

Multilingual AMR Parsing with Noisy Knowledge Distillation*

Deng Cai[♡] Xin Li[♣] Jackie Chun-Sing Ho[♡] Lidong Bing[♣] Wai Lam[♡]

[♡]The Chinese University of Hong Kong

[♣]DAMO Academy, Alibaba Group

thisisjcyd@gmail.com

{xinting.lx, l.bing}@alibaba-inc.com

{schun, wlam}@se.cuhk.edu.hk

Abstract

We study multilingual AMR parsing from the perspective of knowledge distillation, where the aim is to learn and improve a multilingual AMR parser by using an existing English parser as its teacher. We constrain our exploration in a strict multilingual setting: there is but one model to parse all different languages including English. We identify that noisy input and precise output are the key to successful distillation. Together with extensive pre-training, we obtain an AMR parser whose performances surpass all previously published results on four different foreign languages, including German, Spanish, Italian, and Chinese, by large margins (up to 18.8 SMATCH points on Chinese and on average 11.3 SMATCH points). Our parser also achieves comparable performance on English to the latest state-of-the-art English-only parser.

1 Introduction

Abstract Meaning Representation (AMR) (Banasescu et al., 2013) is a broad-coverage semantic formalism that encodes the meaning of a sentence as a rooted, directed, and labeled graph, where nodes represent concepts and edges represent relations among concepts. AMR parsing is the task of translating natural language sentences into their corresponding AMR graphs, which encompasses a set of natural language understanding tasks, such as named entity recognition, semantic role labeling, and coreference resolution. AMR has proved to be beneficial to a wide range of applications such as text summarization (Liao et al., 2018), machine translation (Song et al., 2019), and question answering (Kapanipathi et al., 2020; Xu et al., 2021).

One most critical feature of the AMR formalism is that it abstracts away from syntactic realization and surface forms. As shown in Figure 1, different

*This work was supported by Alibaba Group through the Alibaba Innovative Research (AIR) Program.


English

The boy wants the girl to believe him.

The boy's desire is for the girl to believe him.

The boy wants to be believed by the girl.

AMR



German, Spanish, Italian, and Chinese

Der Junge möchte, dass das Mädchen ihm glaubt.

El chico quiere que la chica le crea.

Il ragazzo vuole che la ragazza gli creda.

男孩要女孩相信他。

Figure 1: An example of AMR. Sentences written in English and other languages share the same meaning and therefore correspond to the same AMR graph.

English sentences with the same meaning correspond to the same AMR graph. Furthermore, there are no explicit alignments between elements (nodes or edges) in the graph and words in the text. While this property leads to a distinct difficulty in AMR parsing, it also suggests the potential of AMR to work as an interlingua (Xue et al., 2014; Hajič et al., 2014; Damonte and Cohen, 2018), which could be useful to multilingual applications of natural language understanding (Liang et al., 2020; Hu et al., 2020). An example is given in Figure 1, we represent the semantics of semantically-equivalent sentences in other languages using the same AMR graph. This defines the multilingual AMR parsing problem we seek to address in this paper.

Multilingual AMR parsing is an extremely challenging task due to several reasons. First, AMR was initially designed for and heavily biased towards English, thus the parsing has to overcome some structural linguistic divergences among languages (Damonte and Cohen, 2018; Zhu et al.,

2019). Second, the human-annotated resources for training are only available in English and none is present in other languages. Moreover, since the AMR graph involves rich semantic labels, the AMR annotation for other languages can be labor-intensive and unaffordable. Third, current modeling techniques focus mostly on English. For example, existing AMR aligners (Flanigan et al., 2014; Pourdamghani et al., 2014; Liu et al., 2018) and widely-used pointer-generator mechanisms (Zhang et al., 2019b; Cai and Lam, 2019, 2020) rely on the textual overlap between English words and AMR node values (i.e., concepts).

Some initial attempts (Damonte and Cohen, 2018; Biloshmi et al., 2020; Sheth et al., 2021) towards multilingual AMR parsing mainly investigated the construction of pseudo parallel data via annotation projection. In this paper, we study multilingual AMR parsing from the perspective of knowledge distillation (Buciluă et al., 2006; Ba and Caruana, 2014; Hinton et al., 2015; Kim and Rush, 2016), where our primary goal is to improve a multilingual AMR parser by using an existing English parser as its teacher. We focus on a strict multilingual setting for developing one AMR parser that can parse all different languages. In contrast to the language-specific (one parser one language) setting, our setting is more challenging yet more appealing in practice. Intuitively, knowledge distillation is effective because the teacher’s output provides a rich training signal for the student parser. We develop both the teacher parser and the student parser with language-agnostic *seq2seq* design and expect the student parser to imitate the behaviors of the teacher parser (i.e., English parser) when processing semantically-equivalent input in other languages. We first show that multilingual *seq2seq* pre-training, including language model and machine translation pre-training, provides an excellent starting point for model generalization across languages. We further capitalize on the idea that the student should be robust to noisy input and introduce noise by machine translation for improving student performance. To migrate the risk that the student learns the mistakes made by the teacher, the student is then fine-tuned with gold AMR graphs.

We present experiments on the benchmark dataset created by Damonte and Cohen (2018), covering four different languages with no training data, including German, Spanish, Italian, and Chinese. To cover as many languages as possible, we also

include the original English test set in our evaluation. On four zero-resource languages, our single universal parser consistently outperforms the previous best results by large margins (+11.3 SMATCH points on average and up to +18.8 SMATCH points). Meanwhile, our parser achieves competitive results on English even compared with the latest state-of-the-art English AMR parser in the literature.

To sum up, our contributions are listed below:

- We study AMR parsing in a strict multilingual setting, there is but one parser for all different languages including English.
- We propose to train a multilingual AMR parser with multiple pre-training and fine-tuning stages including noisy knowledge distillation.
- We obtain a performant multilingual AMR parser, establishing new state-of-the-art results on multiple languages. We hope our parser can facilitate the multilingual applications of AMR.

2 Background

2.1 Prior Work

Cross-lingual AMR parsing is the task of mapping a sentence in any language X to the AMR graph of its English translation. To date, there is no human-annotated X -AMR parallel dataset for training. Therefore, one straightforward solution is to translate the sentences from X into English then apply an English parser (Damonte and Cohen, 2018; Uhrig et al., 2021). However, it is argued that the method is not informative in terms of the cross-lingual properties of AMR (Damonte and Cohen, 2018; Biloshmi et al., 2020). To tackle cross-lingual AMR parsing, most previous work relies on pre-trained multilingual language models and silver training data (i.e., pseudo parallel data).

Pre-trained Multilingual Language Model

Previous work proves that language-independent features provided by pre-trained multilingual language models can boost cross-lingual parsing performance. For example, Biloshmi et al. (2020) use mBERT (Devlin et al., 2019) and Sheth et al. (2021) employ XLM-R (Conneau et al., 2020).

Silver Training Data There are two typical methods for creating silver training examples: (I) *Parsing English to AMR* (Damonte and Cohen,

2018). This approach creates silver training examples for the foreign language X through an external X -EN parallel corpus and an existing English AMR parser. The English sentences of the parallel corpus are parsed using the existing AMR parser. Then resultant AMR graphs are used as pseudo targets. Note that the target side of the constructed X -AMR training corpus is of silver quality. (II) *Translating English to X* (Biloshmi et al., 2020; Sheth et al., 2021). This approach does not exploit external X -EN parallel corpus but makes use of the existing EN-AMR parallel corpus. It uses off-the-shelf machine translation systems to translate the English side of the EN-AMR pairs into the foreign language X . It is worth noting that although the source side of the constructed training examples may contain noise introduced by automatic translation, the target side consists of gold AMR graphs. However, the number of training examples created by this approach is limited by the size of the original English dataset.

2.2 Our Task: Multilingual AMR Parsing

Here, we formally define the task of multilingual AMR parsing. As illustrated in Figure 1, this task aims to predict the semantic graph given the input sentence in any language. Specifically, we consider five different languages: German (DE), Spanish (ES), Italian (IS), Chinese (ZH), and English (EN). The biggest challenge is due to the only access to a set of human-annotated English training examples. Formally, denote $\mathbb{Z} = \{\text{DE}, \text{ES}, \text{IT}, \text{ZH}\}$ as the set of *foreign languages* other than EN. Our goal is to develop a multilingual parser for all languages in $\{\text{EN}\} \cup \mathbb{Z}$. However, there only exists a set of gold EN-AMR training pairs (x, y) where x and y are the English sentence and AMR graph, respectively. For any language $X \in \mathbb{Z}$, there is *no* gold training example.

Following the recent state-of-the-art practice for English AMR parsing (Bevilacqua et al., 2021), we formulate the problem as a *seq2seq* task. The input sentence serves as the source sequence, while the linearization of the AMR graph is treated as the target sequence.

3 Our Parser

3.1 Overview

We choose vanilla *seq2seq* architecture (Vaswani et al., 2017; Bevilacqua et al., 2021) for our multilingual AMR parser to dispose of the need of

explicit word-to-node alignments. Unlike Damonte and Cohen (2018); Sheth et al. (2021), the advantage of alignment-free parsers is that the training is prevented from depending on noisy alignments derived from automatic cross-lingual aligners.

The training of our parser consists of multiple pre-training and fine-tuning stages. First, we initialize both the encoder and decoder of our parser using parameters pre-trained for multilingual denoising autoencoding and multilingual machine translation. We argue that both pre-training stages boost model generalization across languages and the latter is especially beneficial to AMR parsing because translating to a meaning representation resembles machine translation. Then, we fine-tune our parser in two stages. In the first stage, we aim to transfer the knowledge of a high-performing English AMR parser to our multilingual parser via knowledge distillation. Finally, we fine-tune our parser with gold AMR graphs to alleviate the drawback of over-fitting to teacher’s mistakes. Each training stage is detailed in §3.3 and its individual effect is empirically revealed in §5.1.

3.2 Base Model

Model Architecture We consider the standard *Transformer* (Vaswani et al., 2017) for *seq2seq* modeling. The encoder in the Transformer consists of a stack of multiple identical layers, each of which has two sub-layers: one implements the multi-head self-attention mechanism and the other is a position-wise fully connected feed-forward network. The decoder is also composed of a stack of multiple identical layers. Each layer in the decoder consists of the same sub-layers as in the encoder layers plus an additional sub-layer that performs multi-head attention to the output of the encoder stack. See Vaswani et al. (2017) for more details.

Linearization & Post-processing To formulate AMR parsing as a *seq2seq* problem, one needs to first obtain the linearized sequence representation of AMR graphs. To this end, we adopt the fully graph-isomorphic linearization techniques as in Bevilacqua et al. (2021). That is, the graph is recoverable from the linearized sequence without losing adjacency information. We use special tokens $\langle v_0 \rangle, \langle v_1 \rangle, \dots, \langle v_n \rangle$ to represent variables in the linearized graph and to handle co-referring nodes. We make a clear distinction between constants and variables, as variable names do not carry any semantic information. The graph

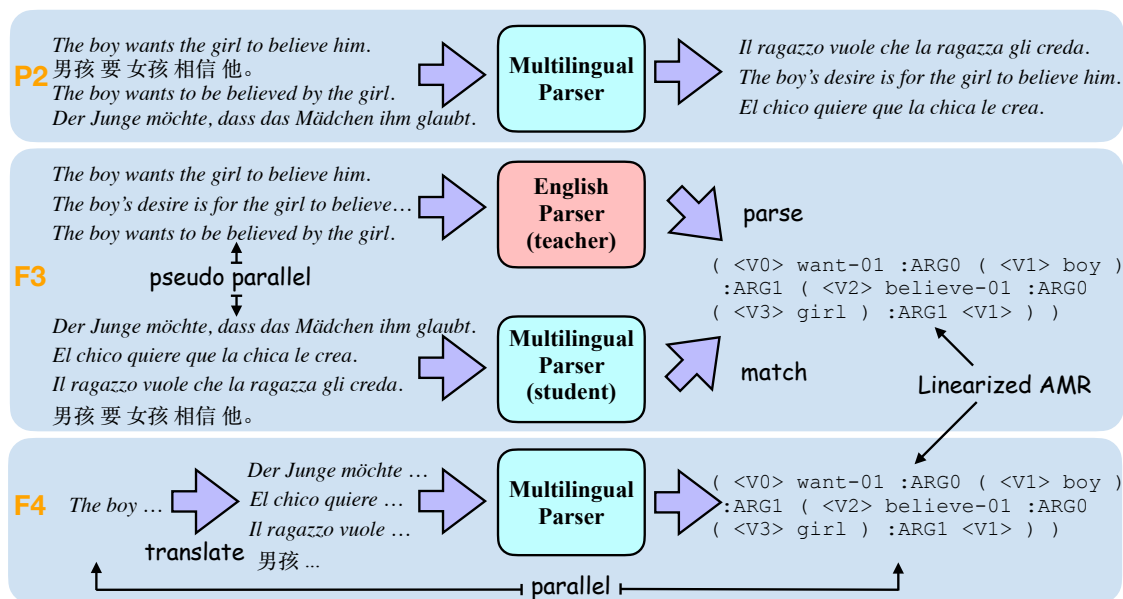


Figure 2: Illustration of different training stages. Stage P1 is omitted for space limit.

is linearized through a depth-first traversal starting from the root. For edge ordering, we use the default order in the release files of AMR datasets as suggested by Konstas et al. (2017). The bottom right of Figure 2 illustrates the linearization result of the AMR graph in Figure 1.

The output sequence of our *seq2seq* model may produce an invalid graph. For example, the parenthesis parity may be broken, resulting in an incomplete graph. To ensure the validity of the graph produced in parsing, post-processing steps such as parenthesis parity restoration and invalid segment removal are introduced. We use the pre- and post-processing scripts provided by Bevilacqua et al. (2021).¹

3.3 Training Stages

We now clarify the four different training stages. The whole training process is referred to as $P1 \rightarrow P2 \rightarrow F3 \rightarrow F4$.

P1: Multilingual Language Model Pre-training

Pre-trained multilingual language representations such as mBERT (Devlin et al., 2019) have greatly improved performance across many cross-lingual language understanding tasks. For cross-lingual AMR parsing, in particular, Biloshmi et al. (2020) used mBERT² (Devlin et al., 2019) while Sheth et al. (2021) employed XLM-R³ (Conneau et al.,

2020) to provide language-independent features. Unlike previous work, we argue that such encoder-only pre-trained models are not the most suitable choice for our *seq2seq* parser. Instead, we adopt mBART, an encoder-decoder denoising language model pre-trained with monolingual corpora in many languages (Liu et al., 2020b), to initialize both the encoder and decoder of our *seq2seq* parser.

P2: Multilingual Machine Translation Pre-training (MMT-PT)

The task of multilingual machine translation (MMT) is to learn one single model to translate between various language pairs. Essentially, natural languages can be considered as *informal* meaning representations compared to *formal* meaning representation such as AMR. On the other hand, AMR can be regarded as a special language. The above observations connect the dots between MMT and multilingual AMR parsing, both of which model the process of digesting the semantics in one form and conveying the same semantics in another form. Therefore, we argue that pre-training our parser using the MMT task should be helpful. In fact, the usefulness of MT pre-training has also been validated in English AMR parsing (Xu et al., 2020). In practice, we directly use the mBARTmmt checkpoint (Tang et al., 2020), an MMT model covering 50 languages that are trained from mBART.

F3: Knowledge Distillation Fine-tuning (KD-FT)

Motivated by the fact that the parsing accuracy on English is significantly better than those

¹<https://github.com/SapienzaNLP/spring>

²bert-base-multilingual-cased

³xlm-roberta-large

on other languages, we propose to reduce the performance gap via knowledge distillation (Kim and Rush, 2016). Specifically, we first pre-train a high-performance AMR parser for English and treat it as the teacher model. By considering our multilingual AMR parser as the student model, the goal is to transfer the knowledge of the teacher model to the student model. Intuitively, when feeding an English sentence to the teacher model, the student model, which receives its translation as input, should imitate the behaviors of the teacher model and make similar predictions. The strategies we adopt to achieve this goal are detailed in §3.4.

F4: Gold AMR Graph Fine-tuning (Gold-FT)

Note that in the knowledge distillation stage, the student parser is trained to match the predictions of the teacher model. A potential risk of such knowledge distillation is that the mistakes made by the teacher model may be propagated to the student model as well. Another fine-tuning stage, which we found useful for alleviating the risk, is to further fine-tune our parser with gold AMR graphs. Following Blloshmi et al. (2020); Sheth et al. (2021), we transform the gold standard English datasets into other languages using MT models for fine-tuning our multilingual AMR parser.

3.4 Knowledge Distillation

Knowledge distillation (KD) refers to a class of techniques for training a *student* model to imitate a *teacher* model for close or even better performance. In contrast to most KD applications that focus on reducing the performance gap caused by architectural differences, our primary goal is to minimize the mismatch of model behaviors across languages. That is, we expect the student and teacher to behave similarly even with *different* input languages.

Recall that we formulate AMR parsing as a *seq2seq* problem with standard maximum likelihood estimation training objective.

$$L_{\text{MLE}} = \sum_{t=1}^{|y|} \log p(y_t | y_{:<t}, x) \quad (1)$$

where y_t denotes the t -th token in the linearized AMR sequence. One natural and common method for KD is to replace the discrete target with the soft token-level distributions provided by the teacher model $p_T(y_t | y_{:<t}, x^*)$.

$$L_{\text{token}} = \sum_{t=1}^{|y|} \text{KL}((p(y_t | y_{:<t}, x), p_T(y_t | y_{:<t}, x^*)))$$

where KL computes the Kullback–Leibler divergence between two distributions. We use x^* and x to highlight that the input sentences are in different languages. The above method is referred to as *token-level* KD as it attempts to match the local token distributions of the teacher model. Opposed to token-level KD, *sequence-level* KD (Kim and Rush, 2016) allows knowledge transfer at sequence-level $L_{\text{seq}} = \text{KL}(p(y, x), p_T(y, x^*))$. Due to the intractability of sequence-level distribution computation, following Kim and Rush (2016), we replace the teacher’s distribution with its mode. Specifically, we use beam search to approximate the teacher’s most probable output, which is then used as the target to train the student model as in Eq. 1.

One appealing property of sequence-level KD is that it does not require gold AMR graphs. Therefore, it can be performed with an external X-EN parallel corpus at scale. However, the inherent noise in the teacher’s output hampers training with the student often being prone to hallucination (Liu et al., 2020a). To alleviate this problem, we propose to also inject noise to the input side of the student model. We find that automatic translation can serve as an effective noise generator for multilingual AMR parsing. That is, instead of using gold translations, we feed automatic machine translations to the student model. We find that the noise introduced by machine translation performs better than random noise likely due to that the translations preserve the most salient semantics.

4 Experimental Setup

4.1 Datasets

Gold Data Following conventions, we use the benchmark dataset created in Damonte and Cohen (2018) as our testbed. This dataset contains human translations of the test set of AMR2.0 dataset (LDC2017T10) in German (DE), Spanish (ES), Italian (IT), and Chinese (ZH). For a more complete multilingual setup, we also include the original English (EN) test set for evaluation. The gold training corpus in our experiments is the training set of AMR2.0, which contains 36, 521 EN-AMR pairs.

Silver Data For other foreign languages (DE, ES, IT, and ZH), we construct silver training data following Blloshmi et al. (2020). Specifically, we use OPUS-MT (Tiedemann and Thottingal, 2020)⁴,

⁴https://huggingface.co/transformers/model_doc/marian.html

an off-the-shelf translation tool, to translate English sentences in AMR2.0 to other foreign languages. To ensure the quality of silver data, we filter out data with less accurate translations via back-translation consistency check. That is, the translation quality is measured by the cosine similarity between the original English sentence and its back-translated counterpart using LASER (Artetxe and Schwenk, 2019). We refer readers to Blloshmi et al. (2020) for an exhaustive description of the data filtering process. Detailed statistics of our training, dev, and test sets are shown in Table 1.

Knowledge Distillation Data For the knowledge distillation stage, we use 320K English sentences in the Europarl corpus (Koehn, 2005), which contains parallel sentence pairs of $En \leftrightarrow DE$, $En \leftrightarrow ES$, and $En \leftrightarrow IT$. Unless otherwise specified, we use sequence-level KD with noisy input from OPUS-MT. Note that essentially our noisy KD only requires monolingual English data. Nevertheless, we choose Europarl following Damonte and Cohen (2018); Blloshmi et al. (2020) and use the gold translations as noise-free input to demonstrate the impact of our noisy KD comparatively (§5.2).

4.2 Settings

We differentiate two settings for training and evaluating multilingual AMR parsing.

- **Language-specific.** For each target language, a language-specific parser is trained.
- **Multilingual.** One single parser is trained to parse all target languages.

While this paper focuses on the multilingual setting, we also report the results of the language-specific parsers in previous work (Damonte and Cohen, 2018; Blloshmi et al., 2020; Sheth et al., 2021) for comparative reference.

4.3 Models

Model Variants Our full training pipeline consists of multiple pre-training and fine-tuning stages. To study the effect of each training stage, we implement a series of model variants:

- **w/o MMT-PT.** To measure the help from MMT-PT, we remove the second pre-training stage (P2). The training process becomes $P1 \rightarrow F3 \rightarrow F4$.

Language	Train	Dev	Test
English(EN)	36,521*	1,368*	1,371*
German(DE)	34,415	1,319	1,371*
Spanish(ES)	34,552	1,325	1,371*
Italian(IT)	34,521	1,322	1,371*
Chinese(ZH)	33,221	1,311	1,371*

Table 1: The number of instances per language and for each data split. * marks gold quality and otherwise silver quality.

- **w/o KD-FT.** To show the benefits from KD, we conduct an ablation experiment where the KD-FT stage (F3) is skipped. The training process becomes $P1 \rightarrow P2 \rightarrow F4$.
- **w/o Gold-FT.** To validate the necessity of the fine-tuning with gold AMR graph, we also report the model results without the final Gold-FT (F4) stage. The training process is then $P1 \rightarrow P2 \rightarrow F3$.
- **w/o MMT-PT & KD-FT.** We exclude both the MMT-PT (P2) stage and the KD-FT (F3) stage. This variant ($P1 \rightarrow F4$) is reminiscent of the best-performing model of Blloshmi et al. (2020) that fine-tunes multilingual language model with silver training data.
- **w/o MMT-PT & Gold-FT.** We also report the model performance without MMT-PT and Gold-FT for reference ($P1 \rightarrow F3$).

Implementation Details Following Bevilacqua et al. (2021), we make slight modifications to the vocabulary of mBART for better suiting linearized AMRs. Specifically, we augment the original vocabulary of mBART with the names of AMR relations and frames occurring at least 5 times in the gold training corpus. The augmented vocabulary allows more compact target sequence after tokenization. As introduced in §3.3, the first two pre-training stages are out of scope for this paper and we directly load pre-trained model checkpoints, mBART⁵ (Liu et al., 2020b) and mBART-mmt⁶ (Tang et al., 2020), from Huggingface’s transformers library (Wolf et al., 2020). At each fine-tuning stage, models are trained for up to 30,000 steps with a batch size of 5,000 graph linearization tokens, with RAdam (Liu et al., 2019) optimizer and a learning rate of $1e-5$. Dropout is set to 0.25. We do model selection according

⁵facebook/mbart-large-50

⁶facebook/mbart-large-50-many-to-many-mmt

Model	SMATCH						
	DE	ES	IT	ZH	EN	AVG _X	AVG
Language-Specific							
Damonte and Cohen (2018)	39.0	42.0	43.0	35.0	-	39.8	-
Biloshmi et al. (2020)	53.0	58.0	58.1	43.1	-	53.1	-
Sheth et al. (2021) [†]	62.7	67.9	67.4	-	-	-	-
Multilingual							
Biloshmi et al. (2020) [†]	52.1	56.2	56.7	-	-	-	-
Biloshmi et al. (2020)	49.9	53.2	53.5	41.0	-	49.4	-
Ours (P1→P2→F3→F4)	73.1	75.9	75.4	61.9	83.9	71.6	74.0
w/o MMT-PT (P1→F3→F4)	72.4	75.6	75.4	60.6	83.3	71.0	73.5
w/o KD-FT (P1→P2→F4)	71.8	74.5	73.8	61.0	82.6	70.3	72.7
w/o Gold-FT (P1→P2→F3)	70.9	74.0	73.1	59.5	82.4	69.4	72.0
w/o MMT-PT & KD-FT (P1→F4)	70.8	73.8	73.2	59.9	81.8	69.4	71.9
w/o MMT-PT & Gold-FT (P1→F3)	70.0	73.3	72.7	58.4	81.4	68.6	71.2
State-of-the-art English-only Parser							
Bevilacqua et al. (2021) [‡]	-	-	-	-	84.3	-	-

Table 2: SMATCH scores on test sets. AVG_X and AVG denote the averages over zero-resource languages (DE, ES, IT, and ZH) and all languages respectively. † indicates that the results do not include ZH. ‡ marks that we report the best score without graph re-categorization considering our models do not use graph re-categorization either.⁷

to the performance on dev sets. At prediction time, we set beam size to 5. The teacher model is separately trained and obtains 84.2 SMATCH score on the English test set, which is close to the recent state-of-the-art result (Bevilacqua et al., 2021). We release our code, data, and models at <https://github.com/jcyk/XAMR>.

5 Experimental Results

The performance of AMR parsing is conventionally measured by SMATCH score (Cai and Knight, 2013), which quantifies the maximum overlap between two AMR graphs. The reported results are averaged over 3 runs with different random seeds.

5.1 Main Results

In Table 2, we present the SMATCH scores of our models and the best-performing models in the current literature. Our model with the full training pipeline achieves new state-of-the-art performances on all the four zero-resource languages, substantially outperforming all previous results. Concretely, the performance gains over the previous best results (Sheth et al., 2021) are 10.4, 8.0, 8.0, and 18.8 SMATCH points on German, Spanish, Italian, and Chinese respectively. This is even more remarkable given that the previous best results are achieved via a set of language-specific parsers, while ours are obtained by one single multilingual parser. Notably, our multilingual parser

also obtains close performance on English to that achieved by the state-of-the-art English-only parser. These results are encouraging for developing AMR parser in a strict multilingual setting (i.e., using one parser for all languages).

The results of our ablated model variants further reveal the source of performance gains. As seen, each of MMT-PT, KD-FT, and Gold-FT make indispensable contributions to the superior performance. Skipping any of them leads to a considerable performance drop and removing two further degrades the model performance. Concretely, the averaged SMATCH score across all languages (AVG) decreases by 0.5 points when removing MMT-PT, which confirms our hypothesis that MMT is a beneficial pre-training objective for multilingual AMR parsing. It is also observed that the AVG score drops down from 74.0 to 72.7 (−1.3 points) when skipping KD-FT. In other words, introducing KD-FT boosts the performance by 1.3 SMATCH points on average. The improvement is striking since **Ours w/o KD-FT** is already a very strong baseline (AVG=72.7). Lastly, by comparing the results of **Ours w/o Gold FT** and **Ours**, we can see that appending Gold-FT to the preceding training stages yields a growth of 2.0 AVG points. This demonstrates that KD-FT alone is not sufficient and fine-tuning with gold AMR graphs has a complementary effect. Another interesting finding is that even our worst-performing variant surpasses previous best methods, which validates that pre-trained encoder-decoder architecture, mBART, is more effective for multilingual AMR parsing than encoder-only pre-trained models used in prior work.

⁷Graph re-categorization is a popular technique for reducing the complexity of AMR graphs, which involves manual efforts for hand-crafting rules. Recent work (Bevilacqua et al., 2021) points out that graph re-categorization may harm the generalization ability to out-of-domain data.

Method	DE	ES	IT	ZH	EN	AVG
<i>tok</i>	71.8	75.1	74.0	60.9	82.7	72.9
<i>seq</i>	73.1	75.9	75.4	61.9	83.9	74.0
<i>tok + seq</i>	73.1	75.8	75.3	61.6	83.9	73.9
<i>seq*</i>	71.9	75.0	74.1	61.2	82.9	73.0

Table 3: Comparison of different KD methods.

5.2 Discussions

Now we delve into more discussions on our key innovation, i.e., the knowledge distillation stage.

Effect of Different Knowledge Distillation Methods As introduced in §3.4, there are two kinds of knowledge distillation (KD) methods for *seq2seq* tasks: token-level KD (*tok*) and sequence-level KD (*seq*). In Table 3, we compare *tok*, *seq*, and their combination (*seq+tok*). For *seq+tok*, we train the student on teacher-generated graphs but still use a token-level KL term between the teacher/student. Note that *tok* can only utilize data with gold AMR graphs (i.e., the constructed silver training data), while *seq* and *seq+tok* leverage additional English sentences. Therefore, we also report the result of *seq* using the same English sentences as *tok*, denoted as *seq**. As seen, *seq* performs much better than *tok* and their combination does not bring further improvement. However, *seq** only gives similar result to *tok*. These results show that training on more data is crucial and using *seq* alone is sufficient for knowledge transfer.

Effect of Noise for Knowledge Distillation

Next, we study the effect of noise during knowledge distillation. Recall that we use automatic machine translation to generate noisy input for the student model. To show that noise is an important ingredient for superior performance, we also conduct experiments where the reference translations in Europarl are used as noise-free input to the student. Also, to show that the noise from MT is non-trivial, we further employ BART-style random noise (Lewis et al., 2020) for comparison. BART-style noise masks text spans in the input and we tune the rate of word deletion. The results are presented in Table 4. We show that MT noise is indeed helpful and its role cannot be replaced by simple random noise.

Effect of Data Sizes for Knowledge Distillation

Lastly, we study the relation between model performance and the size of monolingual data used for KD. Figure 3 shows that the SMATCH scores

Noise	DE	ES	IT	ZH	EN	AVG
None	72.3	75.3	74.8	61.3	83.1	73.4
Word deletion						
10%	72.4	75.1	74.6	61.3	83.5	73.4
15%	72.4	75.1	74.7	61.6	83.3	73.4
20%	72.7	75.6	75.1	61.3	83.7	73.7
25%	72.5	75.2	74.3	61.1	83.3	73.3
30%	72.5	75.3	74.7	61.3	83.5	73.4
MT	73.1	75.9	75.4	61.9	83.9	74.0

Table 4: Comparison of different noise generators. Word deletion k%: randomly mask k% words.




Figure 3: Comparison of different data sizes for KD.

(AVG_x and AVG) grow approximately logarithmically with the data size for KD.

6 Related Work

Cross-lingual AMR Parsing AMR (Banarescu et al., 2013) is a semantic formalism initially designed for encoding the meanings of English sentences. Over the years, a number of preliminary studies have investigated the potential of AMR to work as an interlingua (Xue et al., 2014; Hajič et al., 2014; Anchieta and Pardo, 2018; Zhu et al., 2019). These works attempt to refine and align English AMR-like semantic graphs labeled in different languages. Damonte and Cohen (2018) show that it is possible to use the original AMR annotations devised for English as representation for equivalent sentences in other languages and release a cross-lingual AMR evaluation benchmark (Damonte and Cohen, 2020) very recently. Cross-lingual AMR parsing suffers severely from the data scarcity issue; there is no gold annotated training data for languages other than English. Damonte and Cohen (2018) propose to build silver training data based on external bitext resources and English AMR parser. Blloshmi et al. (2020) find that translating the source side of existing English AMR dataset into other target languages produces better silver training data. Sheth et al. (2021) focus on improving cross-lingual word-to-node alignment for training cross-lingual AMR parsers that

rely on explicit alignment. Our work follows the alignment-free *seq2seq* formulation (Barzdins and Gosko, 2016; Konstas et al., 2017; Van Noord and Bos, 2017; Peng et al., 2017; Zhang et al., 2019a; Ge et al., 2019; Bevilacqua et al., 2021) and we alternatively study this problem from the perspective of knowledge distillation, which provides a new way to enable multilingual AMR parsing.

Knowledge Distillation for Sequence Generation Knowledge distillation (KD) is a classic technique originally proposed for model compression (Buciluă et al., 2006; Ba and Caruana, 2014; Hinton et al., 2015). KD suggests training a (smaller) student model to mimic a (larger) teacher model, by minimizing the loss (typically cross-entropy) between the teacher/student predictions (Romero et al., 2015; Yim et al., 2017; Zagoruyko and Komodakis, 2017). KD has been successfully applied to various natural language understanding tasks (Kuncoro et al., 2016; Hu et al., 2018; Sanh et al., 2019). For sequence generation tasks, Kim and Rush (2016) first introduce sequence-level KD, which aims to mimic the teacher’s actions at the sequence-level. KD has been proved useful in a range of sequence generation tasks such as machine translation (Freitag et al., 2017; Tan et al., 2019), non-autoregressive text generation (Gu et al., 2017; Zhou et al., 2019), and text summarization (Liu et al., 2020a). To the best of our knowledge, our paper is the first work to investigate the potential of knowledge distillation in the context of cross-lingual AMR parsing.

7 Conclusion

We presented a multilingual AMR parser that significantly advances the state-of-the-art parsing accuracies on multiple languages. Notably, the superior results are achieved with one single AMR parser. Our parser is trained with multiple pre-training and fine-tuning stages including a noisy knowledge distillation stage. We hope our work can facilitate the application of AMR in multilingual scenarios.

References

Rafael Anchiêta and Thiago Pardo. 2018. [Towards AMR-BR: A SemBank for Brazilian Portuguese language](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).

Mikel Artetxe and Holger Schwenk. 2019. [Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond](#). *Transactions of the Association for Computational Linguistics*, 7:597–610.

Lei Jimmy Ba and Rich Caruana. 2014. Do deep nets really need to be deep? In *Proceedings of the 27th International Conference on Neural Information Processing Systems-Volume 2*, pages 2654–2662.

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, Sofia, Bulgaria. Association for Computational Linguistics.

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, San Diego, California. Association for Computational Linguistics.

Michele Bevilacqua, Rexhina Blloshmi, and Roberto Navigli. 2021. One SPRING to rule them both: Symmetric AMR semantic parsing and generation without a complex pipeline. In *Proceedings of AAAI*.

Rexhina Blloshmi, Rocco Tripodi, and Roberto Navigli. 2020. [XL-AMR: Enabling cross-lingual AMR parsing with transfer learning techniques](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2487–2500, Online. Association for Computational Linguistics.

Cristian Buciluă, Rich Caruana, and Alexandru Niculescu-Mizil. 2006. Model compression. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 535–541.

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 3799–3809, Hong Kong, China. Association for Computational Linguistics.

Deng Cai and Wai Lam. 2020. [AMR parsing via graph-sequence iterative inference](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1290–1301, Online. Association for Computational Linguistics.

- 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.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Marco Damonte and Shay Cohen. 2020. Abstract meaning representation 2.0-four translations ldc2020t07. *Web Download, Philadelphia: Linguistic Data Consortium*.
- Marco Damonte and Shay B. Cohen. 2018. [Cross-lingual Abstract Meaning Representation parsing](#). 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 1146–1155, New Orleans, Louisiana. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational 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, Baltimore, Maryland. Association for Computational Linguistics.
- Markus Freitag, Yaser Al-Onaizan, and Baskaran Sankaran. 2017. Ensemble distillation for neural machine translation. *arXiv preprint arXiv:1702.01802*.
- DongLai Ge, Junhui Li, Muhua Zhu, and Shoushan Li. 2019. Modeling source syntax and semantics for neural amr parsing. In *IJCAI*, pages 4975–4981.
- Jiatao Gu, James Bradbury, Caiming Xiong, Victor OK Li, and Richard Socher. 2017. Non-autoregressive neural machine translation. *arXiv preprint arXiv:1711.02281*.
- Jan Hajič, Ondřej Bojar, and Zdeňka Urešová. 2014. [Comparing Czech and English AMRs](#). In *Proceedings of Workshop on Lexical and Grammatical Resources for Language Processing*, pages 55–64, Dublin, Ireland. Association for Computational Linguistics and Dublin City University.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. Xtreme: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation. In *International Conference on Machine Learning*, pages 4411–4421.
- Minghao Hu, Yuxing Peng, Furu Wei, Zhen Huang, Dongsheng Li, Nan Yang, and Ming Zhou. 2018. [Attention-guided answer distillation for machine reading comprehension](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2077–2086, Brussels, Belgium. Association for Computational Linguistics.
- Pavan Kapanipathi, Ibrahim Abdelaziz, Srinivas Ravishankar, Salim Roukos, Alexander Gray, Ramon Astudillo, Maria Chang, Cristina Cornelio, Saswati Dana, Achille Fokoue, et al. 2020. Question answering over knowledge bases by leveraging semantic parsing and neuro-symbolic reasoning. *arXiv preprint arXiv:2012.01707*.
- Yoon Kim and Alexander M. Rush. 2016. [Sequence-level knowledge distillation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1317–1327, Austin, Texas. Association for Computational Linguistics.
- Philipp Koehn. 2005. Europarl: A parallel corpus for statistical machine translation. In *MT summit*, volume 5, pages 79–86. Citeseer.
- 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.
- Adhiguna Kuncoro, Miguel Ballesteros, Lingpeng Kong, Chris Dyer, and Noah A. Smith. 2016. [Distilling an ensemble of greedy dependency parsers into one MST parser](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1744–1753, Austin, Texas. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation,](#)

- and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fenfei Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, et al. 2020. Xglue: A new benchmark dataset for cross-lingual pre-training, understanding and generation. *arXiv preprint arXiv:2004.01401*.
- Keixin 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.
- Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. 2019. On the variance of the adaptive learning rate and beyond. In *International Conference on Learning Representations*.
- Yang Liu, Sheng Shen, and Mirella Lapata. 2020a. Noisy self-knowledge distillation for text summarization. *arXiv preprint arXiv:2009.07032*.
- Yijia Liu, Wanxiang Che, Bo Zheng, Bing Qin, and Ting Liu. 2018. **An AMR aligner tuned by transition-based parser**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2422–2430, Brussels, Belgium. Association for Computational Linguistics.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020b. **Multilingual denoising pre-training for neural machine translation**. *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Xiaochang Peng, Chuan Wang, Daniel Gildea, and Ni-anwen Xue. 2017. **Addressing the data sparsity issue in neural AMR parsing**. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 366–375, Valencia, Spain. Association for Computational Linguistics.
- Nima Pourdamghani, Yang Gao, Ulf Hermjakob, and Kevin Knight. 2014. **Aligning English strings with Abstract Meaning Representation graphs**. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 425–429, Doha, Qatar. Association for Computational Linguistics.
- Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. 2015. **Fitnets: Hints for thin deep nets**. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Janaki Sheth, Young-Suk Lee, Ramon Fernandez Astudillo, Tahira Naseem, Radu Florian, Salim Roukos, and Todd Ward. 2021. Bootstrapping multilingual amr with contextual word alignments. *arXiv preprint arXiv:2102.02189*.
- 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.
- Xu Tan, Yi Ren, Di He, Tao Qin, Zhou Zhao, and Tie-Yan Liu. 2019. Multilingual neural machine translation with knowledge distillation. *arXiv preprint arXiv:1902.10461*.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. 2020. Multilingual translation with extensible multilingual pretraining and finetuning. *arXiv preprint arXiv:2008.00401*.
- Jörg Tiedemann and Santhosh Thottingal. 2020. **OPUS-MT – building open translation services for the world**. In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, pages 479–480, Lisboa, Portugal. European Association for Machine Translation.
- Sarah Uhrig, Yoalli Rezepka Garcia, Juri Opitz, and Anette Frank. 2021. Translate, then parse! a strong baseline for cross-lingual amr parsing. *arXiv preprint arXiv:2106.04565*.
- Rik Van Noord and Johan Bos. 2017. Neural semantic parsing by character-based translation: Experiments with abstract meaning representations. *arXiv preprint arXiv:1705.09980*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 6000–6010.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. **Transformers: State-of-the-art natural language processing**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

- Dongqin Xu, Junhui Li, Muhua Zhu, Min Zhang, and Guodong Zhou. 2020. [Improving AMR parsing with sequence-to-sequence pre-training](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2501–2511, Online. Association for Computational Linguistics.
- Weiwen Xu, Huihui Zhang, Deng Cai, and Wai Lam. 2021. [Dynamic semantic graph construction and reasoning for explainable multi-hop science question answering](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1044–1056, Online. Association for Computational Linguistics.
- Nianwen Xue, Ondřej Bojar, Jan Hajič, Martha Palmer, Zdeňka Urešová, and Xiuhong Zhang. 2014. [Not an interlingua, but close: Comparison of English AMRs to Chinese and Czech](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 1765–1772, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Junho Yim, Donggyu Joo, Jihoon Bae, and Junmo Kim. 2017. [A gift from knowledge distillation: Fast optimization, network minimization and transfer learning](#). In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4133–4141.
- Sergey Zagoruyko and Nikos Komodakis. 2017. [Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Sheng Zhang, Xutai Ma, Kevin Duh, and Benjamin Van Durme. 2019a. [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.
- Sheng Zhang, Xutai Ma, Kevin Duh, and Benjamin Van Durme. 2019b. [Broad-coverage semantic parsing as transduction](#). 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 3786–3798, Hong Kong, China. Association for Computational Linguistics.
- Chunting Zhou, Graham Neubig, and Jiatao Gu. 2019. [Understanding knowledge distillation in non-autoregressive machine translation](#). *arXiv preprint arXiv:1911.02727*.
- Huaiyu Zhu, Yunyao Li, and Laura Chiticariu. 2019. [Towards universal semantic representation](#). In *Proceedings of the First International Workshop on Designing Meaning Representations*, pages 177–181, Florence, Italy. Association for Computational Linguistics.