precision	0.126 \pm 0.00	0.126 \pm 0.00	0.129 \pm 0.00
	Smatch recall	0.142 \pm 0.00	0.136 \pm 0.00	0.139 \pm 0.00

Table 6: Right column: results of our system when we abstain from feeding the dependency tree, and only show the sentence together with the candidate AMR.

(0.5 \geq Smatch F1 < 0.75), ‘very good’ (0.75 \geq Smatch F1 < 0.95) and ‘excellent’ (Smatch F1 \geq 0.95). Here, we do not retrain the models with a classification objective but convert the estimated Smatch F1 to the corresponding label. Since the classes are situated on a nominal scale, and ordinary classification metrics would not fully reflect the performance, we also use quadratic weighted kappa (Cohen, 1968) for evaluation.

The results are shown in Table 5. All baselines, including LG-LSTM, are significantly outperformed by our approach, both in terms of macro F1 (+4 points, 10% improvement) and quadratic kappa (+7 points, 13% improvement).

How important is the dependency information?

To investigate this question, instead of feeding the dependency tree of the sentence, we only feed the sentence itself. To achieve this, we simply insert the tokens in the first row of the former dependency input image, and pad all remaining empty ‘pixels’. In this mode, the sentence encoding is similar to standard convolutional sentence encoders as they are typically used in many tasks (Kim, 2014).

The results are shown in the right column of Table 6. The performance drops are small but consistent across all analyzed dimensions, both in terms of error (0 to 2.2% increase) and Pearson’s ρ (1.4 to 2.4% decrease). This indicates that the dependency trees contain information that can be exploited by our model to better judge the AMR quality. We hypothesize that this is due to similarities between relations such as *subj/obj* (syntactic) or *arg0/arg1* (semantic), etc. Yet, we see that this simpler model,

GPU type	GTX Titan		GTX 1080	
method	LG-LSTM	ours	LG-LSTM	ours
avg. ep. time	722s	59s	1582s	64s
avg. W	105	166	45	128
kWh per epoch	0.021	0.003	0.020	0.002

Table 7: Efficiency analysis of two approaches.




Figure 6: Training cost diagram of two approaches.

which does not see the dependency tree, still outperforms the baseline, except in estimated precision, where the error is increased by 2.4%.

Efficiency analysis Recently, in many countries, there have been efforts to reduce energy consumption and carbon emission. Since deep learning typically requires intensive GPU computing, this aspect is of increasing importance to researchers and applicants (Strubell et al., 2019; Tang et al., 2019; Ganguly et al., 2019). To investigate energy consumption of our method and previous work, we monitor their GPU usage during training, assessing the following quantities : (i) avg. time per epoch, (ii) avg. watts GPU usage, (iii) kilowatts per epoch (in kWh).

The results of this analysis are displayed in Table 7 and outlined in Figure 6. Our method consumes approximately 6.6 times less total kWh on a GTX Titan (10 times less on a GTX 1080). Directly related, it also reduces the training time: prior work requires appr. 1500s training time per epoch (GTX 1080), while our method requires appr. 60s per epoch (GTX 1080). The main reason for this is that our model does not depend on recurrent operations and profits more from parallelism.

4 Related work

Quality measurement of structured predictions

Since evaluating structured representations against human annotations is costly, systems have been developed that attempt an automatic quality assessment of these structures. Due to its popularity, much work has been conducted in *machine translation* (MT) under the umbrella of *quality estimation* (QE). QE can take place either on a word-level (Martins et al., 2017), sentence-level (Spe-

cia et al., 2009), or document-level (Scarton et al., 2015). The conference on Machine Translation (WMT) has a long-standing workshop and shared task-series on MT quality assessment (Bojar et al., 2013, 2014, 2015, 2016, 2017; Specia et al., 2018; Fonseca et al., 2019). Quality estimation for neural language generation has been investigated, i.a., by Scarton et al. (2016), and recently by Dušek et al. (2019), who design a model that jointly learns to rate and rank generations, or by Zopf (2018), who predicts pair-wise preferences for generated summaries.

Furthermore, automatic techniques for the quality assessment of syntactic parses have been proposed. For instance, Ravi et al. (2008) formulate the task as a single-variable regression problem to assess the quality of constituency trees. A major difference to our work is that they try to assess the performance of a *single parser*, while we aim at a parser-agnostic setting where candidate parses stem from *different parsers*. Similarly, Kawahara and Uchimoto (2008) predict a binary label that reflects whether the tree-quality lies above a certain threshold (or not). When multiple candidate parses are available, tree ranking methods (Zhu et al., 2015; Zhou et al., 2016) may also be interpreted as some form of parse quality assessment (see Do and Rehbein (2020) for a recent overview). Compared with assessing the quality of (abstract) meaning representations, judging about syntactic trees perhaps is a conceptually slightly simpler task, since the syntactic graphs are more directly grounded in the sentence¹⁰, and therefore it may be easier to judge whether graph components are correct, redundant, missing, or false.

In comparison to MT, automatic quality assessment of meaning representations is insufficiently researched. Opitz and Frank (2019) propose an LSTM based model that performs a multi-variate quality analysis of AMRs (constituting the baseline which we compared against). We believe that quality estimation approaches may also prove valuable for other meaning representation formalisms (MRs), such as, e.g., discourse representations (Kamp and Reyle, 1993; Kamp, 2008; Abzianidze et al., 2019) or universal semantic dependencies (Reisinger et al., 2015; Stengel-Eskin et al., 2020). For example, since the manual creation of MRs

¹⁰In dependency trees, nodes are words; in constituency trees, nodes are (labeled) phrases; in meaning representations, words or phrases may be projected to abstract semantic nodes, or they may be omitted.

is a notoriously laborious task, automatic quality assessment tools could assist humans in the annotation process (e.g., by serving as a cheap annotation quality check or by filtering automatic parses in active learning).

AMR metrics When a gold graph is available, it can be used to compute the canonical AMR metric Smatch (Cai and Knight, 2013) that assesses matching triples. Furthermore, Damonte et al. (2017) have extended Smatch to inspect various aspects of AMR. In this work, we have shown that our model can predict the expected outcomes of these metrics in the absence of the gold graph. Recently, more AMR metrics have been proposed, for example the Bleu-based (Papineni et al., 2002) SemBleu metric (Song and Gildea, 2019), Sema (Anchiêta et al., 2019) or S²match (Opitz et al., 2020), a variant of Smatch. We plan to extend our model such that it also predicts these metrics.

AMR parsing Recent advances in AMR parsing have been achieved by parsers that either predict latent alignments jointly with nodes (Lyu and Titov, 2018), or by transducing a graph from a sequence with a minimum spanning tree (MST) decoding algorithm (Zhang et al., 2019), or by focusing on core semantics in a top-down fashion (Cai and Lam, 2019), or by performing auto-regressive decoding with a graph encoder (Cai and Lam, 2020). Other approaches apply statistical machine translation (Pust et al., 2015) or sequence-to-sequence models, which tend to suffer from data scarcity issues and need considerable amounts of silver data to improve results (van Noord and Bos, 2017c; Konstas et al., 2017). Previously, alignment-based pipeline models have proved effective (Flanigan et al., 2014) or transition-based approaches that convert dependency trees step-by-step to AMR graphs (Wang et al., 2015, 2016; Lindemann et al., 2020).

5 Conclusion

In this work, we have developed an approach to rate the quality of AMR graphs in the absence of costly gold data. Our model imitates a human judge that is confronted, ‘on paper’, with the AMR in its native multi-line Penman format. We saw how this setup allowed efficient AMR processing with convolutions. Our experiments indicate that the method rates AMR quality more accurately and more efficiently than previous work.

Acknowledgments

I am grateful to the anonymous reviewers for their valuable thoughts and comments. Moreover, I am grateful to Anette Frank for her thoughtful feedback on an earlier draft of this paper and her general guidance throughout my studies.

References

- Lasha Abzianidze, Rik van Noord, Hessel Haagsma, and Johan Bos. 2019. [The first shared task on discourse representation structure parsing](#). In *Proceedings of the IWCS Shared Task on Semantic Parsing*, Gothenburg, Sweden. Association for Computational Linguistics.
- Rafael Torres Anchiêta, Marco Antonio Sobrevilla Cabezudo, and Thiago Alexandre Salgueiro Pardo. 2019. Sema: an extended semantic evaluation for amr. In *(To appear) Proceedings of the 20th Computational Linguistics and Intelligent Text Processing*. Springer International Publishg.
- 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. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 178–186.
- Joost Bastings, Ivan Titov, Wilker Aziz, Diego Marcheggiani, and Khalil Sima’an. 2017. [Graph convolutional encoders for syntax-aware neural machine translation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1957–1967, Copenhagen, Denmark. Association for Computational Linguistics.
- Daniel Beck, Gholamreza Haffari, and Trevor Cohn. 2018. [Graph-to-sequence learning using gated graph neural networks](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 273–283, Melbourne, Australia. Association for Computational Linguistics.
- Ondřej Bojar, Christian Buck, Chris Callison-Burch, Christian Federmann, Barry Haddow, Philipp Koehn, Christof Monz, Matt Post, Radu Soricut, and Lucia Specia. 2013. [Findings of the 2013 Workshop on Statistical Machine Translation](#). In *Proceedings of the Eighth Workshop on Statistical Machine Translation*, pages 1–44, Sofia, Bulgaria. Association for Computational Linguistics.
- Ondřej Bojar, Christian Buck, Christian Federmann, Barry Haddow, Philipp Koehn, Johannes Leveling, Christof Monz, Pavel Pecina, Matt Post, Herve Saint-Amand, Radu Soricut, Lucia Specia, and Aleš Tamchyna. 2014. [Findings of the 2014 workshop on statistical machine translation](#). In *Proceedings of the*

- Ninth Workshop on Statistical Machine Translation*, pages 12–58, Baltimore, Maryland, USA. Association for Computational Linguistics.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Shujian Huang, Matthias Huck, Philipp Koehn, Qun Liu, Varvara Logacheva, Christof Monz, Matteo Negri, Matt Post, Raphael Rubino, Lucia Specia, and Marco Turchi. 2017. [Findings of the 2017 conference on machine translation \(WMT17\)](#). In *Proceedings of the Second Conference on Machine Translation*, pages 169–214, Copenhagen, Denmark. Association for Computational Linguistics.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Aurélie Névéol, Mariana Neves, Martin Popel, Matt Post, Raphael Rubino, Carolina Scarton, Lucia Specia, Marco Turchi, Karin Verspoor, and Marcos Zampieri. 2016. [Findings of the 2016 conference on machine translation](#). In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 131–198, Berlin, Germany. Association for Computational Linguistics.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Barry Haddow, Matthias Huck, Chris Hokamp, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Matt Post, Carolina Scarton, Lucia Specia, and Marco Turchi. 2015. [Findings of the 2015 workshop on statistical machine translation](#). In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 1–46, Lisbon, Portugal. Association for Computational Linguistics.
- Johan Bos. 2016. [Expressive power of abstract meaning representations](#). *Computational Linguistics*, 42(3):527–535.
- Johan Bos. 2019. Separating argument structure from logical structure in amr. *arXiv preprint arXiv:1908.01355*.
- Deng Cai and Wai Lam. 2019. [Core semantic first: A top-down approach for AMR parsing](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3797–3807, Hong Kong, China. Association for Computational Linguistics.
- Deng Cai and Wai Lam. 2020. Amr parsing via graph-sequence iterative inference. *arXiv preprint arXiv:2004.05572*.
- Shu Cai and Kevin Knight. 2013. [Smatch: an evaluation metric for semantic feature structures](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 748–752, Sofia, Bulgaria. Association for Computational Linguistics.
- Jacob Cohen. 1968. Weighted kappa: nominal scale agreement with provision for scaled disagreement or partial credit. *Psychological bulletin*, pages 213–220.
- 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, Valencia, Spain. Association for Computational Linguistics.
- Bich-Ngoc Do and Ines Rehbein. 2020. [Neural reranking for dependency parsing: An evaluation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4123–4133, Online. Association for Computational Linguistics.
- Shibhansh Dohare, Harish Karnick, and Vivek Gupta. 2017. Text summarization using abstract meaning representation. *arXiv preprint arXiv:1706.01678*.
- Ondřej Dušek, Karin Sevegnani, Ioannis Konstas, and Verena Rieser. 2019. [Automatic quality estimation for natural language generation: Ranting \(jointly rating and ranking\)](#). In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 369–376, Tokyo, Japan. Association for Computational Linguistics.
- Ronald A Fisher. 1915. Frequency distribution of the values of the correlation coefficient in samples from an indefinitely large population. *Biometrika*, 10(4):507–521.
- 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, Baltimore, Maryland. Association for Computational Linguistics.
- Erick Fonseca, Lisa Yankovskaya, André F. T. Martins, Mark Fishel, and Christian Federmann. 2019. [Findings of the WMT 2019 shared tasks on quality estimation](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 3: Shared Task Papers, Day 2)*, pages 1–10, Florence, Italy. Association for Computational Linguistics.
- A. Ganguly, R. Muralidhar, and V. Singh. 2019. [Towards energy efficient non-von neumann architectures for deep learning](#). In *20th International Symposium on Quality Electronic Design (ISQED)*, pages 335–342.
- Kyle Gorman and Steven Bedrick. 2019. [We need to talk about standard splits](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2786–2791, Florence, Italy. Association for Computational Linguistics.

- Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. To appear.
- Fuad Issa, Marco Damonte, Shay B. Cohen, Xiaohui Yan, and Yi Chang. 2018. [Abstract meaning representation for paraphrase detection](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 442–452, New Orleans, Louisiana. Association for Computational Linguistics.
- Hans Kamp. 2008. *A Theory of Truth and Semantic Representation*, chapter 8. John Wiley & Sons, Ltd.
- Hans Kamp and Uwe Reyle. 1993. *From Discourse to Logic. Introduction to Model-theoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Kluwer, Dordrecht.
- Robert T. Kasper. 1989. [A flexible interface for linking applications to penman’s sentence generator](#). In *Proceedings of the Workshop on Speech and Natural Language*, HLT ’89, pages 153–158, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Daisuke Kawahara and Kiyotaka Uchimoto. 2008. Learning reliability of parses for domain adaptation of dependency parsing. In *Proceedings of the Third International Joint Conference on Natural Language Processing: Volume-II*.
- Yoon Kim. 2014. [Convolutional neural networks for sentence classification](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751, Doha, Qatar. Association for Computational Linguistics.
- Diederik P Kingma and J Adam Ba. 2019. A method for stochastic optimization. arxiv 2014. *arXiv preprint arXiv:1412.6980*, 434.
- Paul Kingsbury and Martha Palmer. 2002. [From TreeBank to PropBank](#). In *Proceedings of the Third International Conference on Language Resources and Evaluation (LREC’02)*, Las Palmas, Canary Islands - Spain. European Language Resources Association (ELRA).
- 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, Vancouver, Canada. Association for Computational Linguistics.
- Guillaume Lample and François Charton. 2020. [Deep learning for symbolic mathematics](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Kexin Liao, Logan Lebanoff, and Fei Liu. 2018. [Abstract meaning representation for multi-document summarization](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1178–1190, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Matthias Lindemann, Jonas Groschwitz, and Alexander Koller. 2020. Fast semantic parsing with well-typedness guarantees. *arXiv preprint arXiv:2009.07365*.
- 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, Melbourne, Australia. Association for Computational Linguistics.
- Manuel Mager, Ramón Fernandez Astudillo, Tahira Naseem, Md Arafat Sultan, Young-Suk Lee, Radu Florian, and Salim Roukos. 2020. [GPT-too: A language-model-first approach for AMR-to-text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1846–1852, Online. Association for Computational Linguistics.
- William C. Mann. 1983. [An overview of the penman text generation system](#). In *Proceedings of the Third AAAI Conference on Artificial Intelligence*, AAAI’83, pages 261–265. AAAI Press.
- André FT Martins, Marcin Junczys-Dowmunt, Fabio N Kepler, Ramón Astudillo, Chris Hokamp, and Roman Grundkiewicz. 2017. Pushing the limits of translation quality estimation. *Transactions of the Association for Computational Linguistics*, 5:205–218.
- 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*, AAAI’16, page 2779–2785. AAAI Press.
- Rik van Noord and Johan Bos. 2017a. [Dealing with co-reference in neural semantic parsing](#). In *Proceedings of the 2nd Workshop on Semantic Deep Learning (SemDeep-2)*, pages 41–49, Montpellier, France. Association for Computational Linguistics.
- Rik van Noord and Johan Bos. 2017b. Neural semantic parsing by character-based translation: Experiments with abstract meaning representations. *arXiv preprint arXiv:1705.09980*.
- Rik van Noord and Johan Bos. 2017c. Neural semantic parsing by character-based translation: Experiments with abstract meaning representations. *Computational Linguistics in the Netherlands Journal*, 7:93–108.

- Juri Opitz and Anette Frank. 2019. [Automatic accuracy prediction for AMR parsing](#). In *Proceedings of the Eighth Joint Conference on Lexical and Computational Semantics (*SEM 2019)*, pages 212–223, Minneapolis, Minnesota. Association for Computational Linguistics.
- Juri Opitz, Letitia Parcalabescu, and Anette Frank. 2020. [Amr similarity metrics from principles](#). *Transactions of the Association for Computational Linguistics*, 8:522–538.
- Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. The proposition bank: An annotated corpus of semantic roles. *Computational linguistics*, 31(1):71–106.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Nima Pourdamghani, Kevin Knight, and Ulf Hermjakob. 2016. [Generating English from abstract meaning representations](#). In *Proceedings of the 9th International Natural Language Generation conference*, pages 21–25, Edinburgh, UK. Association for Computational Linguistics.
- Michael Pust, Ulf Hermjakob, Kevin Knight, Daniel Marcu, and Jonathan May. 2015. [Parsing English into abstract meaning representation using syntax-based machine translation](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1143–1154, Lisbon, Portugal. Association for Computational Linguistics.
- Sujith Ravi, Kevin Knight, and Radu Soricut. 2008. Automatic prediction of parser accuracy. In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 887–896.
- Drew Reisinger, Rachel Rudinger, Francis Ferraro, Craig Harman, Kyle Rawlins, and Benjamin Van Durme. 2015. [Semantic proto-roles](#). *Transactions of the Association for Computational Linguistics*, 3:475–488.
- Carolina Scarton, Gustavo Paetzold, and Lucia Specia. 2016. [Quality estimation for language output applications](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Tutorial Abstracts*, pages 14–17, Osaka, Japan. The COLING 2016 Organizing Committee.
- Carolina Scarton, Marcos Zampieri, Mihaela Vela, Josef van Genabith, and Lucia Specia. 2015. [Searching for context: a study on document-level labels for translation quality estimation](#). In *Proceedings of the 18th Annual Conference of the European Association for Machine Translation*, pages 121–128, Antalya, Turkey.
- Linfeng Song. 2019. [Tackling graphical NLP problems with graph recurrent networks](#). *CoRR*, abs/1907.06142.
- Linfeng Song and Daniel Gildea. 2019. [SemBleu: A robust metric for AMR parsing evaluation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4547–4552, Florence, Italy. Association for Computational Linguistics.
- Linfeng Song, Daniel Gildea, Yue Zhang, Zhiguo Wang, and Jinsong Su. 2019. [Semantic neural machine translation using AMR](#). *Transactions of the Association for Computational Linguistics*, 7:19–31.
- Linfeng Song, Yue Zhang, Zhiguo Wang, and Daniel Gildea. 2018. [A graph-to-sequence model for AMR-to-text generation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1616–1626, Melbourne, Australia. Association for Computational Linguistics.
- Lucia Specia, Frédéric Blain, Varvara Logacheva, Ramón Astudillo, and André F. T. Martins. 2018. [Findings of the WMT 2018 shared task on quality estimation](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 689–709, Belgium, Brussels. Association for Computational Linguistics.
- Lucia Specia, Nicola Cancedda, Marc Dymetman, Marco Turchi, and Nello Cristianini. 2009. Estimating the sentence-level quality of machine translation systems. In *In EAMT*, pages 28–35.
- Elias Stengel-Eskin, Aaron Steven White, Sheng Zhang, and Benjamin Van Durme. 2020. [Universal decompositional semantic parsing](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8427–8439, Online. Association for Computational Linguistics.
- Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2019. [Energy and policy considerations for deep learning in NLP](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 3645–3650. Association for Computational Linguistics.
- Zhenheng Tang, Yuxin Wang, Qiang Wang, and Xiaowen Chu. 2019. [The impact of gpu dvfs on the energy and performance of deep learning: An empirical study](#). In *Proceedings of the Tenth ACM International Conference on Future Energy Systems, e-Energy '19*, pages 315–325, New York, NY, USA. ACM.

- Oriol Vinyals, Łukasz Kaiser, Terry Koo, Slav Petrov, Ilya Sutskever, and Geoffrey Hinton. 2015. [Grammar as a foreign language](#). In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 2773–2781. Curran Associates, Inc.
- Chuan Wang, Sameer Pradhan, Xiaoman Pan, Heng Ji, and Nianwen Xue. 2016. [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. 2015. [A transition-based algorithm for AMR parsing](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 366–375. Denver, Colorado. Association for Computational Linguistics.
- Sheng Zhang, Xutai Ma, Kevin Duh, and Benjamin Van Durme. 2019. [AMR parsing as sequence-to-graph transduction](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 80–94. Florence, Italy. Association for Computational Linguistics.
- Hao Zhou, Yue Zhang, Shujian Huang, Junsheng Zhou, Xin-Yu Dai, and Jiajun Chen. 2016. [A search-based dynamic reranking model for dependency parsing](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1393–1402. Berlin, Germany. Association for Computational Linguistics.
- Chenxi Zhu, Xipeng Qiu, Xinchu Chen, and Xuanjing Huang. 2015. [A re-ranking model for dependency parser with recursive convolutional neural network](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1159–1168. Beijing, China. Association for Computational Linguistics.
- Markus Zopf. 2018. [Estimating summary quality with pairwise preferences](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1687–1696. New Orleans, Louisiana. Association for Computational Linguistics.