

SETUP: Sentence-level English-To-Uniform Meaning Representation Parser

Emma Markle, Javier Gutierrez Bach, Shira Wein

Amherst College

{emarkle26, jgutierrezbach26, swein}@amherst.edu

Abstract

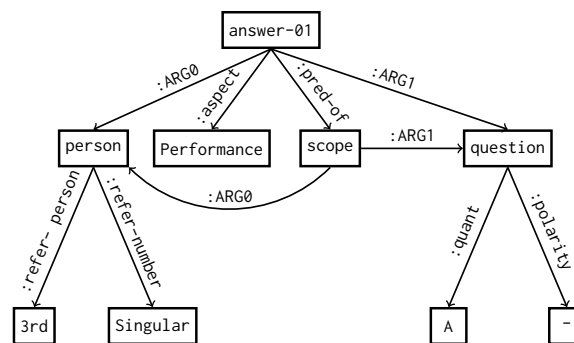
Uniform Meaning Representation (UMR) is a novel graph-based semantic representation which captures the core meaning of a text, with flexibility incorporated into the annotation schema such that the breadth of the world’s languages can be annotated (including low-resource languages). While UMR shows promise in enabling language documentation, improving low-resource language technologies, and adding interpretability, the downstream applications of UMR can only be fully explored when text-to-UMR parsers enable the automatic large-scale production of accurate UMR graphs at test time. Prior work on text-to-UMR parsing is limited to date. In this paper, we introduce two methods for English text-to-UMR parsing, one of which fine-tunes existing parsers for Abstract Meaning Representation and the other, which leverages a converter from Universal Dependencies, using prior work as a baseline. Our best-performing model, which we call SETUP, achieves an AnCast score of 84 and a SMATCH++ score of 91, indicating substantial gains towards automatic UMR parsing.

Keywords: uniform meaning representation, computational semantics, semantic parsing

1. Introduction

Uniform Meaning Representation (UMR; [Van Gysel et al., 2021](#)) is a graph-based semantic framework designed to capture the meaning of text across various languages, accounting for linguistic diversity during the annotation process. An example UMR graph can be seen in Figure 1. UMR builds on the foundations of Abstract Meaning Representation (AMR; [Banarescu et al., 2013](#)), extending its scope to handle multilingual and document-level annotation. Unlike AMR, which was originally developed for English and later adapted cross-lingually ([Wein and Schneider, 2024](#); [Wein, 2024](#)), UMR is designed as a multilingual semantic framework with a unified annotation schema. Its lattice-based annotation design and Stage 0 procedure further make it well-suited for low- and no-resource languages, offering stronger cross-lingual consistency and broader applicability than AMR. UMR additionally encodes aspect, modality, scope, and document-level relations.

AMR has been widely adopted in the NLP community, largely due to the extensive efforts made in developing effective text-to-AMR parsing and AMR-to-text generation models. This work has enabled AMR to be successfully applied in a variety of downstream tasks, including text summarization, machine translation, and information extraction ([Wein and Opitz, 2024](#)). While UMR shows initial promise in improving the performance of downstream applications for low-resource languages ([Wein, 2025](#)), developing text-to-UMR parsers is crucial to exploring the full scope of UMR benefits in multilingual and low-resource settings. Without reliable parsers, UMR’s richer representations cannot be



(a/ answer-01
:ARG0 (p/ person
:refer-person 3rd
:refer-number Singular)
:ARG1 (q/ question
:quant A
:polarity -)
:pred-of (s/ scope
:ARG0 p
:ARG1 q)
:aspect Performance)

Figure 1: UMR graph for the sentence “Someone didn’t answer all the questions” as a graph and in PENMAN notation ([Kasper, 1989](#)), adapted from the UMR Guidelines ([UMR, 2022](#)).

extracted or leveraged in such applications, thereby limiting the utility of UMR. By building text-to-UMR parsers in English, where annotated data and pre-existing AMR models are available, we establish a foundation for transferring these models and approaches to low-resource languages. Additionally, high-quality English AMR parsers are critical for

multilingual or cross-lingual tasks leveraging UMR.

In this paper, we lay the groundwork for expanding the utility of UMR in downstream NLP applications. Specifically, in this work we focus on **sentence-level English UMR parsing**. We constrain our task due to the enormous amount of linguistic information captured in UMR graphs and the complicated nature of this task. Working with English provides a valuable starting point for the more challenging task of low-resource language UMR parsing, where annotated data and pre-existing text-to-AMR parsers (e.g., for Navajo or other indigenous languages) are largely unavailable. By establishing methods and baselines in English, we can better develop techniques that are likely to transfer to low-resource settings in future work.

We explore two approaches to creating text-to-UMR parsers: (1) fine-tuning existing text-to-AMR parsers using UMR data, and (2) converting Universal Dependency (UD) trees into partial UMR graphs following (Gamba et al., 2025), then training a T5 model (Raffel et al., 2023) to expand these partial graphs into complete UMRs. Our contributions include:

1. A fine-grained analysis of the baseline pipeline’s performance (Chun and Xue, 2024) across the data in UMR v1.0 and the newer UMR dataset, UMR v2.0.
2. Five English sentence-level text-to-UMR parser models fine-tuned from existing text-to-AMR architectures.
3. A fine-tuned model that converts partial UMRs to complete UMRs, employing the UD to UMR bootstrap approach by Gamba et al. (2025).

Our best performing model is our fine-tuned text-to-AMR parsers, which we call SETUP: Sentence-level English-to-UMR Parser.¹

2. Background & Related Work

UMR builds upon AMR, a semantic graph framework developed for English and subsequently extended to multiple languages (Wein and Schneider, 2024). Since AMR was designed for English, each language variation of AMR has to account for language-specific features, whereas UMR’s annotation framework is designed to accommodate multilingual representations.

Furthermore, UMR allows for flexible annotation and consistency across languages by employing a lattice-like annotation structure (Van Gysel et al., 2019). This annotation process permits annotators to apply either coarse- or fine-grained annotations,

depending on the language’s features. Additionally, “Stage 0” is a part of the annotation process that allows annotators to establish predicate–argument structures and construct rolesets. This stage is employed for languages without existing rolesets, with particular emphasis on low- and no-resource languages (Vigus et al., 2020). Wein and Bonn (2023) compare the multilingual capabilities of UMR to the cross-lingual adaptations of AMR and find that UMR more effectively captures language-specific phenomena across languages, handling a wider range of morphosyntactic and semantic constructions while maintaining a consistent, unified representation.

While both AMR and UMR capture meaning (Van Gysel et al., 2019), UMR extends AMR in several important ways: it represents tense, aspect, modality, and scope at the sentence level; incorporates document-level information such as temporal and modal dependencies; and facilitates the annotation of coreferential relations. Recent work by Bonn et al. (2023) highlights the key differences between AMR and UMR structures, serving as a resource to provide concrete mappings of the roles and concepts in AMR to those in UMR.

Universal Dependencies (UD; de Marneffe et al., 2021) provide a cross-linguistically consistent syntactic framework that represents grammatical relations in a dependency-tree format. Recent work has leveraged UD trees to bootstrap partial UMR graphs (Gamba et al., 2025), demonstrating that UD-informed approaches can reduce time spent on annotation and improve structural consistency across languages. The design of UD prioritizes cross-lingual comparability, making it a valuable structural foundation that can be utilized in UMR parsing.

Parsing techniques for other semantic representations have also been studied to various degrees. Apart from UMR and AMR, other prominent meaning representation frameworks include Universal Conceptual Cognitive Annotation, Semantic Dependencies, and Discourse Representation Structures. Universal Conceptual Cognitive Annotation (UCCA) Abend and Rappoport (2013) is a compositional-tree-based representation, where the leaves are words and the nodes at other layers represent events and participants. The flexibility of UCCA stems from its multilayered structure allows for multi-sentence and multilingual annotation (Sadeddine et al., 2024). Parsing UCCA can be difficult due to reentrancy resulting in DAG structures, discontinuous structures and non-terminal nodes. TUPA (Hershcovich et al., 2017) handles these issues with a transition-based neural parser by constructing the UCCA graph incrementally using a BiLSTM encoder. More recent approaches to UCCA parsing include Samuel and Straka (2020),

¹Our code and checkpoints are available at https://github.com/ACNLPlab/UMR_Parsing.git

a sequence-to-sequence model that jointly predicts nodes and edges to produce UCCA graphs within a cross-framework architecture. Semantic Dependencies (SDs) is an earlier family of meaning representation with the aim of going integrating semantic structure into syntactic dependency parsing (Oepen et al., 2014). Wang et al. (2021) propose ACE, a neural structured framework that learns to combine multiple pretrained embeddings optimally. The tool applies the combinations in a graph-based parser to obtain strong performance on semantic dependency parsing. Finally, Discourse Representation Structures (DRS) integrate first order logic into meaning representations (Kamp, 1984). It is recursively structured in the form of nested boxes, equipped to encode events, semantic roles, discourse relations, and multiple sentences in a compositional framework. Logical operators are used on the boxes to signal scope. While earlier parsing again use transition-based methods (Evang, 2019), while more recent approaches use character-level neural representations and exhibit performance gain in comparison to word-based encoders for DRS parsing van Noord et al. (2020).

Given that UMR is a multilingual extension of AMR, prior work on AMR parsing serves as an inspiration for UMR parsing and as relevant context. AMR parsing has seen significant advances, driven largely by the release of large-scale AMR corpora such as AMR v2.0 (Knight et al., 2017) and AMR v3.0 (Knight et al., 2020). Recent systems such as SPRING (Bevilacqua et al., 2021) and AMRBART (Bai et al., 2022) treat AMR parsing and generation as dual tasks using pre-trained language models like BART and T5 (Raffel et al., 2023). Models such as BiBL (Cheng et al., 2022) and LeakDistill (Vasylenko et al., 2023) further refine this approach by jointly learning alignments or distilling structural information into encoder representations. These high-performing AMR parsers serve as the foundation for extending semantic parsing to UMR, motivating our exploration of fine-tuning existing AMR parsers on UMR data to accelerate progress toward multilingual and document-level meaning representations.

Efforts towards UMR parsing have been limited to date. First, Buchholz et al. (2024) use documentation resources such as lexical databases and interlinear glossed texts to generate UMR subgraphs for low-resource languages. Using this method, they correctly capture the predicate-argument structures in 109 out of 115 sentences in Aráphaho. Next, given that mapping roles in AMR to those in UMR can be an arduous process as some individual roles can be mapped into various roles in UMR, Post et al. (2024) integrate animacy parsing and logic rules into a neural network to perform role conversion. Their hybrid model, com-

binning logic rules and a neural network, achieves a macro F-score of 63.5, and their model reliant on rules only achieves a weighted F-score of 76.1. Finally, a new UMR parsing tool follows a pipeline approach converts from text to AMR, and subsequently to UMR (Chun and Xue, 2024). This tool combines models that convert text into AMR and uses mapping rules plus UD trees of sentences to generate sentence-level UMRs. Their approach integrates different models specialized in modality, temporal dependencies, and coreference to insert document-level information into the UMR by leveraging alignments between the sentences and the AMRs. They report an AnCast score of 64.2 and a SMATCH score of 72.2 for their best AMR to UMR parser at the sentence level. For document-level parsing, they achieve a comprehensive AnCast score of 61.5.

Additional work on UMR tools has focused on the task of UMR-to-text generation (the reverse direction on text-to-UMR parsing), and begun leveraging UMR for downstream tasks. Markle et al. (2025) introduce the first approaches to generating text from multilingual UMR graphs, finding that fine-tuning existing AMR-to-text models yields the highest qualitative and quantitative results. Furthermore, Wein (2025) investigates UMR-based translation, highlighting how UMR’s cross-lingual abstraction enhances meaning preservation across typologically diverse languages, thereby improving translation quality for low-resource languages. Collectively, these studies highlight an increasing interest in leveraging UMR for multilingual NLP tasks and downstream applications, thereby underscoring the need for reliable UMR parsers.

3. Methods

We build on prior work (Chun and Xue, 2024), using it as a baseline model and evaluating its performance on the newly released UMR v2.0 dataset (Bonn et al., 2025), which will serve as a reference point for the following approaches. Next, we build on the recent UD-to-partial UMR converter (Gamba et al., 2025) to automatically transform dependency trees into partial UMR graphs. These partial UMR graphs capture core semantic information from the text, which we then use to train a model that completes them into fully detailed UMR representations. Finally, we adopt a similar approach to that used in recent UMR-to-text generation work (Markle et al., 2025), where we fine-tune existing text-to-AMR parser models on UMR data, adapting them to capture the specific structures and conventions of UMR graphs.

3.1. Baseline Approach

We use the existing UMR parser method which follows a text-to-AMR-to-UMR pipeline approach (as described in Section 2), as a baseline for English sentence-level UMR parsing (Chun and Xue, 2024). In this architecture, text-to-AMR parsers constitute the first stage of the pipeline, generating AMRs from raw sentences. These AMRs serve as intermediate semantic representations that are subsequently used as input to the UMR conversion stage. Specifically, we use the following text-to-AMR parsing models: AMRBART (Bai et al., 2022), LeakDistill (Vasylenko et al., 2023), amr-lib², SPRING (Bevilacqua et al., 2021), and BiBL (Cheng et al., 2022). Next, the resulting AMRs are passed to LEAMR (Blodgett and Schneider, 2021) to retrieve sentence-level AMR alignments, while UD graphs are generated automatically by the pipeline’s conversion code from the same sentences. Following the baseline approach of Chun and Xue (2024), these alignments and UD structures together serve as inputs to the rule-based UMR conversion component constituting the final stage of the pipeline.

We test this baseline on the original UMR v1.0 dataset, as is evaluated in Chun and Xue (2024), and also evaluate it on the new UMR v2.0 dataset. The format of the new dataset requires some debugging from the pipeline model, in particular with regard to tokenization issues.³

To get a more comprehensive baseline of current state-of-the-art UMR parsing technology, we integrate the neuro-symbolic approach to role conversion from AMR to UMR (Post et al., 2024) into the pipeline. In this process, AMR concepts are classified as animate or inanimate by the parser. Then, the model uses a combination of logical rules and a neural network along with this animacy decision to find a good UMR mapping from the original split role. For example, the `:source` AMR role may be mapped into `:source`, `:destination`, or `:goal` in UMR. The model’s predictions for each role are then integrated into the corresponding UMRs produced by the pipeline, yielding the final UMRs that incorporate the neuro-symbolic role conversion approach.

3.2. Fine-tuning Approaches

Next, we explore two fine-tuning approaches to achieve a text-to-UMR parsing model: (1) fine-

²amr-lib GitHub Repository

³For instance, the parser code maintains “<Architect>” as a single token, which is a common term in the English Minecraft data of UMR v2.0; however, the UD graph expresses this as three different tokens, and therefore, the baseline code is expanded to deal with these tokenization issues.

tuning existing text-to-AMR parsers on UMR data, and (2) converting UDs into partial UMRs following Gamba et al. (2025), then fine-tuning T5 (Raffel et al., 2023) to complete these partial graphs into full UMRs.

For the first approach, we directly fine-tune text-to-AMR parsers on UMR data for 10 epochs with a learning rate of 4e-5. The models we use include:

1. amr-lib: T5 model (Raffel et al., 2023) trained on AMR v3.0 (Knight et al., 2020)
2. SPRING: BART-based model that treats AMR parsing and generation as two complementary, bidirectional tasks, trained on AMR v2.0 (Knight et al., 2017) and AMR v3.0.
3. BiBL: SPRING-based model that aligns AMR graphs with their respective sentences to share learning across parsing and generation, trained on AMR v2.0 and v3.0.
4. LeakDistill: transformer-based model that uses structural adapters to capture graph information, building on word-to-node alignments trained on AMR v2.0 and v3.0.
5. AMRBART: BART-based model trained on AMR v2.0 and AMR v3.0.

In this approach, we fine-tune the existing text-to-AMR models with UMR data, enabling the models to adapt to the UMR structures while retaining the semantic knowledge acquired during AMR training.

Next, we leverage existing work from Gamba et al. (2025), who introduce a method for bootstrapping UMR graphs from UD trees. First, we generate UDs for the sentences in the UMR datasets using Stanford’s Stanza pipeline (Qi et al., 2020), reformat them into the CoNLL-U format, and then convert these UD trees into partial UMR graphs. Then, we train a T5 model on pairs of these partial graphs with their respective sentences and full gold reference graphs, with the objective of training the model to convert the partial graphs into complete UMR graphs.

3.3. Data

The release of UMR v1.0 (Bonn et al., 2024) spans across six languages: English, Chinese, Sanapaná, Aráphaho, Navajo, and Kukama, totaling 2,022 sentence-level UMRs. The English dataset includes fragmented transcripts of a spoken description of a silent film, as well as text from LORELEI news articles. Some examples of these include:

- “and he picks it up,”
- “A-nd then . . . it shifts”
- “or... or before.”

The Chinese data is comprised of sentences from Wikinews. Annotations in Aráphaho, Navajo, and

Model	AnCast	SMATCH	SMATCH++
amrlib	12.560	41.287	43.655
SPRING	34.087	47.966	49.130
BiBL	31.393	49.930	51.960
LeakDistill	19.588	35.912	39.086
AMRBART	47.876	69.289	69.190

Table 1: Overall performance of each AMR model, measured by F-scores, after fine-tuning on UMR v1.0 and comparing their predicted UMRs against the gold references.

Kukama all correspond to narrative texts, whereas the genre of the Sanapaná annotations is not stated (Bonn et al., 2024).

The release of UMR v2.0 (Bonn et al., 2025) has expanded upon the initial six languages, adding Czech and Latin, and now contains 210,237 sentence-level UMRs. However, most of these new UMRs are concentrated in two languages: English, with 29,912 graphs, and Czech, with 175,268. Notably, a large portion of the new English UMRs consists of annotated interactions between two participants, one designated as the “Builder” and the other as the “Architect,” while they play the video game Minecraft, which is a sandbox game where players can build structures. These UMRs capture both spoken dialogue and in-game actions, with a majority of the English sentences in the test set being Minecraft-related, such as:

- “[Builder picks up a purple block at X:1 Y:2 Z:2]”
- “<Architect> oops sorry, I meant behind :)”
- “[Builder puts down a orange block at X:1 Y:2 Z:-2]”
- “<Architect> We will build the ‘I’ portion now with orange blocks”

The other portion of these new English UMRs resembles more fluid spoken English text. Examples include:

- “Likewise when services had to shut down after a day because of snow, the media launched a full blown attack and everyday people were saying how disgusting it was.”
- “The percentage of lung cancer deaths among the workers at the West Groton, Mass., paper factory appears to be the highest for any asbestos workers studied in Western industrialized countries, he said.”
- “Where are those people going to go when the private health insurer disappears?”

3.4. Data Split Selection

In this subsection, we describe our process for preparing the training, development, and test sets,

with the goal of identifying the most effective data split for model training. This is particularly important given the composition of the datasets and their difference in makeup. As shown in Section 3.3, the datasets vary considerably: UMR v1.0 contains short, fragmented spoken dialogue, while UMR v2.0 includes both Minecraft-related interactions and longer, more complex standard English text, which tend to be more structured. Due to variations in style and domain-specific vocabulary, model generalization and performance may be affected. Thus, we create and test different data splits to determine which configuration produces the most generalizable model.

We begin with the UMR v1.0 corpus (Bonn et al., 2024), selecting 143 of the 209 English sentences that have corresponding UMR graphs. To prevent inadvertently rewarding the text-to-AMR models trained on overlapping sentences, we exclude 66 graphs that overlap with AMR sentences used to train the five selected text-to-AMR parser models.

Next, using the UMR v2.0 corpus (Bonn et al., 2025), which includes all data from UMR v1.0, we create three different splits (still excluding the 66 overlapping graphs) to evaluate the effect of different training subsets: (1) all UMR v2.0 data from all languages⁴, (2) all English data from UMR v2.0, and (3) English UMR v2.0 data excluding some of the repetitive Minecraft interactions. In this third split, we retain only 1,000 sentences starting with “[Builder”, as these typically produce highly repetitive action-descriptive sentences (examples shown in Section 3.3). We exclude some of these sentences to encourage the model to generalize beyond repetitive action patterns.

We select our data split by first fine-tuning each existing text-to-AMR architecture on the English UMR v1.0 data to establish baseline performance and select the currently best-performing AMR-to-text model on that small amount of UMR data. We then select the best-performing model, AMRBART (Table 1), and evaluate it on three additional splits derived from UMR v2.0. We find that AMRBART performs best on the full English UMR v2.0 split, suggesting that this configuration offers the most effective balance between data size and domain diversity.

Based on this finding, we adopt a split containing 22,938 training, 6,680 development, and 2,941 test sentences for our main experiments. Finally, we fine-tune all five existing text-to-AMR parsers, as well as our UD pipeline model, on this split; these constitute the models we evaluate in Section 4.2 and form the basis of our final experimental results.

⁴Although our parsing approaches focus exclusively on English, including data from other languages may help the model learn more general graph structures and relations that transfer across languages.

3.5. Metrics

Graph similarity between AMR and UMR graphs can be measured with a variety of standard metrics, including AnCast (Sun and Xue, 2024), SMATCH (Cai and Knight, 2013), and SMATCH++ (Opitz, 2023). AnCast employs an anchor broadcast alignment algorithm designed to circumvent local maxima pitfalls. AnCast decomposes graph comparison into concept F1, unlabeled relation F1, labeled relation F1, and weighted relation F1, providing interpretable scores over nodes and edges via an efficient alignment procedure. SMATCH measures overall graph similarity by computing the maximum F1 overlap of instance and relation triples under a variable alignment, yielding a single structural similarity score. SMATCH++ is an extension of SMATCH with standardized preprocessing, improved alignment guarantees, and fine-grained subgraph scoring, enabling more controlled and diagnostically informative evaluation of semantic graph predictions. We use all three of these evaluation metrics to ensure a robust understanding of parser performance.

4. Results

In this section, we evaluate the performance of our different UMR parsing approaches. We first summarize the performance of the baseline pipeline, noting that it performs worse on UMR v2.0 than on UMR v1.0, due to substantial differences in the text genre. We further analyze the pipeline’s performance across different splits within UMR v2.0, highlighting variations associated with data from different sources. We then report the high-scoring performance of the fine-tuning approach for both the AMR parser architectures trained on UMR data and for the UD-based approach. We evaluate our results using graph similarity metrics, including AnCast, SMATCH, and SMATCH++, and conduct a qualitative analysis of generated UMR graphs, which provides insight into how different approaches and models capture semantic and structural nuances.

4.1. Baseline

Based on the baseline model setup from Section 3.1 and the data split selected in Section 3.4, we run the pipeline approach as our baseline on all English sentences in the UMR v2.0 corpus. The highest AnCast score we obtain is 20.5 for SPRING, the highest SMATCH score is 35.6 for AMRBART and the highest SMATCH++ score is 34.3 for SPRING, as shown in Table 2. These UMR v2.0 scores are considerably lower than the 72.2 SMATCH score of the pipeline approach on UMR v1.0 as reported in Chun and Xue (2024).

Model	AnCast	SMATCH	SMATCH++
amrlib	18.027	33.422	33.832
SPRING	20.484	35.102	34.263
BiBL	17.808	32.236	31.586
LeakDistill	18.754	35.004	32.992
AMRBART	19.878	35.592	34.078

Table 2: For all English sentences from UMR v2.0, the overall performance of each model, measured by F-scores, after converting their generated AMRs to UMRs via the pipeline approach and comparing them against the gold UMRs.

We believe that such a difference in results is due to the stark difference between the data in the first and second UMR corpora. As seen in Section 3.3 UMR v1.0 corpus contains many sentences that are simply fragments of speech. In contrast, the second UMR corpus is completely dominated by the Minecraft sentences Section 3.3. The AMR parsers are not equipped to handle tags indicating dialogue or coordinates tokens, which limits their performance in these contexts.

To explain the discrepancy between the results, we conduct a fine-grained evaluation of the split by evaluating the scores on the Minecraft and non-Minecraft sentences separately. We sort these by separating the UMRs based on whether they contain any “Builder” or “Architect” tags in the sentence. As expected, for the Minecraft data alone, the baseline had very low scores, with the highest scores being 13.3, 30.6, and 30.2 for AnCast, SMATCH, and SMATCH++, respectively, as shown in Table 3. On the contrary, the data which does not include any Minecraft narration performs better: achieving highest scores of 62.2, 64.8 and 61.3 in the same order, as shown in Table 4. These are standard English sentences that do not include any dialogue tags or coordinate descriptions, and are thus more in line with the AMR training data supplied to the text-to-AMR models. Still, these standard English sentences differ from the first corpus. While UMR v1.0 is very conversational and fragmented, the UMR v2.0 non-Minecraft sentences are longer and more complex (see Section 3.3), thus accounting for the smaller difference we see between the UMR v1.0 results and v2.0 results.

Additionally, we use the neuro-symbolic approach by Post et al. (2024) to incorporate the split-role decisions of their model into the UMRs generated by the pipeline approach. The animacy parsing model is trained on five different folds of data, and so we make five different sets of UMRs with the roles predicted by the checkpoints from all folds. We obtain the scores from each of them and average them out to get the final score. However, the new scores lead to a negligible change; for

Model	AnCast	SMATCH	SMATCH++
amrlib	11.390	28.171	29.620
SPRING	12.568	29.322	29.974
BiBL	10.455	26.401	26.886
LeakDistill	11.850	29.864	28.940
AMRBART	13.329	30.600	30.164

Table 3: For English sentences *with Builder/Architect tags* from UMR v2.0, the performance of each model, measured by F-scores, after converting their generated AMRs to UMRs via the pipeline approach and comparing them against the gold UMRs.

Model	AnCast	SMATCH	SMATCH++
amrlib	57.775	62.248	60.351
SPRING	61.907	64.572	61.299
BiBL	62.210	64.843	61.185
LeakDistill	57.010	61.675	58.492
AMRBART	54.065	60.626	58.770

Table 4: Performance of each model, measured by F-scores, after converting their generated AMRs to UMRs via the pipeline approach and comparing them against the gold UMRs for English sentences *without Builder/Architect tags* from UMR v2.0.

example, BiBL achieves a 32.236 SMATCH score with the baseline pipeline and 32.250 after incorporating the neural conversion. As this addition has a minimal effect on overall performance, we exclude these results from further discussion.

4.2. Fine-tuning

In this section, we evaluate how well existing AMR parser architectures adapt to UMR parsing after being fine-tuned on UMR data. Specifically, we assess which AMR-to-text generation model performs best on the designated test split of all English data from UMR v2.0 when fine-tuned on English sentences and their respective UMR graphs, using graph similarity metrics as our primary evaluation criteria. The results in Table 5 demonstrate that BiBL achieves the strongest overall performance among the evaluated models, with scores of 84.348, 88.816, and 90.975 on AnCast, SMATCH, and SMATCH++, respectively.

To complement our experiments with fine-tuned AMR architectures, we also evaluate our UD-based parsing approach. Interestingly, the UD approach performs similarly to the fine-tuned AMR-to-text models in terms of scores, often matching and sometimes even exceeding the results of the UMR fine-tuned SPRING and LeakDistill models. This highlights its viability as a UD-driven method, particularly given that it is not explicitly designed for

UMR. However, a weakness of this approach is T5’s tendency to misalign, omit, or insert extra parentheses, resulting in structurally invalid UMR graphs. For instance, a common error of this approach is the omission of the final parenthesis in long UMR graphs. To address this, we develop a post-processing script to detect and correct parenthesis mismatches. It is also important to note that all the fine-tuning approaches outperform the baseline pipeline approach, which shows that both fine-tuning text-to-AMR parsing models and utilizing UD-driven methods show great promise as state-of-the-art UMR parsing approaches.

To better understand the strengths and weaknesses of these approaches, we examine two sets of generated UMR graphs from our fine-tuning and UD approaches, as well as the gold UMR for the same sentence.

Sentence: <Architect> oops sorry, I meant behind :)

Gold UMR:

```
(s34a / and
 :op1 (s34s / sorry-01
       :mod (s34o / oops
             :mode expressive))
 :op2 (s34m / mean-01
       :ARG0 (s34i / i)
       :ARG2 (s34b / behind-02
             :ARG3 (s34c / cartesian-framework-91
                   :FR (s34r / relative-to-builder))))))
```

Fine-tuned BiBL-generated UMR:

```
(z0 / and
 :op1 (z1 / sorry-01
       :mod (z2 / oops
             :mode expressive))
 :op2 (z3 / mean-01
       :ARG0 (z4 / i)
       :ARG2 (z5 / behind-02
             :ARG3 (z6 / cartesian-framework-91
                   :FR (z7 / relative-to-builder))))
 :op3 (z8 / emoticon
       :value " :)"))
```

UD Approach-generated UMR:

```
(s2a / and
 :op1 (s2s / sorry-01
       :ARG1 (s2o / oops))
 :op2 (s2m / mean-01
       :ARG0 (s2i / i)
       :ARG2 (s2b / behind-02
             :ARG2 s2i)))
```

Analysis of the generated graphs show that both BiBL and the UD approach accurately capture the root, which is and, as well as the main predicates sorry-01 and mean-01. BiBL nearly perfectly reproduces the gold graph, accurately representing the modifier oops with :mode expressive along with the relative structure (:FR relative-to-builder). In this example, BiBL achieves scores of 88.89, 91.89,

and 91.89 for AnCast, SMATCH, and SMATCH++ respectively. The sole deviation is BiBL’s addition of an emoticon node, which is included in the sentence and does not alter the core semantic interpretation of the graph. In contrast, the UD approach misaligns *oops* as the :ARG1 of *sorry-01* rather than the modifier. Moreover, it entirely excludes :mode *expressive*, showing that this approach is susceptible to excluding the nuances of UMR. This is further confirmed by the automatic metrics, where this approach achieves scores of 57.143, 73.333, and 73.33 for AnCast, SMATCH, and SMATCH++, respectively. From this analysis, we find that BiBL produces UMRs that are more closely aligned with the gold reference, whereas the UD approach captures only the core meaning and is prone to omitting finer details, a pattern also reflected in the quantitative results of this example and the overall results.

Sentence: Yes, even flowers have thorns.

Gold UMR:

```
(f / flower
  :mod (e / even)
  :ARG0-of (h / have-03
    :ARG1 (t / thorn)))
```

Fine-tuned BiBL-generated UMR:

```
(c / confirm-01
  :ARG1 (e / even-if
    :op1 (f / flower
      :ARG0-of (t / thorn-01))))
```

UD Approach-generated UMR:

```
(a / and
  :op1 (c / confirm-01)
  :op2 (f / flower
    :mod (e / even)
    :ARG1-of (h / have-03
      :ARG0 (t / thorn))))
```

In our second example, we observe a reversal of performance patterns. The UD approach more closely mirrors the gold UMR, accurately representing even as a modifier of *flower* via :mod, as well as capturing the *have-03* predicate with its correct :ARG0, *thorn*. However, this approach does predict the root to be *and* rather than *flower*. Furthermore, BiBL misidentifies the root as *confirm-01* and incorrectly frames *even* as a conjunction operator *even-if*, losing the modifier relationship entirely. It also fails to recover *have-03* and conflates *thorn* into *thorn-01*. This divergence from BiBL’s otherwise strong performance highlights a key limitation: sentences drawn from non-Minecraft contexts, such as general declarative statements, are underrepresented in the fine-tuning data, leading to degraded performance. Overall, we note that the imbalance between Minecraft dialogue and non-Minecraft data in the training corpus means

Model	AnCast	SMATCH	SMATCH++
amrlib	79.393	85.240	88.842
SPRING	72.553	78.990	81.710
BiBL	84.348	88.816	90.975
LeakDistill	78.108	81.656	81.465
AMRBART	81.700	86.239	88.696
T5: UDPipeline	72.845	80.624	82.853

Table 5: Overall performance of each AMR model, measured by F-scores, after fine-tuning on all English sentences from UMR v2.0 and comparing their predicted UMRs against the gold references

that even the fine-tuned BiBL model is not equally well-equipped to handle all sentence types, with performance on non-Minecraft inputs remaining a limitation of the current approach. Our models demonstrate promising initial performance on English, though our qualitative error analysis reveals avenues for future work.

5. Conclusion & Future Work

In this work, we take steps towards building sentence-level English text-to-UMR parsers by exploring two primary strategies: (1) fine-tuning existing text-to-AMR parsers on UMR data, and (2) converting UD trees into partial UMR graphs and training a T5 model to complete them. First, we find that the current state-of-the-art pipeline approach performs poorly on the UMR v2.0 splits due to the different nature of the data. However, as we evaluate the model on only the non-Minecraft portion of the split, we get more similar results to this approach when done on UMR v1.0. Furthermore, we find that fine-tuning existing text-to-AMR parsers yields improved performance over the baseline, and it is our most successful approach overall. Among these text-to-AMR parsers, BiBL consistently achieves the highest scores across AnCast, SMATCH, and SMATCH++, while AMRBART ranks second. Interestingly, our UD approach, which leverages Universal Dependencies to create partial UMRs and then converts those partial UMRs to complete UMRs, remains competitive. It often surpasses SPRING and rivals LeakDistill, suggesting that UD-driven methods, in combination with sequence-to-sequence models, provide a viable strategy for UMR parsing.

These findings demonstrate that while UMR parsing is a challenging task due to the complexity of representing fine-grained semantic, temporal, modal, and coreference relations across sentences, we are able to achieve state-of-the-art results by leveraging existing AMR and UD tools. More broadly, our work lays the foundation for the

development of UMR parsers as reliable parsing is essential for enabling UMR to be utilized in downstream applications such as machine translation, summarization, and information extraction, and to realize its promise as a cross-linguistic, meaning-based representation that extends beyond English. Focusing on English provides a practical starting point, as it offers high-quality annotated resources and AMR parsers that can be fine-tuned, allowing us to explore approaches to building models and establish strong baselines. These insights and methods can then guide the expansion of UMR parsing to low-resource languages, where annotated data and pre-trained parsers are scarce. By establishing these foundations in English, we provide a pathway to extend UMR’s utility to multilingual and low-resource settings.

Limitations

In order to establish strong foundational work on UMR parsing which leverages existing AMR technologies, we focus on parsing English, sentence-level UMR graphs. As models for UMR improve and datasets expand, additional work on document-level UMR parsing for many languages will be made possible, but will continue to be a challenge for low-resource and morphosyntactically complex languages. While we experiment with five AMR parsers, future work might explore other AMR parsers such as CLAP (Martinez Lorenzo and Navigli, 2024), StructBART (Zhou et al., 2021), or MBSE (Lee et al., 2022). Future work might also extend to automatically parsing the document-level components of the UMR graphs.

AnCast++ (Sun et al., 2025), is a recently introduced extension of AnCast designed to provide a unified evaluation of sentence-level graphs, modal and temporal dependencies, and coreference relations. We do not use AnCast++ in our evaluation as it was released contemporaneous to our work.

Acknowledgments

We thank anonymous reviewers and members of the Amherst College NLP lab for their feedback. This work is supported by the Amherst College HPC, which is funded by NSF Award 2117377.

6. Bibliographical References

2022. Uniform Meaning Representation (UMR) 0.9 specification. <https://github.com/umr4nlp/umr-guidelines/blob/master/guidelines.md>.
- Omri Abend and Ari Rappoport. 2013. [Universal Conceptual Cognitive Annotation \(UCCA\)](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 228–238, Sofia, Bulgaria. Association for Computational Linguistics.
- Xuefeng Bai, Yulong Chen, and Yue Zhang. 2022. [Graph pre-training for AMR parsing and generation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6001–6015, Dublin, Ireland. Association for Computational Linguistics.
- 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.
- 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*.
- Austin Blodgett and Nathan Schneider. 2021. [Probabilistic, structure-aware algorithms for improved variety, accuracy, and coverage of AMR alignments](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3310–3321, Online. Association for Computational Linguistics.
- Julia Bonn, Claire Bonial, Matt Buchholz, Hsiao-Jung Cheng, Alvin Chen, Ching-wen Chen, Andrew Cowell, William Croft, Lukas Denk, Ahmed Elsayed, Eva Fučíková, Federica Gamba, Carlos Gomez, Jan Hajič, Eva Hajičová, Jiří Havelka, Loden Havenmeier, Ath Kilgore, Veronika Kolářová, Lucie Kučová, Kenneth Lai, Bin Li, Jingyi Li, Markéta Lopatková, Marie MacGregor, Marie Mikulová, Jiří Mírovský, Anna Nedoluzhko, Skatje Myers, Michal Novák, Tim O’Gorman, Petr Pajas, Alexis Palmer, Martha Palmer, Jarmila Panevová, Benét Post, James Pustejovsky, Petr Sgall, Jialin Song, Li Song, Magda Ševčíková, Jan Štěpánek, Zdeňka Urešová, Haibo Sun, Yao Sun, Rosa Vallejos Yopán, Jens Van Gysel, Meagan Vigus, Kristin Wright-Bettner, Jiawei Wu, Nianwen Xue, Dan Xing, Keer Xu, Zhixing Xu, Liulu Yue, Daniel

- Zeman, Jin Zhao, Šárka Zikánová, and Zdeněk Žabokrtský. 2025. [Uniform meaning representation 2.0](#). LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL).
- Julia Bonn, Matthew J. Buchholz, Jayeol Chun, Andrew Cowell, William Croft, Lukas Denk, Sijia Ge, Jan Hajič, Kenneth Lai, James H. Martin, Skatje Myers, Alexis Palmer, Martha Palmer, Claire Benet Post, James Pustejovsky, Kristine Stenzel, Haibo Sun, Zdeňka Urešová, Rosa Vallejos, Jens E. L. Van Gysel, Meagan Vigus, Nianwen Xue, and Jin Zhao. 2024. [Building a broad infrastructure for uniform meaning representations](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 2537–2547, Torino, Italia. ELRA and ICCL.
- Julia Bonn, Skatje Myers, Jens E. L. Van Gysel, Lukas Denk, Meagan Vigus, Jin Zhao, Andrew Cowell, William Croft, Jan Hajič, James H. Martin, Alexis Palmer, Martha Palmer, James Pustejovsky, Zdenka Urešová, Rosa Vallejos, and Nianwen Xue. 2023. [Mapping AMR to UMR: Resources for adapting existing corpora for cross-lingual compatibility](#). In *Proceedings of the 21st International Workshop on Treebanks and Linguistic Theories (TLT, GURT/SyntaxFest 2023)*, pages 74–95, Washington, D.C. Association for Computational Linguistics.
- Matthew J. Buchholz, Julia Bonn, Claire Benet Post, Andrew Cowell, and Alexis Palmer. 2024. [Bootstrapping UMR annotations for Arapaho from language documentation resources](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 2447–2457, Torino, Italia. ELRA and ICCL.
- 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.
- Ziming Cheng, Zuchao Li, and Hai Zhao. 2022. [BiBL: AMR parsing and generation with bidirectional Bayesian learning](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 5461–5475, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Jayeol Chun and Nianwen Xue. 2024. [Uniform meaning representation parsing as a pipelined approach](#). In *Proceedings of TextGraphs-17: Graph-based Methods for Natural Language Processing*, pages 40–52, Bangkok, Thailand. Association for Computational Linguistics.
- Marie-Catherine de Marneffe, Christopher D. Manning, Joakim Nivre, and Daniel Zeman. 2021. [Universal Dependencies](#). *Computational Linguistics*, 47(2):255–308.
- Kilian Evang. 2019. [Transition-based DRS parsing using stack-LSTMs](#). In *Proceedings of the IWCS Shared Task on Semantic Parsing*, Gothenburg, Sweden. Association for Computational Linguistics.
- Federica Gamba, Alexis Palmer, and Daniel Zeman. 2025. [Bootstrapping UMRs from Universal Dependencies for scalable multilingual annotation](#). In *Proceedings of the 19th Linguistic Annotation Workshop (LAW-XIX-2025)*, pages 126–136, Vienna, Austria. Association for Computational Linguistics.
- Daniel Hershcovich, Omri Abend, and Ari Rapoport. 2017. [A transition-based directed acyclic graph parser for UCCA](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1127–1138, Vancouver, Canada. Association for Computational Linguistics.
- Hans Kamp. 1984. *A Theory of Truth and Semantic Representation*, pages 1–42. De Gruyter Mouton, Berlin, Boston.
- Robert T. Kasper. 1989. [A flexible interface for linking applications to Penman’s sentence generator](#). In *Speech and Natural Language: Proceedings of a Workshop Held at Philadelphia, Pennsylvania, February 21-23, 1989*.
- Kevin Knight, Bianca Badarau, Laura Banarescu, Claire Bonial, Madalina Bardocz, Kira Griffitt, Ulf Hermjakob, Daniel Marcu, Martha Palmer, Tim O’Gorman, and Nathan Schneider. 2017. [Abstract Meaning Representation \(AMR\) Annotation Release 2.0](#). Technical Report LDC2017T10, Linguistic Data Consortium, Philadelphia, PA.
- Kevin Knight, Bianca Badarau, Laura Baranescu, Claire Bonial, Madalina Bardocz, Kira Griffitt, Ulf Hermjakob, Daniel Marcu, Martha Palmer, Tim O’Gorman, et al. 2020. [Abstract Meaning Representation \(AMR\) Annotation Release 3.0](#). Technical Report LDC2020T02, Linguistic Data Consortium, Philadelphia, PA.

- Young-Suk Lee, Ramón Astudillo, Hoang Thanh Lam, Tahira Naseem, Radu Florian, and Salim Roukos. 2022. [Maximum Bayes Smatch ensemble distillation for AMR parsing](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5379–5392, Seattle, United States. Association for Computational Linguistics.
- Emma Markle, Reihaneh Iranmanesh, and Shira Wein. 2025. [Generating text from uniform meaning representation](#). In *Proceedings of the 14th International Joint Conference on Natural Language Processing and the 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics*, pages 259–271, Mumbai, India. The Asian Federation of Natural Language Processing and The Association for Computational Linguistics.
- Abelardo Carlos Martinez Lorenzo and Roberto Navigli. 2024. [Efficient AMR parsing with CLAP: Compact linearization with an adaptable parser](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 5578–5584, Torino, Italia. ELRA and ICCL.
- Stephan Oepen, Marco Kuhlmann, Yusuke Miyao, Daniel Zeman, Dan Flickinger, Jan Hajič, Angelina Ivanova, and Yi Zhang. 2014. [SemEval 2014 task 8: Broad-coverage semantic dependency parsing](#). In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 63–72, Dublin, Ireland. Association for Computational Linguistics.
- Juri Opitz. 2023. [SMATCH++: Standardized and extended evaluation of semantic graphs](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1595–1607, Dubrovnik, Croatia. Association for Computational Linguistics.
- Claire Benet Post, Marie C. McGregor, Maria Leonor Pacheco, and Alexis Palmer. 2024. [Accelerating UMR adoption: Neuro-symbolic conversion from AMR-to-UMR with low supervision](#). In *Proceedings of the Fifth International Workshop on Designing Meaning Representations @ LREC-COLING 2024*, pages 140–150, Torino, Italia. ELRA and ICCL.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. [Stanza: A Python natural language processing toolkit for many human languages](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2023. [Exploring the limits of transfer learning with a unified text-to-text transformer](#).
- Zacchary Sadeddine, Juri Opitz, and Fabian Suchanek. 2024. [A survey of meaning representations – from theory to practical utility](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 2877–2892, Mexico City, Mexico. Association for Computational Linguistics.
- David Samuel and Milan Straka. 2020. [ÚFAL at MRP 2020: Permutation-invariant semantic parsing in PERIN](#). In *Proceedings of the CoNLL 2020 Shared Task: Cross-Framework Meaning Representation Parsing*, pages 53–64, Online. Association for Computational Linguistics.
- Haibo Sun, Jayeol Chun, and Nianwen Xue. 2025. [AnCast++: Document-level evaluation of graph-based meaning representations](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 19642–19654, Vienna, Austria. Association for Computational Linguistics.
- Haibo Sun and Nianwen Xue. 2024. [Anchor and broadcast: An efficient concept alignment approach for evaluation of semantic graphs](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 1052–1062, Torino, Italia. ELRA and ICCL.
- Jens E. L. Van Gysel, Meagan Vigus, Pavlina Kalm, Sook-kyung Lee, Michael Regan, and William Croft. 2019. [Cross-linguistic semantic annotation: Reconciling the language-specific and the universal](#). In *Proceedings of the First International Workshop on Designing Meaning Representations*, pages 1–14, Florence, Italy. Association for Computational Linguistics.
- Jens EL Van Gysel, Meagan Vigus, Jayeol Chun, Kenneth Lai, Sarah Moeller, Jiarui Yao, Tim O’Gorman, Andrew Cowell, William Croft, Chu-Ren Huang, et al. 2021. Designing a uniform meaning representation for natural language processing. *KI-Künstliche Intelligenz*, 35(3):343–360.
- Rik van Noord, Antonio Toral, and Johan Bos. 2020. [Character-level representations improve DRS-based semantic parsing even in the age of BERT](#).

In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4587–4603, Online. Association for Computational Linguistics.

Pavlo Vasylenko, Pere Lluís Hugué Cabot, Abelardo Carlos Martínez Lorenzo, and Roberto Navigli. 2023. [Incorporating graph information in transformer-based AMR parsing](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1995–2011, Toronto, Canada. Association for Computational Linguistics.

Meagan Vigus, Jens E. L. Van Gysel, Tim O’Gorman, Andrew Cowell, Rosa Vallejos, and William Croft. 2020. [Cross-lingual annotation: a road map for low- and no-resource languages](#). In *Proceedings of the Second International Workshop on Designing Meaning Representations*, pages 30–40, Barcelona Spain (online). Association for Computational Linguistics.

Xinyu Wang, Yong Jiang, Nguyen Bach, Tao Wang, Zhongqiang Huang, Fei Huang, and Kewei Tu. 2021. [Automated concatenation of embeddings for structured prediction](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2643–2660, Online. Association for Computational Linguistics.

Shira Wein. 2025. [Can uniform meaning representation help GPT-4 translate from indigenous languages?](#) In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 278–285, Vienna, Austria. Association for Computational Linguistics.

Shira Wein and Julia Bonn. 2023. [Comparing UMR and cross-lingual adaptations of AMR](#). In *Proceedings of the Fourth International Workshop on Designing Meaning Representations*, pages 23–33, Nancy, France. Association for Computational Linguistics.

Shira Wein and Juri Opitz. 2024. [A survey of AMR applications](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6856–6875, Miami, Florida, USA. Association for Computational Linguistics.

Shira Wein and Nathan Schneider. 2024. [Assessing the cross-linguistic utility of Abstract Meaning Representation](#). *Computational Linguistics*, 50(2):419–473.

Shira B. Wein. 2024. *Development and Evaluation of Cross-Lingual Abstract Meaning Representa-*

tion. Ph.D. thesis, Georgetown University. Copyright - Database copyright ProQuest LLC; ProQuest does not claim copyright in the individual underlying works; Last updated - 2026-01-16.

Jiawei Zhou, Tahira Naseem, Ramón Fernández Astudillo, Young-Suk Lee, Radu Florian, and Salim Roukos. 2021. [Structure-aware fine-tuning of sequence-to-sequence transformers for transition-based AMR parsing](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6279–6290, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

7. Language Resource References